

SANDIA REPORT

SAND2021-16115

Printed December 23, 2021



Sandia
National
Laboratories

Sierra/SD – User's Manual – 5.4

Sierra Structural Dynamics Development Team

Latest Software Release:
5.2-Release 2021-09-15

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185
Livermore, California 94550

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology & Engineering Solutions of Sandia, LLC.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: <http://www.osti.gov/scitech>

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: <https://classic.ntis.gov/help/order-methods>



ABSTRACT

Sierra/SD provides a massively parallel implementation of structural dynamics finite element analysis, required for high-fidelity, validated models used in modal, vibration, static and shock analysis of weapons systems. This document provides a user's guide to the input for **Sierra/SD**. Details of input specifications for the different solution types, output options, element types and parameters are included. The appendices contain detailed examples, and instructions for running the software on parallel platforms.

CONTENTS

0.1.	Release 5.4	3
0.2.	Release 5.2	5
0.3.	Release 5.00	7
0.4.	Release 4.58	10
1.	HOW TO RUN Sierra/SD	13
1.1.	Accessing Sierra/SD	13
1.2.	Modules and Executables	13
1.3.	The Sierra/SD salinas Executable	14
1.4.	MPI Parallel Execution	15
1.4.1.	Number of MPI Processes Needed	16
1.4.2.	Mesh Decomposition	17
1.4.3.	Running the Sierra/SD Executable in Parallel	17
1.4.4.	Post Processing in Parallel	17
1.5.	File system concerns	18
1.6.	Workflow Examples	18
1.7.	Thread Parallelism	20
1.8.	Troubleshooting	20
1.8.1.	Stand-Alone Tools	20
1.8.2.	Using Sierra/SD To Troubleshoot	21
1.8.3.	Modal Analysis	22
1.8.4.	Evaluating Memory Use	22
1.8.5.	Identifying Problematic Subdomains	22
1.8.6.	Limitations of SD Finite Elements	23
1.8.7.	Problematic Elements and Connectivity	23
1.9.	Over-determined Constraints and Loss of Rigid Body Modes	25
2.	GENERAL COMMANDS	26
2.1.	Input deck format	26
2.2.	Input mesh geometry file	29
2.2.1.	Geometry_file	29
2.2.2.	Exodus Database Naming Conventions	30
2.2.3.	ASSEMBLY section	31
2.2.4.	Exodus Naming Limitations	31
2.2.5.	Additional Comments About Output	32
2.3.	Parameters	33
2.4.	Solution Options	44
2.4.1.	Flush	45
2.4.2.	Restart – option	45

	2.4.3.	Solver	50
	2.4.4.	Lumped – option	51
	2.4.5.	Constraintmethod – option	52
	2.4.6.	Scattering – option	52
2.5.	GDSW		53
	2.5.1.	Options	59
	2.5.2.	Diagnostics	60
	2.5.3.	Troubleshooting	62
	2.5.4.	Mathematical Conditioning Issues	66
	2.5.5.	Frequency Response Functions	67
	2.5.6.	Parameters	67
2.6.	Sensitivity		71
	2.6.1.	Attune	72
	2.6.2.	Sensitivity Output	72
2.7.	Coordinate		76
2.8.	Function		84
	2.8.1.	Function Offset/Shifts	85
	2.8.2.	Linear Functions	85
	2.8.3.	Sierra SM Piecewise Linear Functions	87
	2.8.4.	Functions using Tables	88
	2.8.5.	Polynomials	88
	2.8.6.	LogLog Functions	89
	2.8.7.	SamplingRandom	89
	2.8.8.	RandomLib Functions	90
	2.8.9.	Analytic Functions	92
	2.8.10.	Plane Wave (Time Domain)	98
	2.8.11.	Plane Wave (Frequency Domain)	100
	2.8.12.	Planar Step Wave	101
	2.8.13.	Spherically Spreading Wave	102
	2.8.14.	Undex Structural Acoustic Loads	103
	2.8.15.	Fluid Structure Interaction	105
	2.8.16.	Blending	106
	2.8.17.	Matrix-function	109
	2.8.18.	Alternate Table Interface	110
	2.8.19.	Table	111
2.9.	Multipoint constraints		113
3.	SOLUTION CASES		116
	3.1.	Defining Solution Cases	118
	3.2.	Multicase Solutions	119
	3.2.1.	Multicase Options	119
	3.2.2.	Multicase Example	121
	3.3.	Checkout Solution Case	122
	3.4.	CJdamp Solution Case	122
	3.5.	Craig-Bampton reduction Solution Case	123
	3.5.1.	CBModel	126

3.5.2.	Sensitivity Analysis for Craig-Bampton models	131
3.6.	DDAM Solution Case	133
3.7.	DirectFRF Solution Case	136
3.7.1.	Padé Expansion	137
3.8.	DirectFRF-Inverse Solution Case	138
3.8.1.	Load Identification	139
3.8.2.	Material Identification	140
3.8.3.	Multi-Experiment Material Identification	141
3.9.	Dump Solution Case	141
3.10.	Eigen Solution Case	142
3.10.1.	Option nmodes	142
3.10.2.	Solving Singular Systems with Shifts	143
3.11.	Eigen-Inverse Solution Case	146
3.11.1.	Eigenvalue Material Identification	147
3.11.2.	Eigenvector Material Identification	148
3.12.	AEigen Solution Case	151
3.13.	LargestEv Solution Case	153
3.14.	Fatigue Solution Case	153
3.15.	Buckling Solution Case	156
3.16.	ModalFilter Solution Case	157
3.17.	Modal Participation Factor Solution Case	159
3.18.	ModalFrf Solution Case	163
3.19.	ModalFRF-Inverse Solution Case	167
3.19.1.	Load Identification	168
3.19.2.	PSD Load Identification	170
3.20.	ModalRanVib Solution Case	172
3.21.	ModalShock Solution Case	177
3.22.	ModalTransient Solution Case	178
3.23.	ModalTransient-Inverse Solution Case	181
3.24.	QEVF Solution Case	183
3.24.1.	Anasazi	185
3.24.2.	Damped Eigenvalue Problems	186
3.24.3.	SA_eigen	188
3.24.4.	Projection_eigen	191
3.25.	NLStatics Solution Case	192
3.26.	NLTransient Solution Case	194
3.27.	ReceiveSierraData Solution Case	195
3.28.	Statics Solution Case	201
3.29.	Superposition Solution Case	202
3.30.	Tangent Solution Case	203
3.31.	TranShock Solution Case	204
3.32.	Transient Solution Case	206
3.32.1.	nUpdateConstraints Option	208
3.33.	Transient-Inverse Solution Case	210
3.33.1.	Load Identification	211

3.33.2.	Material Identification	212
3.34.	TSRpreload Solution Case	213
3.35.	Residual Vectors Solution Case	215
3.36.	GeometricRigidBodyModes Solution Case	217
3.37.	Waterline Solution Case	219
3.38.	Gap Removal Solution Case	223
3.39.	Inverse Problems	224
3.39.1.	Optimization	226
3.39.2.	Inverse-Problem	230
3.39.3.	Inverse Data Files	231
3.39.4.	Block section for Material Identification	238
3.39.5.	Material section for Material Identification	239
3.39.6.	Loads section for Load Identification	246
3.39.7.	Limitations for Inverse Load Problems	248
3.39.8.	ROL Output for Inverse Problems	248
4.	MATERIALS	249
4.1.	Elastic	249
4.1.1.	Isotropic	250
4.1.2.	Orthotropic	250
4.1.3.	Anisotropic	251
4.1.4.	Lamé Material	252
4.2.	Acoustic	253
4.3.	Linear Viscoelastic	254
4.3.1.	Limitations of Viscoelastic Use.	256
4.4.	Properties	257
4.4.1.	Density	258
4.4.2.	High Cycle Fatigue	258
4.4.3.	S-N curve Definitions	259
4.4.4.	S-N Curve Units	259
4.4.5.	Typical Material Data for Fatigue	260
4.4.6.	Temperature dependence	262
4.4.7.	Spatially Variant Material Properties	263
4.4.8.	Specific Heat	264
4.4.9.	Frequency dependence	265
4.5.	Piezoelectric Material	266
4.6.	Dielectric Material	267
4.7.	Block	268
4.7.1.	Block Parameters	268
4.7.2.	General Block Parameters	271
4.8.	Damping	275
4.8.1.	Nonlinear transient solutions with damping	278
4.8.2.	Nonlinear Distributed Damping	278
5.	ELEMENTS	281
5.1.	Hex8	281
5.2.	Hex20	282

5.3.	Wedge6	282
5.4.	Wedge15	283
5.5.	Tet4	283
5.6.	Tet10	283
5.7.	Two-Dimensional Shell and Membrane Elements	283
5.7.1.	QuadT, Quad8T, and Tria6	284
5.7.2.	QuadM	286
5.7.3.	Nquad/Ntria	287
5.7.4.	TriaShell	289
5.7.5.	Tria3	290
5.7.6.	Stiffness Scaling	291
5.7.7.	Shell Coordinate Systems	291
5.7.8.	Layered Shells	293
5.7.9.	Offset Shells	295
5.7.10.	Spatially Dependent Properties via Exodus Attributes	296
5.8.	HexShell	297
5.9.	Beam2	301
5.10.	Nbeam	305
5.11.	TiBeam	308
5.12.	Truss	309
5.13.	Ftruss	309
5.14.	ConMass	310
5.15.	Spring	312
5.15.1.	Spring Parameter Values	312
5.16.	RSpring	313
5.17.	Spring3 - nonlinear cubic spring	313
5.18.	Dashpot	314
5.19.	SpringDashpot	315
5.20.	Hys	317
5.21.	Joint2G	319
5.21.1.	Specification	319
5.21.2.	Constitutive Behavior	320
5.22.	Line Weld	328
5.23.	Gap element	331
5.24.	Gap2D	334
5.25.	GasDmp	335
5.26.	Nmount	336
5.27.	Rrod	338
5.28.	RBar	339
5.28.1.	Interaction of RBars	340
5.29.	RBE2	340
5.30.	RBE3	340
5.31.	Superelement	342
5.32.	Dead	348
5.33.	Rigidset	348

5.34.	Rrodset	350
5.35.	Tied Joint	351
5.35.1.	Input Specification	351
6.	BOUNDARY CONDITIONS AND INITIAL CONDITIONS	358
6.1.	Boundary conditions	358
6.1.1.	Prescribed Displacements and Pressures	361
6.1.2.	Prescribed Voltage	362
6.1.3.	Prescribed Accelerations	362
6.1.4.	Prescribed Frequency-Varying Displacements	364
6.1.5.	Node_List_File	364
6.1.6.	Nonreflecting Boundaries	365
6.1.7.	Impedance Boundary Conditions	366
6.1.8.	Slosh	366
6.1.9.	Infinite Elements	367
6.1.10.	Perfectly Matched Layers	369
6.1.11.	Periodic Boundary Conditions	371
6.1.12.	Usage Guidelines	372
6.2.	Exodus Mesh Boundary Condition Input	373
6.2.1.	SpatialBC Functions	374
6.2.2.	Input an Acoustic Point Source from a Volume	375
6.2.3.	Input an Acoustic Point Source from a Node Set	376
6.2.4.	ReadSurface functions	377
6.2.5.	ExodusRead functions	378
6.2.6.	In Core Transfer Functions	378
6.3.	Loads	380
6.3.1.	Load	381
6.3.2.	Scale Factors for the Load	381
6.3.3.	Sideset Loading	382
6.3.4.	Spatial Variation	384
6.3.5.	Required Section	384
6.3.6.	Follower Stiffness	384
6.3.7.	Acoustic Loads	385
6.3.8.	Thermal Loads	388
6.3.9.	Energy Deposition Input and Loads	392
6.3.10.	Consistent Loads	392
6.3.11.	Pressure_Z	392
6.3.12.	Surface Charge	393
6.3.13.	Static Loads	393
6.3.14.	Time Varying Loads	394
6.3.15.	Random Pressure Loads	394
6.3.16.	Frequency Dependent Loads	398
6.3.17.	Modal Force Loading	398
6.3.18.	Rotational frames	399
6.3.19.	Rigid Body Filter for Input	403
6.3.20.	RanLoads	404

6.4.	Initial Conditions	406
6.4.1.	Reading Initial Conditions from the Mesh File	407
6.4.2.	Setting Initial Conditions in the Input Deck	407
6.5.	Use cases for initial acceleration	408
7.	OUTPUTS	410
7.1.	Exodus	410
7.1.1.	Surface Projection of Element Variables	412
7.1.2.	Database Name	413
7.1.3.	Properties	413
7.1.4.	Maa	414
7.1.5.	Material	414
7.1.6.	material_direction	414
7.1.7.	Kaa	414
7.1.8.	Faa	415
7.1.9.	MLumped	415
7.1.10.	MPhi	415
7.1.11.	ElemEigChecks	416
7.1.12.	ElemQualchecks	416
7.1.13.	Displacement	417
7.1.14.	Velocity	417
7.1.15.	Acceleration	417
7.1.16.	Strain	418
7.1.17.	Strain = GP	419
7.1.18.	Stress	419
7.1.19.	Principal Stresses	420
7.1.20.	Von Mises stress	420
7.1.21.	Signed Von Mises Stress	421
7.1.22.	Rainflow Cycle Counting	421
7.1.23.	Fatigue Damage	421
7.1.24.	Stress = GP	422
7.1.25.	Vrms	424
7.1.26.	Rotational_displacement	424
7.1.27.	Rotational_acceleration	424
7.1.28.	Energy	424
7.1.29.	GEnergies	424
7.1.30.	Globals	425
7.1.31.	Block_Energies	425
7.1.32.	Mesh_Error	426
7.1.33.	MFile	426
7.1.34.	Force	428
7.1.35.	Constraint force	428
7.1.36.	Reaction Force	428
7.1.37.	Right-hand side	428
7.1.38.	EForce	429
7.1.39.	Line_Weld	431

7.1.40.	Relative_Displacement	431
7.1.41.	Residuals	432
7.1.42.	TIndex	433
7.1.43.	EOrient	433
7.1.44.	Pressure	434
7.1.45.	NPressure	434
7.1.46.	APressure	435
7.1.47.	APartVel	435
7.1.48.	Constraint_Info	435
7.1.49.	Statistics	436
7.1.50.	KDiag	437
7.1.51.	ADiag	438
7.1.52.	ddamout	438
7.1.53.	Temperature	439
7.1.54.	Voltage and Charge	439
7.2.	User Output	440
7.2.1.	Block Statistics	440
7.2.2.	Closest Distance	441
7.2.3.	Derived Variables	442
7.3.	History	443
7.3.1.	Global History Output Near a Location	445
7.4.	Frequency	446
7.5.	Linesample	448
7.6.	Stresses and Strains	449
7.6.1.	Stress/Strain Truth Table	449
7.6.2.	Solid Elements	451
7.6.3.	Shell Elements	451
7.6.4.	Beam Elements	452
7.7.	Echo	453
7.7.1.	Mass Properties	455
7.7.2.	Multipoint constraints	455
7.7.3.	ModalVars	456
7.7.4.	Subdomains	456
7.7.5.	Memusage	457
8.	CONTACT	458
8.1.	Tied Surfaces	458
8.1.1.	Contact Normal Vectors	459
8.1.2.	Mortar Methods	461
8.1.3.	Node to Face	461
8.2.	Contact Definition	464
8.2.1.	Defining Contact Surfaces	465
8.2.2.	Setting up Contact Interactions	467
8.2.3.	Gap removal and element quality	469
8.2.4.	Examples	470
8.2.5.	Notes and Usage Guidelines	471

8.2.6.	Differences Between SM and SD Defaults	472
8.3.	Lofted Surfaces and Gap Removal	472
8.3.1.	Example	473
8.3.2.	Projection Approach	473
8.4.	Moving MPCs	475
9.	EXAMPLE INPUT FILES	477
9.1.	Eigenvalue problem	477
9.2.	Anisotropic Material	478
9.3.	Multiple materials	479
9.4.	Modaltransient	481
9.5.	ModalFrf	483
9.6.	Directfrf	484
9.7.	Statics	486
Bibliography		487
Index		493

LIST OF FIGURES

Figure 1-1.	Single <i>Spring</i> element.	24
Figure 1-2.	Truss Decomposition Issues.	24
Figure 2-3.	Example MFile Format Results.	42
Figure 2-4.	Notes on Eigen Restart.	49
Figure 2-5.	Eigen Sensitivity Example Data	73
Figure 2-6.	Semi-Analytic Methods for Sensitivity Analysis.	75
Figure 2-7.	Coordinate system definition vectors: note that XZ POINT determines the $\tilde{X}$ axis, but need not lie along it, which is the case depicted in this figure.	79
Figure 2-8.	Coordinate System Examples	80
Figure 2-9.	Conical Coordinate System Definition at $\tilde{X}\tilde{Z}$ plane.	80
Figure 2-10.	Ellipsoidal Coordinate System Using axis_stretching=(2.0, 1.1, 1.0). $\hat{r}$ is red, $\hat{s}$ is green, and $\hat{t}$ is blue.	81
Figure 2-11.	Coordinate system behavior near the $\tilde{Z}$ axis. $\hat{r}$ is red, $\hat{s}$ is green, and $\hat{t}$ is blue.	82
Figure 2-12.	Full model with different coordinate system uses	84
Figure 2-13.	Linear function "ignored_point".	86
Figure 2-14.	Linear function "extrap_fun".	86
Figure 2-15.	Linear function #5. "multiple_fun".	87
Figure 2-16.	RandomLib Temporal Interpolation.	92
Figure 2-17.	Spherical Wave Geometry.	103
Figure 2-18.	Fluid-Structure Interaction (FSI) Infrastructure.	106
Figure 2-19.	Illustration of first crossing blended function.	107
Figure 2-20.	Illustration of Nth crossing blended functions.	108
Figure 3-21.	Craig-Bampton Reduction	130
Figure 3-22.	Superposition Data Flow Diagram.	202
Figure 3-23.	Waterline Coordinate Definition	221
Figure 3-24.	Net Force vs depth for a Rigid Body	222
Figure 3-25.	Sample Data Truth Table Input for Acoustic Problem	233
Figure 3-26.	Example Data Truth Table for Structural-only problem	233
Figure 3-27.	Example data truth table for Structural Acoustics	233
Figure 3-28.	Sample Real Data File Input for an acoustics-only problem	234
Figure 3-29.	Sample Real Data File Input for a structural-only problem	235
Figure 3-30.	Sample Real Data File Input for a data type moduli problem	236
Figure 3-31.	Sample Transient Data File Input for a structural-only problem	237
Figure 3-32.	Example of <i>ROL_Messages.txt</i> file for Inverse Problem Solution	249
Figure 4-33.	S-N Curve for Steel Sheet	261
Figure 5-34.	QuadT Element.	284
Figure 5-35.	Quad8T Element.	285

Figure 5-36. Tria6 Element.	285
Figure 5-37. Function for nquad_eps_max.	289
Figure 5-38. Projection of global coordinate system to shell.	292
Figure 5-39. Rotation of global coordinate system about axis prior to projection to shell.	292
Figure 5-40. Rotation of local system about normal after to projection to shell.	293
Figure 5-41. Stacking arrangement for a multi layer shell element.	294
Figure 5-42. Offset examples for shell of thickness 0.1.	296
Figure 5-43. Beam Orientation and Local Coordinate System.	303
Figure 5-44. Beam Offset and Local Coordinate System.	304
Figure 5-45. Nbeam Orientation, Offset and Local Coordinate System.	306
Figure 5-46. Hys element parameters.	318
Figure 5-47. Iwan Constitutive Model.	322
Figure 5-48. Hysteresis Microslip Variation with β	323
Figure 5-49. Hysteresis Macroslip Variation with β	324
Figure 5-50. Reduced Iwan Load Displacement Curve.	325
Figure 5-51. Eplas Model.	326
Figure 5-52. Line weld definitions for attaching the purple and cyan shells. The line weld follows the red beam, which is contiguously meshed with the purple shell.	328
Figure 5-53. Line weld Joint2G connections. The Joint2G elements are initially zero-length, and the green and red nodes coincident, but they are depicted with a finite deformation to delineate the individual components.	329
Figure 5-54. Gap element Force-Deflection Curve.	332
Figure 5-55. Mass bouncing off a Gap.	333
Figure 5-56. Gap2D force diagram.	335
Figure 5-57. Nmount Orientation.	338
Figure 5-58. Rigidset/TiedJoint Centernode Connection.	349
Figure 5-59. Rrodset Constraints.	351
Figure 5-60. Tied Joint Geometry.	351
Figure 5-61. Tied Joint Surface Normal Definition.	354
Figure 5-62. Construction of tied joint with side=average.	355
Figure 5-63. Construction of tied joint with side=rigid.	355
Figure 5-64. Construction of tied joint with side=rrod.	355
Figure 6-65. Example of Interpolation of Exodus Data to Analysis Steps.	374
Figure 6-66. Coordinate Frame Projection for Traction.	384
Figure 6-67. RandomPressure Loading Approximations.	395
Figure 7-68. Convergence of maximum stress at element centroids and surfaces.	412
Figure 7-69. Example KDiag output.	438
Figure 7-70. Tria3 Stress Recovery.	452
Figure 8-71. Shell Normal in Contact or Tied Interactions.	460
Figure 8-72. Search Tolerance definition.	462
Figure 8-73. Normal Definitions on Faceted Geometry.	462
Figure 8-74. Smoothing Parameters for Surface Normal Vectors.	463
Figure 8-75. Lofted Constraint Example.	473

LIST OF TABLES

Table 1-1.	Command Line Options.	15
Table 1-2.	Determining Number Of Processors Needed.	16
Table 1-3.	Elements using other elements condition number.	23
Table 2-4.	Comment String Options.	27
Table 2-5.	Available keywords in the Parameters section.	34
Table 2-6.	Available keywords in the Parameters section related to code coupling and handoff.	35
Table 2-7.	Developer keywords in the Parameters section.	35
Table 2-8.	Some useful combinations of units.	35
Table 2-9.	Eigenvector Normalization Methods.	41
Table 2-10.	Sierra/SD Solution Options.	45
Table 2-11.	Supported restart capabilities for transient integrators in Sierra/SD	49
Table 2-12.	Restart File Format and Names.	50
Table 2-13.	GDSW Section Options. (Basic)	54
Table 2-14.	GDSW Section Options 1. (Advanced)	55
Table 2-15.	GDSW Section Options 2. (Advanced)	56
Table 2-16.	GDSW Section Options. (Supplemental Output)	62
Table 2-17.	GDSW Section Options (Advanced).	68
Table 2-18.	GDSW Section Print Options.	69
Table 2-19.	GDSW Section Options (Helmholtz).	70
Table 2-20.	GDSW Section Options (Solvers).	70
Table 2-21.	Sensitivity Analysis Keywords.	71
Table 2-22.	Sensitivity Analysis Solution Type Availability.	75
Table 2-23.	Coordinate Names for history files.	83
Table 2-24.	SamplingRandom function parameters.	90
Table 2-25.	RandomLib function parameters.	91
Table 2-26.	Predefined Analytic Input Variables.	98
Table 2-27.	Planar Step Wave Parameters.	102
Table 2-28.	Spherical Wave Parameters.	102
Table 2-29.	Undex Load Parameters	104
Table 2-30.	Blended Function Parameters.	106
Table 2-31.	Table Section Options.	112
Table 2-32.	General MPC commands.	114
Table 2-33.	MPC Equation lines.	114
Table 3-34.	Eigen Solution Types.	116
Table 3-35.	Modal Solution Types.	117
Table 3-36.	Direct Solution Types.	117
Table 3-37.	Preprocessing and Postprocessing Solution Types.	118

Table 3-38.	Multicase Options.....	120
Table 3-39.	Checkout Solution Case Parameters.	122
Table 3-40.	CJdamp Solution Case Parameters.	122
Table 3-41.	CBR Solution Case Parameters.	123
Table 3-42.	CBModel Parameters.....	127
Table 3-43.	Data output for Craig-Bampton Reduction.	129
Table 3-44.	DDAM Solution Case Parameters.	133
Table 3-45.	DirectFRF Solution Case Parameters.	136
Table 3-46.	DirectFRF-Inverse Solution Case Parameters.	138
Table 3-47.	Dump Solution Case Parameters.	141
Table 3-48.	Eigen Solution Case Parameters.	142
Table 3-49.	Eigen-Inverse Solution Case Parameters.....	146
Table 3-50.	AEigen Solution Case Parameters.	151
Table 3-51.	AEigen Verbosity Table.....	152
Table 3-52.	LargestEv Solution Case Parameters.....	153
Table 3-53.	Fatigue Solution Case Parameters.....	153
Table 3-54.	Buckling Solution Case Parameters.....	156
Table 3-55.	ModalFilter Solution Case Parameters.	157
Table 3-56.	Modal Filter Keywords.	159
Table 3-57.	MPF Solution Case Parameters.	159
Table 3-58.	MPF Summary data	161
Table 3-59.	ModalFrf Solution Case Parameters.	163
Table 3-60.	ModalFRF-Inverse Solution Case Parameters.	167
Table 3-61.	ModalRanVib Solution Case Parameters.	172
Table 3-62.	ModalRanVib Output to Exodus File	175
Table 3-63.	ModalShock Solution Case Parameters.....	177
Table 3-64.	ModalTransient Solution Case Parameters.....	178
Table 3-65.	ModalTransient-Inverse Solution Case Parameters.	181
Table 3-66.	QEVF Solution Case Parameters.....	183
Table 3-67.	A 2005 Comparison of Quadratic Eigenvalue Problem Methods. Although Ceigen and Anasazi have changed substantially since 2005, this table has not been updated to reflect those changes.....	185
Table 3-68.	Options for qevp Anasazi Solutions.	185
Table 3-69.	ceigen Solution Case Parameters.	186
Table 3-70.	Ceigen Tests.	187
Table 3-71.	SA_Eigen Options.	189
Table 3-72.	Verification Summary for SA_Eigen.	190
Table 3-73.	Projection_Eigen Options.....	192
Table 3-74.	NLStatics Solution Case Parameters.....	192
Table 3-75.	NLTransient Solution Case Parameters.....	194
Table 3-76.	ReceiveSierraData Solution Case Parameters.	195
Table 3-77.	Nodal data used in receive_sierra_data	197
Table 3-78.	Element data used in receive_sierra_data . Only the Neo-Hookean and hyperfoam Lamé models are supported; see Section 4.1.4.	198

Table 3-79.	Data Transferred in <code>receive_sierra_data</code> specific to <code>orthotropic_layer</code> Material. Some of this data is transferred automatically. Other data requires the from_transfer keyword in the material definition.	199
Table 3-80.	List of valid label keys for <code>receive_sierra_data</code>	201
Table 3-81.	Statics Solution Case Parameters.	201
Table 3-82.	Superposition Solution Case Parameters.	202
Table 3-83.	Tangent Solution Case Parameters.	203
Table 3-84.	TranShock Solution Case Parameters.	204
Table 3-85.	Transient Solution Case Parameters.	206
Table 3-86.	Transient-Inverse Solution Case Parameters.	210
Table 3-87.	TSRpreload Solution Case Parameters.	213
Table 3-88.	Residual Vectors Solution Case Parameters.	215
Table 3-89.	GeometricRigidBodyModes Solution Case Parameters.	217
Table 3-90.	Waterline Solution Case Parameters.	219
Table 3-91.	gapremoval Solution Case Parameters.	223
Table 3-92.	Optimization Section Parameters	229
Table 3-93.	Optimization Section Parameters for Line Search	229
Table 3-94.	Optimization Section Parameters for Trust Region	229
Table 3-95.	Inverse-Problem parameters	230
Table 3-96.	Block Section Parameters for Material Inversion	238
Table 3-97.	Material Section Parameters for Material Inversion	240
Table 3-98.	Block Section Parameters for Material Inversion	241
Table 3-99.	Table of Supported Material Parameters for Inverse Methods	244
Table 3-100.	Parameters in material section for Damage Identification	245
Table 3-101.	Parameters in inverse-problem section for Damage Identification	245
Table 3-102.	Loads Section Parameters for Force Inversion	246
Table 3-103.	Inverse Load Type Options	247
Table 4-104.	Material Stiffness Parameters.	249
Table 4-105.	Material Section Parameters for Fatigue Parameters.	258
Table 4-106.	Common Unit Scalings using Fatigue_Stress_Scale	260
Table 4-107.	Element Attributes.	270
Table 4-108.	General Block Parameters.	272
Table 4-109.	Non-Structural Mass Units.	274
Table 4-110.	Unhandled Corner Cases.	274
Table 4-111.	Combining NSM with <code>Density_Scale_Factor</code>	274
Table 4-112.	DAMPING Section Options.	276
Table 5-113.	QuadT, Quad8T, Tria6 Inputs.	286
Table 5-114.	QuadM inputs.	286
Table 5-115.	Nquad/Ntri inputs.	288
Table 5-116.	TriaShell input options.	290
Table 5-117.	Shell parameters that can be set via <code>attributes</code>	296
Table 5-118.	HexShell Verification Summary.	301
Table 5-119.	Attributes for Beam2.	304
Table 5-120.	Attributes and Parameters for Nbeam.	308
Table 5-121.	Ftruss Attributes and Parameters.	310

Table 5-122.	SpringDashpot Parameters.	316
Table 5-123.	Older Iwan 4-parameter model.	321
Table 5-124.	Revised Iwan 4-parameter model.	321
Table 5-125.	Line weld output: note that these are intended to exactly match the equivalent outputs in Sierra/SM	330
Table 5-126.	Built in Nmount Models.	337
Table 5-127.	Rbar Exodus Attributes.	339
Table 5-128.	Acceptable names of matrices within DMIG input files.	347
Table 5-129.	Rigidset Parameters.	349
Table 5-130.	Tied Joint Parameters.	352
Table 5-131.	Tied Joint, “Normal” and “Side” dependencies.	356
Table 6-132.	Dirichlet Boundary Enforcement Keywords.	360
Table 6-133.	Limitations for Prescribed Acceleration Boundary Conditions.	364
Table 6-134.	Available parameters for the infinite element section.	368
Table 6-135.	PML Element Parameters.	370
Table 6-136.	Parameters for Periodic Boundary Conditions.	373
Table 6-137.	ReadNodal function parameters.	376
Table 6-138.	Load Specification Keywords.	383
Table 6-139.	Random Pressure Inputs.	396
Table 6-140.	Rotating Frame Parameters.	400
Table 7-141.	OUTPUT Section Options.	411
Table 7-142.	Exodus Property Output Options.	413
Table 7-143.	Hex20 Gauss Point Locations	423
Table 7-144.	Data Files Written Using the MFile Option.	427
Table 7-145.	Typical Output.	429
Table 7-146.	FRF Output.	430
Table 7-147.	ModalRanVib Exodus Output.	430
Table 7-148.	ModalRanVib Frequency Output.	430
Table 7-149.	Typical Output.	431
Table 7-150.	ModalRanVib Frequency Output.	432
Table 7-151.	ModalRanVib Exodus Output.	432
Table 7-152.	TIndex parameters.	433
Table 7-153.	Element Orientation Outputs.	434
Table 7-154.	Element Orientation Interpretation.	434
Table 7-155.	Supported Statistical Data types	437
Table 7-156.	Selected Dynamic Matrix Definitions.	438
Table 7-157.	Variables that are output from ddam analysis.	439
Table 7-158.	Frequency Value Specification Methods.	447
Table 7-159.	Element Stress Truth Table.	450
Table 7-160.	Echo Section Options.	454
Table 8-161.	Tied Surface Parameters	463
Table 8-162.	Coordinates of Face (red) and Nodes (blue).	473
Table 8-163.	Conventional Constraint Equations.	474

This page intentionally left blank.

Acknowledgments

The **Sierra/SD** software package is the collective effort of many individuals and teams. A core Sandia National Laboratories based **Sierra/SD** development team is responsible for maintenance of documentation and support of code capabilities. This team includes Dagny Beale, Gregory Bunting, Mark Chen, Nathan Crane, David Day, Clark Dohrmann, Sidharth Joshi, Payton Lindsay, Julia Plews, Brian Stevens, and Johnathan Vo

The **Sierra/SD** team also works closely with the Sierra Inverse and Plato teams to jointly enhance and maintain several capabilities. This includes contributions from Wilkins Aquino, Brett Clark, Murthy Guddati, Sean Hardesty, Chandler Smith, Benjamin Treweek, and Timothy Walsh.

The **Sierra/SD** team works closely with other Sierra teams on core libraries and shared tools. This includes the DevOps, Sierra Toolkit, Solid Mechanics, Fluid Thermal Teams. Additionally, analysts regularly provide code capabilities as well as help review and verify code capabilities and documentation. Individuals directly contributing to the **Sierra/SD** documentation or code base during the last year include Samuel Browne, Victor Brunini, Pete Coffin, Jonathan Clausen, Gabriel de Frias, Mike Glass, Jacob Koester, Mark Merewether, Scott Miller, Kendall Pierson, Vincent Pericoli, Timothy Smith, and Alan Williams.

Historically dozens of other Sandia staff, students, and external collaborators have also contributed to the **Sierra/SD** product its documentation.

Many other individuals groups have contributed either directly or indirectly to the success of the **Sierra/SD** product. These include but are not limited to;

- Garth Reese implemented the original **Sierra/SD** code base. He served as principal investigator and product owner for **Sierra/SD** for over twenty years. His efforts and contributions led to much of the current success of **Sierra/SD**.
- The ASC program at the DOE which funded the initial **Sierra/SD** (Salinas) development as well as the ASC program which still provides the bulk of ongoing development support.
- Line managers at Sandia Labs who supported this effort. Special recognition is extended to David Martinez who helped establish the effort.
- Charbel Farhat and the University of Colorado at Boulder. They have provided incredible support in the area of finite elements, and especially in development of linear solvers.
- Carlos Felippa of U. Colorado at Boulder. His consultation has been invaluable, and includes the summer of 2001 where he visited at Sandia and developed the HexShell element for us.
- Danny Sorensen, Rich Lehoucq and other developers of ARPACK, which is used for eigenvalue problems.

- Esmond Ng who wrote *sparspak* for us. This sparse solver package is responsible for much of the performance in **Sierra/SD** linear solvers.
- The *metis* team at the university of Minnesota. *Metis* is an important part of the graph partitioning schemes used by several of our linear solvers. These are copyright 1997 from the University of Minnesota.
- Padma Raghaven for development of a parallel direct solver that is a part of the linear solvers.
- The developers of the SuperLU Dist parallel sparse direct linear solver. It is used through GDSW for a variety of problems.
- Leszek Demkowicz at the University of Texas at Austin who provided the HP3D³⁰ library and has worked with the **Sierra/SD** team on several initiatives. The HP3D library is used to calculate shape functions for higher order elements.

This work was supported by the Laboratory Directed Research and Development (LDRD) program.

Release Notes

0.1. *Release 5.4*

New or Improved Features

The closest distance user output functionality has been expanded to also be able to compute a nodal field which is the closest distance at each node. See Section [7.2.2](#).

The TSR Preload verification document was added to the Verification Manual.

New commands were added to allow scaling and shifting of XY functions. See Section [2.8.1](#).

New user output options for computing the max, min, average, or standard deviation of element variables. See Section [7.2.3](#).

Sierra/SD now works consistently with the **sierra** script `-pre` option. This option will automatically detect the mesh file used by the **Sierra/SD** analysis and decompose it in parallel if this has not yet been done.

It is now possible to provide alternative names for exodus fields (such as displacement, stress, etc.) when importing data from a previous solution. See Section [3.27](#).

Material properties can now be defined as general functions of element or nodal variables. See Section [4.4.7](#).

Deprecated Features

The **Sierra/SD** syntax for including files in an input file, such as `#include material.inp` or `include blocks` is deprecated, in favor of the Aprepro syntax used by other Sierra codes. For example `{include(waveguide.inp)}`. Currently, 111 tests use the deprecated syntax, and 4 tests use the correct syntax.

Other features that have been newly deprecated or previously deprecated and removed:

- Constraint output syntax `slave_constraint_info`
- Contact `master` and `slave` options. Use `side b` and `side a` instead.
- Solution case `QmodalFrf`
- Echo `none`, `mpc_nonorthog` options
- GDSW parameter `atLeastOneIteration`
- RBE3 option `method old|new`.
- non-aprepro input deck file includes
- Option `thickness varies by element`

- File `sierra_input_file` command
- Old `Coordinate` section syntax. Use `begin rectangular coordinate system`, `begin cylindrical coordinate system` etc. instead.
- Material name option.

Bug Fixes

Maintenance and development prioritizes use cases. Input is implemented using Functions. Some Functions can use data from either the input deck or the Exodus mesh file. These Functions can interpolate, extrapolate and integrate (once or twice).

Changes were made so that the Linear, 1D Table, SM and Spatial BC Functions have consistent behavior. Spatial BC Functions include Random Lib Functions. SM refers to Sierra SM Piecewise Linear Functions. In several cases, the behavior changed from possibly inconsistent to consistent. The Functions are now unit tested and documented.

Table Functions are linear Functions optimized for constant time step size. A Table Function dimension can be 1,2,3, or 4. One dimensional Table Functions have high priority. The handling of edge cases for 1D Table Functions improved. It is possible, albeit unlikely, that simulation results will change. Higher dimensional Table Functions have the old behavior. They are complex. Their handling of edge cases is not known by current developers. The priority of 2D Table Functions requires clarification. 3D and 4D Table Functions have low priority.

The behavior of SpatialBC and RandomLib Functions was improved. Previously, these Functions were incorrect in a way that would have artificially forced the value to be zero at $t = 0$, and also would have shifted all the values as if the origin was at $t = 0$. Regression test results were changed (see sha a743640134329a08ddeab8ab69 in the Test repository). Spatial BC Functions have low priority.

Contact tolerances for shell-to-shell contact have been improved. Previously tolerances only used half the shell thickness in the search tolerances and could miss contact.

The contact `cutoff` which is used to exclude nodes from contact based on field variable values now only considers the field value on the node of the node-on-face constraint. This is more consistent with how Sierra/SM writes such variables, the handoff of which is the primary use case for this capability.

Fixed a bug where the attribute order could be scrambled in the output results mesh. This would cause issues if that output mesh was used as the input of a subsequent SD analysis.

Now using a `solver_options` for global GDSW parameters will override the any global options specified. Previously the global options would always be used.

0.2. Release 5.2

Sierra/SD is considered production ready for the ATS-2 platform which utilizes GPUs. Most commonly used solution cases such as Eigen and transient show substantial speedup using this platform's GPUs. Detailed documentation on how to run **Sierra/SD** on ATS-2 can be found in the online wiki "Running Sierra/SD on ATS-2". Contact sierra-help@sandia.gov to locate this information.

A capability has been implemented for in-core transfer of loads from a simultaneously running SPARC analysis. See Section 6.2.6.

Several improvements have been made to automatically handle potentially over-determined constraints produced by self contact or contact between surfaces that also share nodes.

Thanks to Dagny Beale for reviewing the Coordinates section (2.8).

User mounts may now be incorporated into **Sierra/SD** at runtime.

The `nskip` parameter can now be set independently in the linesample input.

Analytic functions can now be spatially dependent functions of coordinates or other variables. See Section 2.8.9 for more details.

Rain flow cycle counting of stress cycles and fatigue estimation from these cycles is now available in transient solution. See Section 7.1.22. Additionally an example of this capability has been added to the "How To" manual.

Sierra/SD will disallow use of certain reserved words for thing like block names as they would cause ambiguity in parsing. For example a block named "end". See Section 2.2.4 for more details on how to determine reserved names and also how to override this error handling for legacy models.

Several improvements have been made to the eigen value sensitivity output. Previously the column labels for output did not accurately describe what these outputs represented.

Many improvements have been made to robustness for output and history files. This was accomplished making it possible to request every output in any solution case. If additional issues are seen, then please report them to sierra-help@sandia.gov.

A new QUIVER video is available on solver debugging. Additionally, training videos recorded from a several day in-person class are available on the Computational Simulation Knowledge Base website. Contact sierra-help@sandia.gov to locate them.

The input deck is now by default echoed to the **Sierra/SD** log file. This echoing can be turned off by specifying "input off" in the echo block.

The default solver shift for eigen solution has been changed from "-1" to "-1.0e+6". Ideally the shift should be set to the negative of the first eigen value of for the system (the square of two pi times the first flexible frequency.) The new default is much more likely to be in the right ballpark for most engineering systems. However, consideration still needs to be

given to shift for hard to solve problems. Additionally, the default maximum size of type 1 constraints was raised from 100 terms to 250 terms to improve solver robustness.

The input syntax to create exodus assemblies within the **Sierra/SD** input deck was renamed from `MESH GROUP` to `BEGIN ASSEMBLY` for greater consistency with other Sierra applications.

The TriaShell mass matrix calculation has been improved. Previously it was missing the inertia terms relating to the shell thickness. In practice the effect of these terms is usually slight as shell structures should be relatively thin compared to element size.

Deprecated Features

Features that have been newly deprecated or previously deprecated and removed:

- Material syntax `name=string`, `T_current=real`, `s_isotropic`
- Constraint output syntax `slave_constraint_info`
- Tied data and contact syntax `master`, `slave`
- Joint2g option `Shys`
- Solution case `QmodalFrf`
- Solution case `Qmodaltransient`
- `ncfout` (output4, op4) utility,
- Solution parameter `no_symmetrize_struc_acous` (has been replaced by `symmetrize_struc_acous`)
- RBE3 option `method old|new`.
- GDSW parameter `atLeastOneIteration`
- Old `Coordinate` section syntax. Use `begin rectangular coordinate system`, `begin cylindrical coordinate system` etc. instead.
- Echo `none`, `mpc_nonorthog` options
- File `sierra_input_file` command
- Coupling to the deprecated “Aero” application
- Runtime compiled functions (RTC)

0.3. Release 5.00

A new chapter of *User's Manual* describes how to run Sierra/SD §1.1 and other fundamental usability issues. Also, Sierra/SD tutorials for new users are available <https://snl-wiki.sandia.gov/display/CKB/How-to+articles>

An **Exodus** hierarchical data structure, the Assembly,⁶⁴ has been developed to make building models easier. Sierra/SD has full support for Assemblies §2.2.2

The GDSW linear solver is compatible with certain Graphical Processing Units. Users may notice more precise timing information of GDSW. This release includes a full featured interface §2.5 to the SuperLU Dist parallel sparse direct linear solver.

The default GDSW preconditioner is designed for poorly conditioned problems, and is a less efficient option for well conditioned problems. An adaptive block diagonal preconditioner §2.5.6 can be used by setting `preconditioner_type` to 3. The options §2-15 `identify_low_quality_elements` and `max_element_condition` enrich the diagonal preconditioner based on mesh element quality.

Sierra/SD and Sierra/SM use consistent input syntax and share source code for coordinate systems §2.7. Coordinate system definition options have been added such as defining coordinate system by nodesets. Ellipsoidal coordinate systems are now available.

The syntax for spatial boundary condition functions that read data from the mesh is more consistent. It is necessary to specify the variable that the function reads. For consistency with the methods for input from the mesh, the keyword `variable_name` has been replaced by the keyword `exo_var`.

Traction input is supported. The `ReadSurface` function §6.2.4 will input spatially dependent and time dependent data from the mesh file. The data is automatically integrated to determine the corresponding forcing function.

Huge improvements in the performance of the Modaltransient §3.22 and ModalFrf §3.18 methods have made it possible to remove the Qmodaltransient and QmodalFrf methods. The output capabilities of Modaltransient and ModalFrf solutions have been extended to reproduce all the capabilities of Qmodaltransient and QmodalFrf.

Another capability to compute all the eigenvalues and eigenvectors for problems with modest numbers of dofs has been extended to handle several important cases such as models with a singular mass matrix and models with multipoint constraints.

The geometric rigid body mode solution case now works for models with any Sierra/SD element including Superelements (c.f. Jira Compsimhd-5610)

Users may use the input **Exodus** mesh to set spatially dependent material properties fields §4.4.7.

Acoustic boundary conditions §6.3.7 include the `point_volume_acceleration`. An acoustics problem can use void elements in the structural mesh to represent air. The void elements move with the solid elements. The flux air volume causes an acoustic pressure

wave to propagate out. One approach is to map the volume flux term in the solid to the closest acoustic element. A new approach is to map the volume flux term to the void elements representing the air.

HexShell element §5.8 input syntax is compatible with new coordinate frame syntax. Users must specify the thickness direction of the elements. One way to do this is through a coordinate frame. The frame may also be specified by name instead of the id, which is clearer.

All Sierra/SD elements support stress output. The corresponding issues are Jira Compsimhd-9550 and Nesm-2853.

The lumped mass matrix, **MLumped**, can now be written to the **Exodus** file. Also, the modes shapes from the Eigen solution case multiplied by the mass matrix, **MPhi**, can be written to the **Exodus** file too. The corresponding ticket is Jira Compsimhd-9537.

Principal stress §7.1.19 is output to the **Exodus** file using the **principal_stresses** keyword. Stress output includes a signed von Mises stress §7-141 and §7.1.21. It has been added to support fatigue estimation using Rain Fall cycle counting.

If rigid body filtering is applied using **FilterRbmLoad**, then Force output includes the filtered forces using syntax consistent with Sierra/SM.

Coordinate frames specified in MPCs §2-32 can be defined by either the coordinate id or the coordinate name.

The output for the local or element coordinate frame using **EOrient** §7.1.43 is consistent for all element types. The local coordinate frame can also be output using §7.1.6. This capability is required for anisotropic materials §4.1.3.

begin contact definition §8.2.3 is more general. Sliding is supported by using the **frictionless** friction model. The capability provides a linearization of equivalent Sierra/SM syntax. In response to Jira Compsimhd-9517, Contact search failure, gap removal diagnostic information includes an Element Quality table. This applies to **tied data** too. Note that Sierra/SD reports an element condition number (greater than or equal to one), and Sierra/SM reports the element quality (less than one), which is multiplicative inverse of the condition number.

Sierra/SD has an undocumented interface to a code named Nemo. This release cycle, Sierra/SD added support for two messages. The **FailedElementIds** message with **mpi_tag** 1601 reports an integer boolean for every shared element face. It was added for consistency with Sierra/SM. A face is either inactive, 0, or active 1. Sierra/SD reports all elements as active. The **SidesetIds** message with **mpi_tag** 1701 reports the integer sideset global ID associated with each coupled face. It is used for diagnostic information between the codes. The IDs are the global ID numbers of the **Exodus** mesh. Further work is needed on the case of an element face belonging to multiple coupled sidesets.

Deprecated Features

The Sparsepak sparse serial linear solver is deprecated and has been removed. GDSW with the new option `krylov_solver = none` produces bit-wise equivalent results.

Most deprecated features are ready to be removed immediately after release: material syntax `name=string`, `T_current=real`, `s_isotropic`, constraint syntax `slave_constraint_info`, `master`, `slave`, property line `Shys`, solution case `QmodalFrfr`, solution case `Qmodaltransient`, `ncfout` (output4, op4) utility, `no_symmetrize_struc_acous` (to be replaced by `symmetrize_struc_acous`).

Certain deprecated features may be removed before the next release: RBE3 option `method old|new` and `rtc` user functions.

Newly deprecated features include: `elmat`, old syntax for coordinate frames, constraint method transform.

The SD coordinate system syntax is changing and the old syntax is deprecated. To aid in the syntax transfer, a Python 2 script named `convertCoordinateSystems.py` is available to convert existing coordinate blocks in an input deck to the new syntax.

```
$ python convertCoordinateSystems.py FILE_1 FILE_2 FILE_3 ...
```

Files are overwritten by this script. A message is printed for each file touched. All inputs other than `coordinate` definitions remain untouched. Comments are preserved everywhere in the file except within the coordinate system definition. It accepts any number of file arguments, so you can chain it together other Linux tools,

Example: In the current directory, convert any *.inp file containing "coordinate":

```
$ python convertCoordinateSystems.py $(grep -l -i coordinate ./*.inp)
```

Example: Below the current directory, convert any *.inp file, recursively:

```
$ python convertCoordinateSystems.py $(find . -name *.inp)
```

Bug Fixes

Loads are applied correctly at the junction between dead element blocks and element blocks that are not dead for any partition of the mesh into subdomains. Jira Compsimhd-9411 had reported an error in the case in which, for one of the subdomains sharing the junction, all the elements attached to the junction were in dead element blocks. The `parallelWithDeadBlocks` nightly verification test was added.

An Air Force code verification project found defects with Sierra/SD layered shell elements (Jira Compsimhd-9822). First in a shell element with one layer the `layer` keyword is redundant and optional. The incorrect behavior, now fixed, was that a single-layer shell with a non-trivial fiber orientation and without the `layer` keyword, Sierra/SD incorrectly used trivial fiber orientation. A unique attribute value can be specified for each element in

an element block in the **Exodus** file. If specified in the input deck, then an element attribute becomes a block parameter. For block parameters a single value is applied to all the elements in the element block. For layered shell elements thickness and fiber orientation attributes in the **Exodus** file were considered to be block parameters. One value was applied to all the elements in the element block. This is now fixed. The layerShells test was added to cover these cases.

Jira Compsimhd-9343 does not involve a code defect, but is notable due to reporting incorrect results for a transient simulation. It is important to bear in mind that some solution cases are reliable only if used correctly. In this case users are expected to select sufficiently small step sizes and linear solver tolerances.

Multiple users had trouble using the Seacas tool Aprepro. All issues have been resolved.

Issues using `tied data` and begin contact definition have continued, leading to usability improvements of begin contact definition. The gap removal QUIVER video

<https://snl-wiki.sandia.gov/display/CKB/How-to+articles> is on `tied data` and begin contact definition.

0.4. Release 4.58

Deprecated Features

Algorithmic terminology referencing master/slave has been changed. The purpose is both to increase technical clarity and use more inclusive terminology. Most of these changes effect documentation only. One behavior change is that the output variable 'slave_constraint_info' has been deprecated use 'constraint_info' instead. Updates have also been made to the exodus field names 'constraint_info' output produces. Additionally, master/slave references in the 'contact definition' and 'periodic bc' capabilities have been updated to be consistent with Sierra/SM terminology in the equivalent capabilities.

Previously deprecated features have been removed.

- Output warninglevel cannot be an integer,
- solution methods blk_eigen and old_transient,
- feti-dp linear solver & option rbm_tol_svd,
- history and frequency output option auto-join,
- patch negative elements (parameter),
- exodus_file & restart_file functions,
- Eigen option nmodes_acoustic,
- spherical_wave option point_source_origin,
- planar_step_wave option C0,

- tiedjoint options edge tolerance and Newton tolerance.
- Boundary conditions smooth angle, smoothing distance, smoothing resolution,
- echo timing.

New Features

Sierra/SD supports coupled electro-mechanical physics in order to simulate the electro-mechanical behavior of piezoelectric materials when subjected to an electric field or mechanical stress. This support includes static, transient, Eigen, and direct frequency response solution methods, piezoelectric and dielectric material models, and electric potential and surface charge boundary conditions. For the inverse solution methods, electro-mechanical physics enables the user to define experimental data in terms of measured voltages.

A new option has been added to **contact definition** to remove contact constraints based on an exodus mesh variable. The target use case is removal of open frictional constraints based on a solid mechanics preload. See Section [8.2.1](#) for details.

Output of reaction and constraint forces for static and transient analyses is a production feature.

A new output **relative_disp** is available to track the relative displacement between the ends of a Joint2G. The output contains a statistical description of the gap for modal random vibration or a numerical value for the gap in Eigen, static, and transient analysis. See Section [7.1.40](#).

An XML description of the valid commands can be generated with a **-create-xml** command line argument. The primary purpose of this file is to provide a description of valid command lines to SAW.

A conical coordinate system is available. See [2.7](#).

It is possible to connect multiple tied joints to the same block.

All element types may use temperature dependent material properties. Previously only solid elements could use the temperature dependent properties.

A new capability was added to scale density by block. See [4-108](#).

The line weld has an option to either remove or not remove gap in the constraints it generates.

Syntax that matches the Sierra/SM piecewise linear function is supported in Sierra/SD .

The documentation has been reorganized following.⁶⁹ The Table of Contents and Index have been rewritten to make finding information easier, especially for new users.

Bug Fixes

The cbmap MATLAB format superelement output has been corrected. Previously this map was missing rows corresponding to internal degrees of freedom.

Multiple bugs have been corrected in the Qmodaltransient and QmodalFrf capability relating to scrambled output or inconsistent use of modal filters. The Qmodal capability is expected to be deprecated in a future release in favor of a faster Modaltransient and ModalFrf capabilities.

Several robustness improvements have been made to the **FilterRbmLoad** option. It works properly with both static and transient analysis.

The buckling capability works with gravity loads and constraints.

Modal random vibration properly handles coordinate system in the frequency file outputs.

The element quality metric printed to the log file has been improved. Tetrahedron quality was off by a factor of three, and it could be misreported for elements with small volume.

A `geometric_rigid_body_modes` analysis interacts more correctly with the Eigen solution.

An error was fixed where a single layer hexshell with temperature dependent material properties would use incorrect material properties.

The Navy triangular shell finite element incorrectly transmitted loads. Only the Navy uses the element. Correcting the element lead to re-blessing at least five regression tests.

Syntax and Defaults Changes

Modal random vibration truncation method default has changed from displacement to none. The displacement truncation could truncate modes important to the solution in some cases, such as modes that are localized to a small portion of the domain.

1. HOW TO RUN Sierra/SD

Sierra/SD provides a massively parallel implementation of structural dynamics finite element analysis. This capability is required for high fidelity, validated models used in modal, vibration, static and shock analysis of weapons systems. General capabilities for modal, statics and transient dynamics are provided. The **Sierra/SD** software evolved from the “Salinas” package.

Sierra/SD tutorials for new users are available
<https://snl-wiki.sandia.gov/display/CKB/How-to+articles>

This section describes the command line arguments and workflows typical of **Sierra/SD** analysis. This section is primarily applicable to analysts at Sandia. Sections 2 through 8 describe **Sierra/SD** capabilities invoked via the input deck. Section 9 provides example problems.

1.1. Accessing Sierra/SD

At Sandia Sierra is installed on many Unix-based High Performance Computing (HPC) systems. This includes dedicated and shared blades, shared computational clusters (e.g. cee-compute) and large queued clusters (Institutional Clusters such as Eclipse). To run Sierra applications you need to have access to one of these machines. You also must request “SIERRA Analysts Code Access” through WEBCARS.

1.2. Modules and Executables

Configuration of Sierra applications is controlled via modules. The first step to running a Sierra application is to load the appropriate module for the desired version. For example:

```
workstation_prompt> module load sierra
```

Several modules are commonly used.

- **sierra**: This is the latest released version of Sierra and is generally the recommended version to be used.
- **sierra/X.YY**: This will load a specific Sierra version, for example “sierra/4.58”. Sierra releases are done twice a year and generally the previous four releases are available.
- **sierra/sprint**: This is the latest “sprint snapshot” version of Sierra. A new sprint snapshot is installed every three weeks. Sprint snapshots contain the latest developed features but have less quality assurance testing than the main releases. Sprint snapshots should primarily be used to do beta evaluation of newly developed features.

- **sierra/daily**: This is the development version of Sierra built on the previous night. No quality assurance is done on this version. This version should only be used when directly working with the development team to help evaluate a bug fix.

1.3. *The Sierra/SD salinas Executable*

The primary executable for running structural dynamics analysis is called “salinas”. It uses a text format input deck to configure a simulation. This users manual primarily describes the input deck options and format. An example invoking a basic serial **Sierra/SD** analysis is:

```
workstation_prompt> module load sierra
workstation_prompt> salinas -i input.inp
```

The commonly used command line arguments to the salinas executable are given in Table [1-1](#).

Table 1-1. – Command Line Options.

String	Descriptor
-h	Prints help message
-i	Path and name of the input deck. The input deck is traditionally given the extension <code>.inp</code> . However, any name is allowed. If no '-i' is given the last argument of the line is assumed to be the input deck name.
-l or -o	Name for the output log file. By default, output to <code><input_deck_name>.rslt</code> in serial or <code><input_deck_name>_0.rslt</code> in parallel.
-n	Do not overwrite any existing output log and diagnostic <code>.dat</code> files. Do write a new file appending a 1, 2, 3, etc. to the name.
-d	Change working directory for the run. Equivalent to <code>cd</code> to that directory and running <code>salinas</code>
-define	Define variables for Aprepro pre-processing of the input deck. See Section 2.1 for details
--check-syntax	Do initial syntax checking of the input deck and then stop before reading mesh or doing any large calculations. This option can be used to debug model input in serial prior to submitting to a queued system.
-beta	Enable use of 'beta' capabilities. These are capabilities in active development, are subject to change without notice, and may have less rigorous testing.
-nt	Number of threads to use per MPI rank. Using more threads than physical cores (i.e., hyper-threading) may result in decreased performance. This option is only supported on a few platforms. Ask the Sierra/SD team for more information.
-ndevices	Number of GPUs to use per hardware node. This option is in development and not ready for general use.
-create-xml	Create an xml format command specification. Primarily for interacting with graphical input file creation utilities such as SAW.

1.4. *MPI Parallel Execution*

Sierra/SD may be run either in serial or in parallel with MPI. In parallel **Sierra/SD** has demonstrated efficiency to thousands of processor cores.¹⁵ Some basic work-flow examples for parallel execution are provided below. Information on using on the current platforms is available at High Performance Computing website <https://computing.sandia.gov>

Additional steps are necessary to execute in parallel instead of serial. The following examples use the input deck `input.inp` and the **Exodus** mesh file `mesh.exo`.

1.4.1. Number of MPI Processes Needed

MPI is an optimization that can decrease run time. The configuration of Sierra depends on the architecture. A central processing unit (CPU) has multiple cores and possibly a graphical processing unit (GPU) accelerator. These notes only discuss the CPU. A user must know one thing about the CPU, the number of cores per node (or processor). One way to determine the cores per node is to use the `lscpu` command and look for **Core(s) per socket**. Another way is `cat /proc/cpuinfo`. Look for **cpu cores**.

Running **Sierra/SD** in parallel requires the user to specify how many MPI ranks will be used. Choosing the number of cores so that the number of elements per subdomain approximately 5000 typically works fine. Generally a computational “node” will run several MPI ranks and many computational nodes may be used simultaneously. Selecting the number of ranks to use is based on three concerns. First, each computational node has finite memory and enough nodes must be used to fit the problem in memory. Second, use of more MPI ranks and more computational nodes can reduce time to solution. Third, efficient use of machines should be considered as no analysis will achieve perfect parallel efficiency and using many computational cores on a small problem may provide sub optimal machine utilization. The appropriate number of MPI ranks to use is primarily determined by how many degrees of freedom (DOF) are in the model. Though machines are variable Table 1-2 can be used as a rough guide help guide selection of the number of MPI ranks and thus computational cores to use.

mem/core	dofs per core	Num Cores Needed
1GB	15,000	$dofs/15,000$
2GB	30,000	$dofs/30,000$
4GB	50,000	$dofs/50,000$
8GB	70,000	$dofs/70,000$
16GB	90,000	$dofs/90,000$

Table 1-2. – Determining Number Of Processors Needed.

Memory use depends on a variety of factors including the element type used, the solution strategy and the output processing. The numbers in the table are generally conservative. Fortunately, in most cases it is not critical unless there is not enough memory on the compute node.

For high memory use analyses it may be necessary to run **Sierra/SD** with half the maximum number of cores available per computational node. The reason is that using half MPI ranks doubles the memory assigned to each rank. As there is significant memory overhead for an MPI rank, memory limited analyses will benefit more by adding more memory to each rank than spreading the mesh to a larger number of ranks.

1.4.2. Mesh Decomposition

If running in parallel, prior to running the **salinas** executable first the **Exodus** mesh file must be parallel partitioned. Mesh decomposition involves dividing the input mesh file into several pieces, one of which will be read by each MPI rank. Several tools exist for this mesh decomposition step, the recommended tool is **stk_balanceparallel** **computing!stk_balance**. For more details on **stk_balance** see.⁷⁰

In the next example **stk_balance** is executed on a non queued blade.

```
workstation_prompt> launch -n 16 stk_balance mesh.exo split
```

This command will decompose the mesh file into 16 parts with names like **mesh.exo.16.00**, **mesh.exo.16.01**, etc. and place those spread files into the directory 'split'. Storing the decomposed mesh files in a separate directory is optional. The location of the input mesh file is specified via commands in the input deck as described in Section 2.2.

1.4.3. Running the Sierra/SD Executable in Parallel

With multiple ranks, the **Sierra/SD** executable must also be run in parallel.

An example of running **Sierra/SD** in parallel on a non queued blade would be

```
workstation_prompt> launch -n 16 salinas -i input.inp
```

See Section 1.6 for more examples on different systems.

1.4.4. Post Processing in Parallel

In parallel **Sierra/SD** creates several output files as described in Section 7.

- **log**: In serial the default log file name is **<input_deck_name>.rslt**. In parallel the default name is **<input_deck_name>_0.rslt**. This is a text output file giving vital information about the run. The '**_0**' indicates the is file written by processor zero. Information relevant to the global model behavior is accumulated to processor zero for output.
- **Mesh based Exodus**: The output name will depend on the structure of the input deck, but usually follows a pattern like **<mesh_name_base>-out.<mesh_name_extension>**. This includes mesh based information such as nodal displacement. Similar to the decomposed input mesh this output **Exodus** file will be written as a decomposed file with one portion per MPI rank. Some post processing tools such as Ensign or Paraview can read the decomposed **Exodus** file directly. For other tools the file must first be combined to a single file with the Seacas tool 'epu'.³¹

- **Diagnostics:** This could include detailed information from the solver (`dd_solver.dat`) or other text and MATLAB format output files.

1.5. *File system concerns*

Sierra/SD can output large files and works best when writing to fast output systems. Running analyses on scratch drives such as `’/nscratch’` on HPC system or local `’/scratch’` on blades are ideal. Large shared disks such as `’/gpfs1’` are viable though may incur some runtime overhead. Drives that are not mounted by computational nodes of HPC systems such as `/tmp` should be avoided.

1.6. *Workflow Examples*

There are many ways to run **Sierra/SD** depending on the analysis size and machine used. Several common examples are presented.

The next commands run a serial job on a blade and examine the log file.

```
workstation_prompt> module load sierra
workstation_prompt> salinas -i input.inp
workstation_prompt> cat input.rslt
```

Run a parallel job using four processors of a blade. Look at exodus output results via the textual output query tool `explore`.⁶³

```
workstation_prompt> module load sierra
workstation_prompt> launch -n 4 stk_balance mesh.g
workstation_prompt> launch -n 4 salinas -i input.inp
workstation_prompt> module load seacas
workstation_prompt> epu --auto mesh-out.g.4.0
workstation_prompt> explore mesh-out.g
```

Run a parallel job on a queued system using 72 MPI ranks and 2 computational nodes manually calling `sbatch`, then combine to a single file, all with a single queue submission. In this case assuming a machine like Eclipse where each computational node can run up to 36 MPI ranks. The Seacas tool `’epu’`³¹ can run in parallel, but is usually run on only a few MPI ranks.

```
Contents of submit.sh
#!/bin/bash
launch -n 72 stk_balance -i mesh.g
launch -n 72 salinas -i input.inp
launch -n 8 epu --auto mesh-out.g.72.00
```

```
workstation_prompt> module load sierra
workstation_prompt> module load seacas
workstation_prompt> /usr/bin/sbatch --account=<acct_id> --nodes=2
--time=4:00:00 ./submit.sh
```

Run a parallel job on a queued system using 8 MPI ranks spread to 4 computational nodes get access to more memory per computational process. In this case assuming a machine like Eclipse where each computational node defaults to 36 ranks.

```
Contents of submit.sh
#!/bin/bash
launch -n 8 stk_balance -i mesh.g
launch -n 8 salinas -i input.inp
launch -n 8 epu --auto mesh-out.g.8.0

workstation_prompt> module load sierra
workstation_prompt> module load seacas
workstation_prompt> /usr/bin/sbatch --account=<acct_id> --nodes=4
--time=4:00:00 ./submit.sh
```

Run a parallel job on a queued system using the “sierra” script. If the mesh file was not yet decomposed for parallel execution the sierra script will automatically decompose it prior to running the **salinas** executable. Epu would run in serial on the HPC head node. Note the sierra script can be used to submit **Sierra/SD** analyses. However, the sierra script ‘-post’ option is currently not functional with **Sierra/SD** (for other Sierra tools the post option automatically combines output mesh files, similar to the below “epu” command.) The sierra script can potentially be helpful to translate the number of MPI ranks to the correct number of nodes to use on a given machine. Additionally, the same sierra script submission command can be used on queued HPC systems and non queued blades and clusters.

```
workstation_prompt> module load sierra
workstation_prompt> sierra -j 72 -T 4:00:00 salinas -i input.inp
workstation_prompt> module load seacas
workstation_prompt> epu --auto mesh-out.g.72.00
```

Run coupled **Sierra/SD** Fuego analysis on a non queued system for fluid structure coupling. In this case salinas will use 8 MPI ranks and fuego will use 16 MPI ranks for a total use of 24 MPI ranks. Note, this is an advanced use case.

```
workstation_prompt> module load sierra
workstation_prompt> launch -n 8 stk_balance salinas_mesh.g
workstation_prompt> launch -n 16 stk_balance fuego_mesh.g
workstation_prompt> launch -n 8 salinas -i salinas_input.inp
--fuego-coupling : -np 16 fuego -i fuego_input.inp
```

1.7. *Thread Parallelism*

In addition to decomposition based MPI parallelism, **Sierra/SD** also supports thread parallelism on some platforms (currently Trinity and GCC development platforms). Threads are activated by the command line option “-nt <numThreads>”. The ‘numThreads’ given will be the number of OpenMP threads to use on each MPI rank. Threaded execution is most valuable on large models. Using a mixture of thread parallelism and MPI parallelism can give optimal performance when the number of MPI processes required would otherwise be very large. As a rule of thumb thread parallelism will provide benefit when exceeding about 200 MPI processes or when more cores are required than mpi ranks to obtain more memory. When using thread parallelism, the number of threads used times the number of mpi ranks used should be setup to be equal the total number of processor cores available on compute nodes.

Note: while the number of threads used in **Sierra/SD** is controlled by the command line option “-nt”, it is recommended that the user also set the environment variable ‘OMP_NUM_THREADS’ to be the same value. While **Sierra/SD** doesn’t depend on ‘OMP_NUM_THREADS’, there might be other aspects of your workflow that would, and so we recommend setting both to be consistent. In fact, **Sierra/SD** will output a warning if ‘OMP_NUM_THREADS’ does not match the value set by “-nt”.

1.8. *Troubleshooting*

A variety of issues can cause an analysis to fail. There are still bugs in the **Sierra/SD** software, and these will continue to be found. However, most problems are identified with problems in the model or other input to the software. This section may help to identify these issues with the goal of completing the analysis properly. The best troubleshooting strategy is to try to eliminate the modeling issues, and only then treat the problem as a potential bug.

Users can troubleshoot **Sierra/SD** issues through stand-alone tools or using **Sierra/SD** capabilities. The following sections will describe some ways to do this. First stand-alone tools are described. Second ways of using **Sierra/SD** capabilities to troubleshoot are described.

Queuing systems are inevitable in high performance computing. Any extra step lowering the risk of there being a mistake in the input file is worthwhile.

1.8.1. *Stand-Alone Tools*

Currently, two tools exist which can help the user debug their mesh file, i.e., **Exodus** file: **Explore** and **cubit**.

1.8.1.1. Explore Explore is an ACCESS/SEACAS utility that can be used to interrogate the **Exodus** file. One of the commands in **Explore** that can be used is **check**. It is used as follows:

```
prompt> explore cube.exo
.
.
EXPLORE> check

Database check is completed

EXPLORE>
```

If there are any warning or errors, they will appear before the **Database check is completed** message.

1.8.1.2. Cubit The Cubit team has developed a GUI-based tool named **cubit**. It can be used to visualize mesh quality parameters of the **Exodus** file. For questions about **Cubit**, please contact the Cubit team at **cubit-dev@sandia.gov**.

1.8.2. Using Sierra/SD To Troubleshoot

The checkout solution method [3.3](#) checks the input for errors. It skips matrix formation and solves. The dump solution method [3.9](#) also check the input for errors. The dump solution does form matrices, but like checkout skips the solves.

Syntax checking [2.3](#) is helpful.

The user has to take additional steps before executing the parallel version of **Sierra/SD**. One of the steps is to run **stk_balance** to decompose the finite element model. Using the **Exodus** file and the load balancing file, the next step is to run **nem_spread** to create the partitioned files on the parallel platform where **Sierra/SD** will be executed. Finally, the commands needed to run **Sierra/SD** on the parallel platform need to be learned so that execution of **Sierra/SD** can begin. Many of these steps can cause frustration to the user, but problems with any of these steps are often easily addressed.

All of these steps are worthwhile before running a large model on a queued system. Running the large model on a smaller computer (without a queue) to confirm that the simulation fails in the expected way (by running out of virtual memory) is also worthwhile.

The output includes the Approximate Matrix infinity norm ratios table. This table shows information about the stiffness, mass and dynamic matrices. The value shown is the ratio of the largest diagonal entry to the smallest diagonal entry. If the matrix has a 0 on the diagonal, then 0 is shown.

1.8.3. Modal Analysis

It is possible for the `eigen` solution method (discussed in Section 3.10) to diverge.

Section 2.3 mentions the `eig_tol` parameter. The default value is about 10^{-16} (for the Eigen solution case). Adequate modes can be determined with much larger values of `eig_tol`. Values such as 10^{-8} are reasonable.

At times an analyst may choose to use the `eigen` method to diagnose a problem. This can be done by using as large a value of `eig_tol` as is needed. The number of modes would also be as small as needed. After the issue is resolved, if eigenvectors are still needed, don't forget to reset `eig_tol` to a small value.

The modal analysis algorithm depends on a linear solver, and assumes that the linear systems are solved accurately. The shift is almost always negative. Increasing the magnitude of a negative shift forms linear systems that are easier to solve, and makes the eigenvalue problem harder to solve. The iterative linear solver (GDSW) parameters trade off between speed and accuracy. For example, a small value of the `solver_tol` increases accuracy and computational expense. Solving the linear system more accurately makes the eigenvalue problem easier to solve.

If changing the shift does not resolve your issue. And increasing `eig_tol` does not resolve your issue, you may consider inspecting convergence information from the code that solves the eigenvalue problem. To do this, increase `info` from 0 to 2 or 3. This is not supported with the `sierra` script.

```
workstation_prompt> launch -n 2048 salinas -info=2 -i deck.inp >& arpackinfo
```

These eigenvalues are shifted and inverted $\tau = 1/(\lambda - \text{shift})$, $\lambda = \text{shift} + 1/\tau$, $\lambda = \omega^2$, frequency = $\omega/(2\pi)$.

1.8.4. Evaluating Memory Use

The **Sierra/SD** software tends to use a lot of memory. Matrices are generated and solved, and while this is often the fastest method of solution, it results in large memory demands. Parallel computing has its own issues for memory use.

Memory use diagnostics can be requested in the “ECHO” section of the input (see Section 7.7).

1.8.5. Identifying Problematic Subdomains

The partition of a finite element mesh into subdomains may be visualized by rejoining the partitioned files into a new file using `epr` with the `-add_processor_id` option.

1.8.6. Limitations of SD Finite Elements

Element quality is a function of both the element geometry and the type of element. The Hex8, Wedge6, Tet4, Tria3, and QuadT elements implement a condition number. Note that the QuadT consists of two triangular elements.

Overall second and higher order finite element methods are much less sensitive to element quality than low order methods. However, there is a tipping point. A high order element will invert before⁶⁶ the corresponding low order element. With that in mind, it is noteworthy that **Sierra/SD** ignores this issue.

element	uses condition number of	element
Hex20		Hex8
Wedge15		Wedge6
Tet10		Tet4
CuTet10		Tet4
Tria6		Tria3
TriaShell		Tria3
NLayeredShell		QuadT
QuadS_GY		QuadT
NQuad		unknown

Table 1-3. – Elements using other elements condition number.

These approximations are optimistic. Although the NLayeredShell uses a bilinear mapping, the QuadT uses two affine mappings.

Some **Sierra/SD** elements that have condition numbers and can invert do not implement a condition number. Instead, **Sierra/SD** reports a condition number of 1. All **Sierra/SD** HexShell, InfiniteElement, PHex, PTet, PWedge and PmIIsoSolid elements have the optimistic condition number 1. Rigid, Superelement, 1 dimensional and 0 dimensional elements do not have a condition number, and **Sierra/SD** reports the condition number 1.

1.8.7. Problematic Elements and Connectivity

Many problems are caused by “bad” elements. Following are a few issues that come up periodically.

Rotational Invariance can be lost for certain elements such as springs if they are not of zero length. The spring shown in Figure 1-1 is invariant to rotation about the x-axis, but not invariant to rotation about y or z . If we consider an undeformed rotation about the center of the beam along the z axis we would find that $u_y(1) < 0$ and $u_y(2) = -u_y(1)$. If the spring has $KY \neq 0$, then this undeformed rotation results in

strain energy, $E = 2 K_Y u_y^2$. Thus, the rigid body rotation is no longer a zero energy mode.

This is important for a variety of line type elements including spring, Joint2G and gap elements.


Figure 1-1. – Single *Spring* element.



Bad element shape is a major source of problems. For example, we have examined models that have “triangles” where one side is 1/200th the length of the other sides. This produces poor element matrices. In some cases this can destroy the condition of the entire system. Such elements can sometimes be found using the `kdiag` output option described in Section 7.1.50.

Decomposition weakness The broken figure is an issue for trusses (or rods) and some other elements. The truss in the top part of Figure 1-2 is self-sustaining when made of truss elements. However, because truss elements have no rotational stiffness, the decomposed model in the lower part of the figure contains mechanisms. Note that there is no way to decompose the model without introducing such mechanisms.

Figure 1-2. – Truss Decomposition Issues.



Truss elements can be used in **Sierra/SD**. Simulation issues have been traced to poor decompositions in the past. In practice these issues have not come up with `stk_balance` (see 1.1).

Poor Connectivity A structure that has poorly connected regions can be difficult to analyze. If elements have not been properly equivalenced, there can be thousands of zero energy modes in the model. **Sierra/SD** can identify up to a few dozen redundant modes in the best of cases.

Poor Units In models with shell elements, rotational degrees of freedom are active. The rotations scale like the displacements divided by the length. An unfortunate choice of length unit will cause ill conditioning. In theory a unit of length that is much too large could lead to inaccurate rotations. And a unit of length that is much too small could lead to inaccurate displacements. This issue has not been investigated. A smaller linear solver tolerance would generally mitigate the issue.

1.9. *Over-determined Constraints and Loss of Rigid Body Modes*

Building models with `tied surfaces` or other multi-point constraints (MPCs) may involve more than one iteration. Computing eigenvalues is a common way to check a model. Gap Removal is a light-weight solution method [3.38](#) enabling users to obtain information about their tied surfaces earlier in the process, before having to do all the work necessary to successfully set up, submit, and run the massively parallel simulation of the modes.

Solving eigenvalue problems has issues of its own. The eigenvalue decomposition of a model with no prescribed boundary conditions would generally be expected to produce six rigid body modes. A variety of internal constraint issues can cause loss of these rigid body modes. Such artificial constraint of a model can substantially degrade the quality of the overall solution. **Sierra/SD** has several internal correction algorithm and warnings to detect and repair issues, such as `GDSW con_tolerance` described in [Section 2.5](#). However, care should be taken to minimize the complexity of constraint equations. Generally having different types of constraints using the same nodes can be potentially problematic. The following are the types of constraint combinations are known to be particularly problematic:

- Intersection of tied data or contact constraints where one node is tied to two or more faces.
- Chains or cycles tied/contact constraints. This could involve a node on side 'A' of an interface being tied to faces on side 'B' and simultaneously nodes on side 'B' being tied to the faces of side 'A'. This is especially problematic in self contact type situations (a tied surface folding over on itself.)
- Intersection of tied/contact constraints and rigid set constraints where a tied node of the node-face constraint is also part of the rigid set.
- Manually specified MPCs that form long chains, cycles, or extensive intersections.
- Intersection of complex manually specified MPCs with any other type of constraint such as rigid set or tied data.

2. GENERAL COMMANDS

No one seems to know what makes a command general. That is why the explanation is missing here.

2.1. *Input deck format*

The input is a finite element mesh (in **Exodus** format) paired with an input deck setting both the materials used in the model and what to simulate. Once mesh and materials are set, any type of analysis can be performed. An input deck (or text input file) consists of sections. The order of the sections does not matter, nor does the order of the directives within a section. Commands are case-insensitive.

The Sierra tools related to **Sierra/SD** include the Cubit geometry and mesh generation toolkit, the ACCESS/SEACAS tool suite for manipulating exodus files and the Nasgen program for convert NASTRAN files to both the exodus file and the text input file.

Traditionally the sections of an input file are ordered so that directives that are most likely to change are near the top. The solution section selects an analysis type. The file section sets the **Exodus** file name. An **Exodus** file for a structural model typically contains many element blocks. The input file must have a block section for each element block. And a material section specifies each of the materials used in one or more element blocks.

Typically, the input file has an extension of **.inp**, although any extension is permitted. If the **.inp** extension is used, **Sierra/SD** may be invoked on the input without specifying the extension.

The input file is logically separated into sections. Each section begins with a keyword (**Solution** , **block** , etc), and ends with the reserved word **end** . Words within a section are separated with “white space” consisting of tabs, spaces, and line feeds. Comments are permitted anywhere within the file, and follow the C++ convention, i.e. a comment begins with the comment character (i.e. **//**) and ends with the end of the line. ¹

Comments

Several options are available for a comment specifier. These are listed in Table 2-4. For any string used to specify a comment, all characters following the comment string are skipped.

¹To be safe, define comments as one of the allowable comment characters (i.e. **//**) followed by a space.

Table 2-4. – Comment String Options.

String	Descriptor
//	C++ style comments.
/*...*/	C++ style block comments.
#	Hash character. Used in scripts and Sierra input.
\$	Dollar sign. Aprepro/Seacas default.

Meta-characters

Sierra/SD supports the use of several meta-characters in the input file. They are listed below, along with what they represent.

%P The number of processors for the run.

%B The base file name. For example, if the input file is `example.inp`, %B will be replaced with **example**.

%T The **Sierra/SD** start time. This is the same time that is echoed to the screen at the beginning of the run, but in a shorter and more file-friendly format. For example, if the **Sierra/SD** start time was **Tue Sep 24, 2019, 16:05:41**, %T will be replaced with **2019_09_24_16_05_41**.

Skipping Sections

Occasionally an entire section may need to be commented out. This may be done using “//” on each line of the section, or surrounding the section with the block comment characters “/*” and “*/”. A third way to comment out an entire section is to begin it with double “\$\$” characters. In the following, block 1 is commented out, and block 4 is active.

```
$$ BLOCK 1           // this section skipped
  material 1
END

BLOCK 4              // this section valid
  material 1
END
```

Except for special cases such as file names, the input file is case-insensitive. Either the single quote ' or the double quote " may be used throughout the input to keep multiple words together (e.g. a name or title). Quotes may be nested by using both single and double quotes, as in either 'a string with "embedded" quotes' or "a string with 'embedded' quotes".

Preprocessing with Aprepro

The Algebraic PREPROcessor, **aprepro**, in the ACCESS/SEACAS tool suite can be run either stand alone, or as part of the analysis. The use of Aprepro is enabled by default, but can be disabled with the **-noaprepro** command line flag. Aprepro can be used for a variety of purposes.

1. Define variables on the command line. This is especially useful for automated runs such as optimization and uncertainty quantification.
2. Simplify input by allowing algebraic expressions, e.g. $Y = \{ 4 * 3 \}$.
3. Automatically include text of other files.
4. Manage various systems of units, e.g. $Y = \{ 10 * \text{psi} \}$.

For details on Aprepro in general, and for stand alone documentation, please refer to the Seacas documentation.⁶² All the rules for command line substitution apply to the built-in capability. Definition of command line variables during analysis requires specification of a command line argument, **-define**, as used in the example.

```
sierra salinas --define "Eval=10E6 NuVal=0.30" -i example.inp
```

In this example, the text “Eval” in the input file, **example.inp** will be replaced with “10E6”. Likewise, “NuVal” will be replaced by 0.30.

Including Files

The input parser supports nested includes using Aprepro. Files may be included to any depth.

```
{include("english_materials")}
```

The **include** command may occur anywhere on the line (though for readability and consistency we recommend that it be the start of the line). Case sensitivity will be preserved.

Input Summary

Summarizing, a minimum of two files are needed to run **Sierra/SD**, namely, a text input file, e.g. **example.inp**, and an **Exodus** input file,⁵⁵ e.g. **example.exo**, which contains the finite element model. The finite element model is specified in **example.inp** as the geometry file (see Section 2.2).

Each of the **Sierra/SD** input sections is described in the following section.

Integer List An integer list may be required as a parameter for a number of keywords. The list is of a format similar to that of MATLAB. A simple list such as “1,2,3,4” is possible. One may also provide a sequence such as “1:4” which is completely equivalent to the previous example. A *step* value may also be provided, as in “2:2:20”. The second term between the colons is the step. For this example, we list all even values between (and including) 2 and 20. Such combinations can also be combined, as in “1,2,3:2:7,11,13,17,19”.

It is recommended that if such lists have spaces they be placed in (single or double) quotes, i.e. '1 2' is preferred, but 1 2 without quotations is also acceptable.

2.2. *Input mesh geometry file*

Disk files names are specified in the **file** section. The parameters for the **file** section are,

Option	Description
geometry_file	Indicates which Exodus file to use
user subroutine file	File containing user subroutines, currently relevant only to user mount element

2.2.1. **Geometry_file**

The geometry file is used for input of the mesh geometry including the nodes, elements, connectivity and attributes. It is typically a binary **Exodus** file.

2.2.1.1. Multiple processors In a multiprocessor environment, the file name is determined by appending the “dot qualified” processor number and processor id to the geometry file specification.² For example, if the user specifies,

```
geometry_file='temp1/example.par'
```

There are 4 processors, then the following files will be opened.

```
temp1/example.par.4.0
temp1/example.par.4.1
temp1/example.par.4.2
temp1/example.par.4.3
```

²In other words, the user specifies the path name of the first parallel file, but omits the processor count information. This method permits specification of the file name independent of the number of processors used.

In this example the input file is not in the current working directory.

Specifying the **Exodus** mesh input file in a way that does not depend on the number of input files (or MPI ranks) saves time in the long run. If the input file is `example.exo` and the partitioner is configured to generate the files 'example.exo.4.0',... then the same **Sierra/SD** text input file works in serial and in parallel.

2.2.1.2. Single processor On a single processor, the file is not “spread”, and the full file path is provided. A representative serial `file` section follows.

```
file
    geometry_file exampleg.exo
end
```

Note:

- If the file name is not included in quotes, it will be converted to lower case.
- A single processor run, even using mpi protocol, will not append the number of processors and processor ID to the file name.
- Section 1.1 shows the steps involved in the parallel execution of **Sierra/SD**.

2.2.2. Exodus Database Naming Conventions

When the mesh file has an **Exodus** format, there are three basic conventions that apply to user input for various command lines. The conventions concern side sets (surfaces), node sets and blocks.

First, for a mesh file with the **Exodus** format, the **Exodus** side set is referenced as a surface. In SIERRA, a surface consists of element faces plus all the nodes and edges associated with these faces. A surface definition can be used not only to select a group of faces but also to select a group of edges or a group of nodes that are associated with those faces. For nodal boundary conditions that use the surface specification, all the nodes associated with the faces on a specific surface will have this boundary condition applied to them.

A group of elements can also be used to select other mesh entities. In SIERRA, a block consists of elements plus the associated faces, edges, and nodes. The block and surface concepts are similar in that both have associated derived quantities.

Blocks, sidesets, or nodesets can be grouped together into **assemblies**.⁶⁴ However, an assembly must contain entities of the same "type", where "type" is either a block, sideset, nodeset, or another assembly. An assembly cannot contain itself, either directly or indirectly. Currently, not all tools support assemblies. One way to create them is with `io_modify`. In addition, assemblies can be created directly in the **Sierra/SD** input deck using the **ASSEMBLY** section (section 2.2.3).

The specification for a surface identifier can have the form `surface_integerid`, where `integerid` is the integer identifier for the surface. Surfaces can also be identified by the `integerid` alone. For example, if the side set identifier is 125, the surface could be referred to as `surface_125` or simply 125. It is also possible to name a surface in some mesh generation programs, and that name can be used in the input file. The specification for element blocks and nodesets are defined similarly (`block_integerid/integerid/blockName` and `nodelist_integerid/integerid/nodesetName`). Assemblies can also be specified similarly (`assembly_integerid/assemblyName`), but specification of assemblies by `integerid` alone is not allowed. Assemblies can be substituted anywhere in the input deck where a list of blocks, sidesets, or nodesets is expected. If the assembly itself contains sub-assemblies, the set of unique leaf blocks, sidesets, or nodesets will be used, depending on context.

2.2.3. ASSEMBLY section

Assemblies can be created using the **ASSEMBLY** section of the **Sierra/SD** input deck as shown below (syntax 2.1). The assemblies created in this way can then be used elsewhere in the **Sierra/SD** input deck. If an integer is given as the assembly name, then the integer will be used as the assembly identifier and the assembly name will be `assembly_integerid`, where `integerid` is the assembly identifier. Note that a **ASSEMBLY** section can only refer to assemblies that either already exist in the input mesh or were defined in a prior **ASSEMBLY** section. If an assembly already exists, either on the input mesh or from a previous **ASSEMBLY** section, defining it again using a new **ASSEMBLY** section will result in modifying the assembly so that it contains the union of the entities in the original and the new instance. Additionally, a warning will be printed in the **Sierra/SD** results file about modifying an existing assembly.

```
BEGIN ASSEMBLY <string>
// exactly one of the following:
    block = <list(block)>
    block = all
    sideset = <list(sideset)>
    sideset = all
    nodeset = <list(nodeset)>
    nodeset = all
END
```

Syntax 2.1. ASSEMBLY example

2.2.4. Exodus Naming Limitations

Sierra/SD supports the use of mesh names in the input deck and lists of mesh objects. As a result, it becomes ambiguous which keywords are mesh names and which are intended as

input syntax. For example, a nodeset named "fixed" creates ambiguity when the user specifies `nodeset 1 fixed`. This could mean that nodeset 1 is fixed in space, or it could be a list of nodesets "1" and "fixed".

To resolve this ambiguity, **Sierra/SD** maintains a list of reserved keywords which may not appear in the mesh. This list has over 400 entries, and changes with each release. The rules of this list are simple; any valid syntax which appears adjacent to a nodeset list may not appear in the mesh's nodesets. The same rule applies to sidesets and blocks. Since assemblies are valid in all three contexts, they inherit the restrictions of all three.

There are no keywords in the **Sierra/SD** input deck which begin or end with an underscore. Therefore, `"_fixed_"` is a valid name for any mesh object in any context. Similarly, `"subAssy_5_front"` is obscure enough to be safe from this restriction.

Pure integer values are not allowed as mesh names for the same reasons. This is not a restriction on names like `"Block_5"`, but blocks named `"5"` are not allowed.

A warning is issued if reserved characters are found in the mesh names. These warnings are not fatal, but having reserved characters in the mesh names severely limits our ability to parse them in the input deck. The reserved character list includes comment characters (`//`, `/`, `*`, `*`, `/`, `#`, and `$`), and special-use characters (`=`, `'`, `"`, and `:`), which have extra meaning in **Sierra/SD** input decks.

Since the list of reserved keywords is based on the allowable syntax in **Sierra/SD**, it will change based on the current version of the code. The current list of reserved keywords for the version of the code you're using can be printed to screen at any time using the `"-keywords"` command-line option. Additionally, the ambiguous naming checks can be changed to a warning or ignored entirely using the `reserved_keywords` parameter (section 2.3).

2.2.5. Additional Comments About Output

A text log or *results* file can be written for the run. Details of the contents of the results file are controlled in the **echo** section (see Section 7.7). The results file name is determined by the name of the input file, and will be in the same directory as the input text file, regardless of whether **Sierra/SD** is being executed in serial or parallel. However, if executing in parallel, using the **subdomains** option in the **echo** section allows control of the number of *results* files. For example, if running on 100 processors, up to 100 result files may be output. Using **subdomains** `"0:2"` will only output three files, from subdomains 0, 1, and 2. The default is to output a *results* file only for processor zero. The results file name uses the base name of the input, with an extension of `.rslt`. In a parallel computation, the results file names use the base name of the input file, followed by an underscore and the processor number, then followed by the `.rslt` extension.

For calculations in which geometry based output requests are included (see Section 7), an output **Exodus** file will be created. An **Exodus** file is a binary file (NetCDF). It contains

the original geometry and the requested output variables. The output **Exodus** file name is determined from the geometry file name. The base name of the output is taken from the geometry file by inserting the text “-out” before the file name extension. The output **Exodus** file will be written to the same directory where the geometry file is stored. If executing **Sierra/SD** on a parallel machine, the **Exodus** output files should be written to the raid disks for reasonable performance.

2.3. *Parameters*

This *optional* section provides a way to input parameters that are independent of the solution method or solver. Only one **parameters** section is recognized in each file. The parameters and their meanings are listed below and in Table 2-5 and 2-6. For reference, Table 2-7 are parameters occasionally used by developers, but not recommended for common use.

wtMass This variable multiplies all mass and density on the input, and divides out the results on the output. It is provided primarily for the English system of units where the natural units of mass are units of force. For example, the density of steel is $0.283\text{ lbs}/\text{in}^3$, but “lbs” includes the units of

$$g = 386.4\text{ in}/\text{s}^2.$$

Using a value of WtMass of 0.00259 (1/386.4), density can be entered as 0.283, the outputs will be in pounds, but the calculations will be performed using the correct mass units. **Sierra/SD**, like most finite element codes, does not manage the *units* of the analysis. The selection of a consistent set of units is left to the analyst. For example, if the analyst uses the SI system (Kg, m, s) the units of pressure must be Pascals. Frequencies are reported in Hz. For micromachines these units are awkward. It may be better to use units of grams, millimeters and microseconds. The analyst must ensure that all material properties and loads are converted to these units.

Some examples of useful units are shown in Table 2-8.

NegEigen Unconstrained structures have zero energy modes which may evaluate to small negative numbers due to machine round off. The eigenvalues and associated frequencies are reported as negative numbers in the results files. However, many post processing tools (such as *Ensign*) require non-negative frequencies. By default, **Sierra/SD** converts all negative eigenvalues to near zero values in the output **Exodus** files ¹. To retain the negative eigenvalues in the output file, select parameter **NegEigen**.

¹Because many post processing tools are written for transient dynamics, they expect monotonically increasing, positive values for the time. Since eigenvalues are written in the time columns of the output file, they are converted to be monotonically increasing, positive values. Note that the numerically computed eigen frequencies are also stored as global variables in the file

Keyword	Type	Default	Description
AllowExodusAttributes	<i>bool</i>	true	Seek undefined attributes on the mesh
AllowExodusDistFacts	<i>bool</i>	true	Apply distribution factors from the mesh
AllowInvalidExodusParts	<i>bool</i>	false	Ignore error due to requesting an invalid Exodus part (e.g. sideset/block)
ComplexStress	<i>yes/no</i>	no	output complex stress in FRF solutions
condition_limit	<i>Real</i>	1e6	element quality output control
constraint_correction	<i>yes/no</i>	no	orthogonalize constraints to RBMS
defaultSpecificHeat	<i>Real</i>	none	material specific heat
DoInitialMassSolve	<i>bool</i>	true	transients: perform initial mass solve
eig_tol	<i>Real</i>	auto	Eigenvalue tolerance
eigen_norm	<i>string</i>	mass	“visualization” or “unit”
energy_time_step	<i>int</i>	1	input of energy data
energy_exo_var	<i>string</i>	TEMP	Exodus energy variable name
FilterRbmLoad	<i>noFiltering</i> <i>allStructural</i> <i>rotationOnly</i>	noFiltering	rigid body components of loads to filter
ignore_gap_inversion	<i>bool</i>	false	Ignore element quality changes due to gap and overlap removal
Info	<i>int</i>	1	screen output control
MatrixFloor	<i>Real</i>	0	control of matrix fill
MaxmpcEntries	<i>int</i>	10,000	maximum # entries in any mpc
MaxResidual	<i>Real</i>	1	maximum residual for eigen
MortarMethod	<i>string</i>	dual	dual or standard mortar method
MFile_format	<i>string</i>	sparse_function	control output format for MATLAB
NegEigen	<i>none</i>		negative eigenvalue flag
nonlinear_default	<i>yes/no</i>	yes	nonlinear element blocks
num_rigid_mode	<i>int</i>	0	number of system rigid body modes
output_sideset_data	<i>bool</i>	false	Output element values on all sidesets
RandomNumberGenerator	<i>string</i>	“rand”	or “test”
RbmTolerance	<i>Real</i>	1e-10	tolerance for rigid body zero
RemoveRedundancy	<i>yes/no</i>	yes	filter constraints
reorder_Rbar	<i>none</i>		constraint reordering flag
RequireMatchedBlocks	<i>string</i>	ignore	“ignore”, “warn”, or “fatal”
SkipmpcTouch	<i>none</i>		control of MPCS
syntax_checking	<i>string</i>	fatal	“ignore”, “warn”, or “fatal”
reserved_keywords	<i>string</i>	fatal	“ignore”, “warn”, or “fatal”
restart_consistent ...	<i>string</i>	fatal	“ignore”, “warn”, or “fatal”
TangentMethod	<i>string</i>	element	method of computing tangent
WtMass	<i>Real</i>	1	Mass multiplier

Table 2-5. – Available keywords in the Parameters section.

Keyword	Type	Default	Description
nodesets_with_disp	<i>list</i>	1 TEMP new	nodesets with prescribed displacements
thermal_time_step	<i>int</i>		input of thermal data
thermal_exo_var	<i>string</i>		Exodus temperature variable name
mpmd_transfer_version	<i>old/new</i>		In core transfer coupling library
mpmd_transfer_type	<i>fuego/copy</i>		Which code transfer is
mpmd_transfer_sidesets	<i>sparc/interpolation</i> <i>string_list</i> <i>string_list</i>		occurring from Surfaces on which transfer is occurring

Table 2-6. – Available keywords in the Parameters section related to code coupling and handoff.

Keyword	Arg	Default	Description
OutputInitialTime	<i>boolean</i>	false	write output at t=0
SharedAcousticStiffness	yes/no	no	share stiffness across processors
linear_solver_warn_factor	<i>Real</i>	10	GDSW warnings
linear_solver_bailout_factor	<i>Real</i>	100	GDSW fatal errors

Table 2-7. – Developer keywords in the Parameters section.

Developer keywords in the Parameters section. These keywords may be used, but their use is discouraged, and they are not fully tested.

Table 2-8. – Some useful combinations of units.

length	mass	time	WtMass	density	force	modulus	internal mass
<i>m</i>	Kg	sec	1	Kg/m^3	<i>N</i>	N/m^2 or Pa	Kg
ft	slug	sec	1	slug/ft^3	<i>lbf</i>	lb/ft^2	slug
ft	<i>lbm</i>	sec	1/32.2	lbm/ft^3	<i>lbf</i>	lb/ft^2	slug
in	<i>lbm</i>	sec	1/386.4	lbm/in^3	<i>lbf</i>	psi	lbm/386.4
<i>mm</i>	μg	μs	1	Kg/m^3	<i>N</i>	MN/m^2 or MPa	μg
<i>mm</i>	g	sec	1	g/mm^3	μN	N/m^2 or Pa	gram
<i>mm</i>	mg	sec	1/1000	g/cm^3	μN	N/m^2 or Pa	gram

output_sideset_data This option turns on sideset output for any requested element output. Currently, sideset output is only enabled for (centroid) volumetric stress and strain (sections 7.1.16 and 7.1.18).

eig_tol This is the tolerance used in the **eigen** solution method for eigenvalues. The default is machine precision. This parameter can sometimes be loosened to aid convergence. See Section 1.8.3 for details.

nonlinear_default In nonlinear transient dynamics or nonlinear statics, computing the fully nonlinear response of all the elements in the mesh can be computationally expensive. In some cases it is unnecessary. For example, for a simulation that only involves Joint2g elements and solid (3D) elements, the analyst may determine that the nonlinear effects of the solid elements are negligible. In such cases, it is advantageous to be able to control the nonlinear response of elements block-by-block. And there is a block-level parameter described in Section 4.7.2 that selects the optional nonlinearities for specific blocks. Instead of entering this parameter for each block, **nonlinear_default** sets the default for all blocks. If **no**, then all blocks default to linear behavior, unless specified otherwise in the BLOCK section. If **yes**, then all elements default to nonlinear behavior. Note that the block-level flags override the **nonlinear_default** keyword. There are two possible cases for this keyword.

nonlinear_default=no All elements default to linear behavior.

nonlinear_default=yes All elements default to nonlinear behavior.

As noted in Section 4.7.2, there are limitations for using linear materials in nonlinear analysis.

TangentMethod The tangent stiffness matrix may be used in a full Newton update in nonlinear statics or transient dynamics (see Sections 3.25 and 3.26). By default, each of the elements can compute its own tangent stiffness matrix. There are cases (particularly when elements are under development) when it is better to use a tangent matrix computed from finite difference methods. There are three possible values for this keyword.

TangentMethod=element The standard element method.

TangentMethod=difference Use finite difference.

TangentMethod=compare Use the standard method, but also compute the matrix by the difference method. Output of the difference of every element matrix in the model will be sent to the results file. ²

info

²In parallel solutions the results file is written only for the first processor unless the “subdomains” option is specified in the echo section (7.7).

GDSW!prt_debug Option **info** selects diagnostic information for standard out. There are four different levels. Each level increasingly allows more output to standard out. The GDSW option “prt_debug” overrides “info” for GDSW output. The four levels of control are:

- 0. Silent** – output warnings and std error to the screen (untested)
- 1. Normal** – output the data most analysts would use (untested)
- 2. Detailed** – Convergence, solution addressing issues (unimplemented, tested).
- 3. Debug** – Silent, normal, detailed and diagnostic information (tested).

Info is regularly used for debugging eigenvalue problems. Option **prt_debug** is defined in Section 2.5.6. Option **prt_debug** causes GDSW to write load balance information to a file named **subdomainData.dat**. Its contents are described in Section 2.5.3

Example of usage:

```
Parameters
  info=0
End
```

This sets the “info” control level to Silent.

syntax_checking **Sierra/SD** has the ability to check input files for syntax and spelling errors. This option controls this behavior. By default, a violation is printed to the screen and execution is terminated. If the user wishes, violations can be printed while execution continues, or the checking for violations can be disabled completely.

The three levels of control are:

- ignore** silently ignores those entries.
- warn** provides a warning for those entries.
- fatal** stops the analysis. This is the default.

reserved_keywords **Sierra/SD** has the ability to check exodus entities for potentially ambiguous or confusing names. This option controls this behavior. By default, a violation is printed to the screen and execution is terminated. If the user wishes, violations can be printed while execution continues, or the checking for violations can be disabled completely. Note that this option will only control the behavior when checking ambiguous names: we will always issue a warning if certain reserved characters are found within a name. For more information, see section 2.2.4.

The three levels of control are:

- ignore** silently ignores those entries.
- warn** provides a warning for those entries.
- fatal** stops the analysis. This is the default.

restart_consistency_checking by default **Sierra/SD** will check that the restart data file is consistent with the analysis mesh. These checks include confirming identical node ids, element id, node coordinates, and element connectivity. By default, if the restart file is inconsistent **Sierra/SD** will output a fatal error and terminate execution. Generally such restart inconsistencies indicate a serious problem. For example trying to map the restart data from the wrong analysis onto a new mesh will produce nonsensical behavior. Use **restart_consistency_checking** to convert these checks to a warning or skip them.

The three levels of control are:

ignore silently ignores restart consistency issues.

warn provides a warning for restart consistency issues.

fatal stops the analysis for restart consistency issues. This is the default.

SkipmpcTouch **Sierra/SD** uses a unique method of determining an active degree of freedom set. Unlike codes like NASTRAN which use an automatic single point constraint method, **Sierra/SD** loops through all elements and activates only degrees of freedom that are required for elements. Multipoint constraints pose a particular problem because some codes (like NASTRAN) may include multipoint constraints to unused degrees of freedom. Since these are eliminated with the autospc, this poses no problem to these codes, but may confuse **Sierra/SD** significantly. On the other hand, usually degrees of freedom associated with MPCs should be included in the active set, and leaving them out can produce errors.

As a stopgap measure, we provide the parameter **SkipmpcTouch** . If this parameter is set, no degrees of freedom will be activated through multipoint constraints.

condition_limit Element quality checks are important for evaluating the effectiveness of the mesh. By default, elements with moderately bad topology are reported. However, sometimes there are so many of these warnings, that the bad elements may get missed. The **condition_limit** parameter permits user control of the reporting. Setting this parameter to a larger number will eliminate message from marginal elements. Element checking can also be disabled (see the **ElemQualChecks** parameter in the **outputs** Section 7.1.12). The default value is 1e6.

*Note that the **condition_limit** parameter is ignored if the **ElemQualChecks** option has been disabled in the **output** section (7.1.12).*

reorder_Rbar An RBar is a type of rigid element. This option reorders all Rbar elements to minimize the number of them connected to a single node. Having many Rbar elements connected to the same node results in a problematic matrix sparsity pattern which can reduce efficiency. Reducing the number of connections can shorten run time.

If redundant Rbars are present (i.e. connections forming a cycle), they are removed.

Specify `reorder_Rbar yes` or `reorder_Rbar no`. The default value is `yes`, which means `Rbars` will be reordered.

thermal_time_step For thermal analysis solution procedures (i.e. statics or transient dynamics with a **thermal_load** body load), or for any solution procedure that uses temperature dependent material properties, the temperature distribution of the structure must be read in from the **Exodus** file. Typically, the input **Exodus** files in this case would be the output files from a thermal analysis, and thus would contain the necessary temperature data. Since such an analysis could contain several time steps of temperature data, the parameter **thermal_time_step** allows the analyst to select which set of temperature data is to be read into **Sierra/SD**. The following gives an example.

```
Parameters
  thermal_time_step 10
End
```

In this case the user would be requesting that the temperature data corresponding to the 10th time step be read into **Sierra/SD**.

energy_time_step This variable is identical to the “`thermal_time_step`” above, but applies to cases where the energy density is input and must be converted to a temperature. Either energy density or temperature can be input, but not both.

thermal_exo_var If a material specifies a coefficient of thermal expansion and a reference temperature, then the corresponding thermal strain is computed, and the corresponding thermal load 6.3.8 is applied. Otherwise, a material property may be an arbitrary function of temperature. Temperature dependent materials 4.4.6 are supported. Temperature dependent viscoelastic materials 4.3 are also supported, using the corresponding input syntax. In either case to read the temperature from the input **Exodus** file, the name of the temperature field has to be specified. And that is what **thermal_exo_var** is for. The default variable name is 'TEMP'. In the following example the name is changed to temperature.

```
Parameters
  thermal_exo_var "temperature"
End
```

energy_exo_var This variable is identical to the “`thermal_exo_var`” above, but applies to cases where the energy density is input and must be converted to a temperature. Either energy density or temperature can be input, but not both. The only difference is that the energy density will be divided by the specific heat to arrive at the temperature.

```
Parameters
  energy_exo_var "EDEP"
End
```

mpmd_transfer_version Defines which version of the runtime application coupling library to use. Currently, the **old** option is used for coupling to Nemo, and the **new** option is used to couple to SPARC or Fuego. See Section 6.2.6.

mpmd_transfer_type Defines which code **Sierra/SD** is coupling to during the run. The **fuego** are **copy** options are equivalent and one of them must be used when coupling to Fuego. The **sparc** and **interpolation** are equivalent and one of them must be used when coupling to SPARC. See Section 6.2.6.

mpmd_transfer_sidesets Defines the sidesets on which coupling are occurring. When data is mapped from one code to the other, the transfer will only occur on these sidesets. See Section 6.2.6.

FilterRbmLoad Establishes a filter for rigid body components of the input load. The options are described in the table. It defaults to unfiltered. The parameter may need to be combined with the RbmTolerance and solver parameters. The **FilterRbmLoad** parameter is only supported for transient and static solution cases. For other solution cases this parameter will have no effect on the solution.

During rigid body mode filtering a net force in a rigid body mode will be counter-balanced by a force with a distribution defined by the rigid body mode shape times the mass matrix. For example to counter-balance a net force in X direction effectively a gravity load would be applied to the body where the sum of the forces of that gravity load is equal and opposite to the net force to be balanced. For rigid body filtering a well-defined mass matrix is required for both transient and static solution.

The rigid body load filtering should only be used of a model that has exactly the six standard rigid body modes.

Option	Description
NoFiltering	skip all RBM filtering for the load
AllStructural	apply filtering to all 6 structural RBM
RotationOnly	apply filtering to rigid body rotation only

See Section 6.3.19 for more details about the use of this option.

RbmTolerance Rigid body filters depend upon accurate rigid body modes. The application checks the matrix product of the stiffness matrix to ensure that these vectors are in the null space of the stiffness matrix. If any of the requested vectors are not in the null space, the application terminates. The default is $1e-10$. RbmTolerance provides user control of the threshold for that error. It depends on the stiffness matrix, K , and a rigid body vector, ϕ_r ,

$$tolerance = \frac{\|K\Phi_r\|_2}{\|K_d\|_\infty \|\Phi_r\|_2}$$

MatrixFloor Primarily a debugging option. The nearly zero terms in a matrix can be removed using this parameter. Values below this floor are eliminated from the

matrix. This can reduce fill, but if used improperly too much of the matrix can be affected. It can be important when running on different platforms, where round off can affect the matrix fill, and make it difficult to compare solutions. This is a relative value, so 1.0E-6 would remove terms in the matrix that are a million times less than the largest term. Default is zero.

defaultSpecificHeat The specific heat is used to convert energy to temperature with the following equation, $\tilde{E} = C_v \Delta T$, see Section 4.4.8 for details. The specific heat is set on a material by material basis. A default value for the specific heat can be set using the keyword **defaultSpecificHeat** which can then be overridden by defining the specific heat for a material. A fatal error is given if the specific heat is not defined for a block of material containing energy input and the **defaultSpecificHeat** keyword can be used to avoid this.

MaxmpcEntries Soft limit on the number of mpc entries in any single multipoint constraint. Normally the default will be sufficient, but large RBE3 type entries may exceed this in rare cases. The limit is there to avoid errors reading the input, and because such large constraints can consume memory.

eigen_norm Eigenvectors may be arbitrarily normalized. Three common approaches are listed in Table 2-9. All methods retain orthogonality of the eigenvectors, but the normalization differs. The default, mass normalization, is most commonly used as it ensures that the inner products of eigenvectors with the mass matrix is identity. However, this normalization is not well suited to output visualization. The “visualization” normalization mimics what is automatically done in MSC/Patran, and should provide a reasonable visualization without rescaling each mode. In “visualization” normalization, the maximum *translational* displacement is normalized to be less than 10 percent of the maximum model extent, while also insuring that the model rotation remains below 1 radian. Unit normalization ensures that the largest value of the eigenvector is one. ¹ A global variable, *EigenVectScale*, provides the scale factor by which the mode was scaled.

Method	Algorithm	Comment
Mass	$\phi_i^T M \phi_i = 1$	Default. Simplifies numerics
Visualization	$\max(\phi_i) = (\text{model size})/10$	Simplifies visualization
Unit	$\max(\phi_i) = 1$	

Table 2-9. – Eigenvector Normalization Methods.

constraint_correction Ensure that each multipoint constraint generated is orthogonal to all rigid body modes. This is useful for lofted surfaces. If the surfaces are tied as if

¹The “unit” method of normalization computes $\max(\phi)$, which is computed *only* on translational displacement degrees of freedom. Note also that only displacements are renormalized. No effort is made to renormalize element variables such as strains, stresses or energies. Thus, if these are requested in an eigen analysis, they will not be consistent with the renormalized eigenvectors, but will retain mass normalized values.

they were coincident, the constraints are incorrect, and eliminate some or all the rigid body modes as worked out in subsection Orthogonality of MPC to Rigid Body Vectors in Section Linear Algebra Issues of the Theory Manual.

Parameters

```
Constraint_Correction=yes
End
```

MFile_Format Most of our matrix data can be written as MATLAB readable files. By default, these are written as sparse matrices, as functions. Other formats are also available. The “full” format does not use the sparse methods and is thus compatible with *Octave* or other tools. Alternatively, the “3column” format can be used. In this format, the file is loaded using the MATLAB “load” command. The data is then converted to a sparse matrix using the MATLAB “sparse” command. The “3column” format may be significantly faster in some cases, but it does require more user interaction. Figure 2-3 compares a simple example for the three formats. In all cases, the matrix symmetry is the same. A fourth format, “CSV”, is also available for compatibility with other external tools. ²

Sparse_Function	Full	3column
function s=Kssr() s=[1 1 0.11 1 2 0.12 2 2 0.22]; s=sparse(s(:,1),s(:,2),s(:,3));	function s=Kssr() s=zeros(2,2); s(1,1)=0.11; s(1,2)=0.12; s(2,2)=0.22;	1 1 0.11 1 2 0.12 2 2 0.22

Figure 2-3. – Example MFile Format Results.

RemoveRedundancy RemoveRedundancy affects node/face constraints created by 'Tied Data' or 'Contact Definition'. Redundant constraints cause most solvers to fail. Redundant constraints are often introduced when two surface pairs are tied next to each other, but there are a variety of sources for these redundancies. Exact redundancies are always automatically eliminated, but that is often not sufficient. This parameter removes constraints when a node is tied to more than one face or if the node shows up both as a node of a node-face constraint and attached to a face of a different node-face constraint. By default, it is “true”.

RandomNumberGenerator The default random number generator, “rand”, is the standard generator available from system libraries. It should be the best random

²Note that the CSV format should be readable by Microsoft Excel, but there are often limits on the number of columns that can be read.

number generator in terms of the quality implementation. In a few cases the analyst may want a more repeatable random number generator, that is independent of the platform. The “test” random number generator can be used in this case. It is *not* recommended for general use, and the statistics of the generator are not well-established.

MortarMethods Two mortar methods are available in **Sierra/SD**: standard and dual (see⁴⁸). By default, the dual method is selected as it is almost always more efficient in memory use.

ComplexStress Most often, analysts do not want output of stress variables in frequency response function analysis. Such output is complex, and huge volumes can be generated. Selecting “ComplexStress=yes”, along with “stress” in the echo section permits output of this data. The default is “ComplexStress=no”.

num_rigid_mode Is used to signal to the linear solver that the system is singular and that the singularity is associated with structural and/or acoustic rigid body modes. This is used, for example, in the solution of statics problems without any essential boundary conditions or frequency response analysis with the modal acceleration method. Where possible, other methods should be used to eliminate the singularity. For example, in eigen analysis a negative shift is recommended. Currently, allowed values for this parameter are 1 (acoustic mode only), 6 (structural modes only), or 7 (structural and acoustic modes). We also note that when using the **FilterRbmLoad** parameter, it is necessary to specify **num_rigid_mode** to correspond to the number of rigid body modes that will be filtered. For example, if **FilterRbmLoad** was set to **AllStructural**, then **num_rigid_mode** should be set to 6.

ignore_gap_inversion Initial overlap removal is another name for gap removal. Gap removal may change element quality. A fatal error occurs if element quality gets much worse. To ignore this poor element quality set the **ignore_gap_inversion** parameter to true. The gap removal solution case 3.38 making debugging easier.

DoInitialMassSolve Determining the initial acceleration is necessary³⁵ for quadratic convergence. By default, a consistent initial acceleration is found through an initial mass solve, see Section 3.32.1.2. Some combinations of MPCs can result in a singular mass matrix, leading to errors in the initial mass solve. To skip the initial mass solve use the command,

```
DoInitialMassSolve=false
```

OutputInitialTime In transient simulations, the output can be written at time t=0, prior to the first time step by using the command

```
OutputInitialTime=true
```

The default is set to false. This will not write the initial time step if reading from a restart file because the time is not 0.

linear_solver_warn_factor The GDSW linear solver approximates the solution u of a linear system, $Ku = f$ so that the predicted norm of the k -th residual $f - Ku_k$ is less than a user specified tolerance times the initial residual norm $f - Au_o$ (the default is $1.e - 6$). However, the true residual may be larger, and a careful analyst must check the true residuals. To automate the task of the checking, a warning message is written if the true residual is more than *linear_solver_warn_factor* times the predicted residual. The default value is 10.0.

`linear_solver_warn_factor=10.0`

linear_solver_bailout_factor The GDSW linear solver approximates the solution u of a linear system, $Au = f$ so that the predicted norm of the k -th residual $f - Au_k$ is less than a user specified tolerance times the initial residual norm $f - Au_o$ (the default *solver_tol* is $1.e - 6$). However, the true residual may be larger. A fatal error is thrown if the true residual is more than *linear_solver_bailout_factor* times the predicted residual. The default value is 100. For instance, in an ill-conditioned problem as the *solver_tol* is decreased to improve resolution, one may choose to increase the bailout factor to mitigate the risk of a fatal error.

`linear_solver_bailout_factor=100.`

UseAnalystNodeMap This feature is no longer supported.

RequireMatchedBlocks Sierra/SD always requires that each **Exodus** block have a corresponding entry in the `.inp` file. This parameter controls behavior if there are blocks defined in the `.inp` file that do not appear in the mesh.

ignore silently ignores those entries.

warn provides a warning for those entries. This is the default.

fatal stops the analysis.

nodesets_with_disp This parameter, when used in conjunction with the **nUpdateConstraints** transient option (see section 3.32.1), enables specifying prescribed displacements over a subset of all nodes via nodeset output.



`nodesets_with_disp` is currently BETA release.
Enable with the “`--beta`” command-line option.

2.4. Solution Options

The options described in Table 2-10 and in the following paragraphs are part of the **Solution** section in the input file. None of the keywords are required. Note that in multicas solutions most of these parameters may be applied separately within the subcase (see Section 3.2.1).

Table 2-10. – Sierra/SD Solution Options.

Parameter	Type	Default	Description
restart	none read write auto and restart output.	none	Controls both restart input from argument can be given
to change the restart input file can be given to read a restart	name. Addition- ally, file that was written on are being used	An optional an op- tional a differ- ent for the	num_procs argument number of processors than current run.
lumped lumped_consistent			Lump the mass matrix special mass matrix for reducing dispersion er- ror
solver	<string> for example GDSW.	auto Default is	Select solver to use for case, to pick a solver based on model size and structure
ConstraintMethod scattering	Lagrange penalty false true	Lagrange false	Method for applying MPCs Separate incident and scattered acoustic loads. symmetric structural acoustics
symmetrize... ..._struc_acous		formulation	

2.4.1. Flush

The parameter **flush** controls how often the **Exodus** output file buffers should be flushed. Flushing the output ensures that all the data that has written to the file buffers is also written to the disk. This parameter also controls the frequency of output of restart information if requested. Too frequent buffer flushes can affect performance. However, in a transient run, data integrity on the disk can only be assured if the buffers are flushed. A **flush** value of -1 will not flush the **Exodus** output file buffer until the run completes. The default value is to flush the buffers every 50 time steps.

2.4.2. Restart – option

Option **restart** controls restart file // 48:202 processing. Restart files permit the solution to be saved for later use. Only a limited capability is provided, but it is intended to meet most of the typical needs for structural dynamics. Note that except for eigen restarts, the restart files are independent of the **Exodus** output, but the restart options may significantly affect the **Exodus** outputs. Eigen restarts directly use the **Exodus** output. Application of restarts in specific sections is detailed in the following paragraphs. See Figure 2-4 for details on eigen restart.

There are four values associated with this option.

none indicates that restart files will be ignored. They will be neither read, nor written. Existing restart files will not be altered in any way. Restart=none is the default selection if no restart options are entered in the solution block.

read indicates that existing restart files will be read, but no output restart files will be written. If the restart files do not exist, a fatal error will result.

write indicates that existing restart files will be ignored, but restart files will be written.

auto is a combination of read and write. However, unlike read, the existence of previous restart files is optional, i.e. there will be no error message if there are no existing restart files. Invalid restart files will produce a warning, but not a fatal error.

The restart **read** and **auto** options can take an optional **from** clause to define the file containing the restart data. Examples are given below.

Restarts are designed to ensure accuracy of the solution. However, restarts in **Sierra/SD** are not transparent in the sense that there will be small differences in two solutions to a problem when one solution involves a restart. Restarts may also have an expense. For example, GDSW uses an acceleration technique where the values of previous solutions are used as a starting place for new solves. The information associated with previous solutions is not stored in the file. Also, restarts discard Lagrange multipliers.

For transient dynamics, the state of the machine at the most recent time step is recorded. To avoid problems with corruption of a database, the three nodal vectors (disp, velocity, acceleration) and applicable element data are recorded at each time step, but on alternate locations in the file. If previous **Exodus** files exist, they will be appended. Data is written at the same interval as the **Exodus** output.

When restarting a multicas solution, the current time is used to determine which case the restart will begin. For example, assume the following solution block is defined.

```
Solution
  case one
    transient
    restart=auto
    time_step 1e-6
    nsteps 200
  case two
    transient
    restart=auto
    time_step 1e-5
    nsteps 300
End
```

If restarting at time=1E-4, case “one” has a final time value of $T_F = T_0 + 200 \cdot 1E-6 = 2E-4$, assuming $T_0=0$. Since time < T_F , case ‘one’ will restart the solution. If restarting

at `time=2E-4`, case 'one' will not perform any calculations and restart will begin with case 'two'.

An important note for restarts is care must be taken to ensure the restart files are given the expected names. By default, restart file names are defined based on the geometry file and the case name. The following example uses the default behavior and will write a restart file with the name `mesh-one.rst_trans`.

```
Solution
  case one
    transient
    restart=write
    time_step 1e-6
    nsteps 200
End
File
  geometry_file = 'mesh.g'
End
Output
End
```

The default restart (read) input deck would then be as follows. The mesh name and case name will be used to define the restart input file name `mesh-one.rst_trans`.

```
Solution
  case one
    transient
    restart=read
    time_step 1e-6
    nsteps 200
End
File
  geometry_file = 'mesh.g'
End
```

The default restart file names work fine if you are running a model multiple times in one directory, with the same number of processors, say to computing additional transient time steps or eigenvalues with each restart. Alternatively a **from** clause can be used with the **read** or **auto** restart options to explicitly define the restart input filename.

It is possible to read a restart file that is decomposed for a different number of processors than are being used for the run by setting **num_procs**. Restart files will always be written using the same base name and paths as the **outputs** files.

For the following example assume the analysis run is submitted on 4 processors. The combination of the restart **from** and **num_procs** clauses will explicitly set the name and path from which to read the parallel restart files to

'../origRun/myRestartFile.rst_trans.16.00', '../origRun/myRestartFile.rst_trans.16.01',
'../origRun/myRestartFile.rst_trans.16.02', etc.

As the output **database name** is set **Sierra/SD** will write restart to the file name
'outputDir/myNewFile.rst_trans.4.0', 'outputDir/myNewFile.rst_trans.4.1',
'outputDir/myNewFile.rst_trans.4.2', and 'outputDir/myNewFile.rst_trans.4.3'. Restart
and output files will always be written decomposed on the same number of processors that
are being used by the analysis.

```
Solution
  case newCase
    transient
    restart=read from '../origRun/myRestartFile.rst_trans' num_procs 16
    time_step 1e-6
    nsteps 200
End
File
  geometry_file = 'new_mesh.g'
End
Output
  database name = 'outputDir/myNewFile.exo'
End
```

The restart options **from** and **num_procs** can be used to execute the common use case of N
to M restart. For example it can be efficient to execute an expensive eigen decomposition
once on hundreds of processors then read that eigen decomposition on a handful of
processors to do many subsequent modal transient or modal random vibration analyses.

2.4.2.1. Restart solution support Restarts are not supported in all solutions types.
They are supported for the following.

- Eigen
- transient
- NLtransient
- modaltransient
- Q EVP (with Anasazi and sa_eigen methods)

Note that none of the modal solutions except modaltransient support restart. Typically
most of the computation time for these solutions is in the eigen analysis. It is
recommended that the multicas solution be used, with a restart in the eigen analysis
portion of the solution.

Table 2-11 illustrates the current supported restart options for transient analysis. As shown here, one can restart a **modaltransient** analysis from a **transient** analysis, and vice versa.

Integrator	NLtransient	transient	modaltransient
NLtransient	Tested	?	?
transient	Tested	Tested	Tested
modaltransient	Tested	Tested	Tested

Table 2-11. – Supported restart capabilities for transient integrators in **Sierra/SD**.

If the Q EVP solution restarts at all, then every requested mode must be provided. Therefore, you cannot ask for more modes than are in the restart file and have **Sierra/SD** calculate only the additional modes. You can however save previously calculated modes for use in follow on solution methods.

All restart files are in the exodusII format. This permits evaluation and manipulation of the data using standard tools. Details of the file names and formats used for restart are included in Table 2-12.

The eigen solution case avoids the complexity of writing a special restart file. Instead the **Exodus** output is reused. Restart does require that originally displacements were output to the **Exodus** file. However, the use of the standard file format facilitates error diagnosis. Reuse of the **Exodus** output does impact the options available as follows.

NONE ensures that no data will be read from existing **Exodus** files. However, *output, written in the normal fashion, will replace existing files*. As a consequence, the results of the previous analysis will be modified.

auto Data is read from available input files. After a run, data is written to the appropriate files. They will be updated to include additional calculations.

READ Read from the output of the previous modal or transient analysis. In the case of a transient analysis, however, no output is written. Also, if no previous transient analysis wrote a restart file, then there will be nothing to read in this case.

WRITE No restart files are read. Results are output after analysis. For modal analyses, this option is effectively identical to **none**. For transient analyses, however, either **write** or **auto** is required to write a file that can be used to restart.

Figure 2-4. – Notes on Eigen Restart.

The **auto** option makes it difficult to know whether or not restart occurred. To clarify if restart is being used, it is best to explicitly define the restart **read**, **write**, **none** options.

2.4.2.2. Restarts with virtual nodes and elements Currently, Sierra-SD creates virtual nodes and elements when a model has features like Tied Joints, infinite elements, and

Solution	filename	Details
eigen	<i>example-out.exo</i>	Use the standard Exodus output for restart. Displacements <i>must</i> be written or no restart is possible. Other variables (such as strain energy) may also be written.
qevp	<i>example-out.exo</i>	Uses standard Exodus output for restart. No additional modes may be computed, but the superposition is possible.
transient NLtransient modaltransient	<i>example-out.rst_trans.exo</i>	The two most recent time steps are written. They are only written at the “flush” interval.

Table 2-12. – Restart file format and names. Details of the restart files for a file named “example.exo” for various solutions.

superelements. These virtual nodes and elements are created internally during execution of the code, and result in the analysis having more degrees of freedom than are specified in the **Exodus** file.

When the restart files described in Table 2-12 are written to disk, these virtual elements and nodes will also be written, resulting in **Exodus** files that have more nodes and elements than were in the original **Exodus** file. The restart **Exodus** output files can be visualized in Enight and the data can be processed in MATLAB. Thus, the restart files provide this additional benefit of being able to post-process the virtual nodes and elements. For example, one may wish to visualize the Rbars created for a Tied Joint in Enight. This can be done by visualizing the restart files.

2.4.3. Solver

As **Sierra/SD** evolves, various solvers are available for computation of the solution. Each solver brings with it different capabilities and sometimes unwanted features. Currently, available solvers are listed in the following.

auto Use the best known solver. Generally this is recommended. The matrix of solvers versus solution types is messy, and generally the best solution will be found by using this option. For example, there is no need to change the solver as you move from serial to parallel solutions.

GDSW The Generalized Dryja, Smith, Widlund (GDSW) solver is based on a domain decomposition preconditioner which combines overlapping Schwarz and iterative substructuring concepts.²⁵ The GDSW solver is well suited to solving problems with large numbers of constraint equations. It has also been observed to be competitive with other parallel solvers, even for problems with only a small number of

constraints. The GDSW solver is currently under development and supported by the Sierra-SD team. The most recent development efforts for the GDSW solver have been focused on implementation and testing of a new Helmholtz solver for direct frequency response analysis.

Generally no user input is required for specification of a solver. Indeed, up to version 1.0.5 of **Sierra/SD**, only one solver was ever available at any time (i.e. we built separate executables if another solver was desired). Usually the specification can be left off, or specified as **auto**. If a solver is requested and unavailable, a warning will be issued, and **auto** will be selected.

The solver may be specified as a default (above the **case** keywords as detailed in Section 3.2.1), or it may be individually specified within the case framework. The default value is **auto**. In the example shown below GDSW will be used for the eigen analysis, and the **auto** selection for the for transient dynamics and direct frequency response. If “input_summary” is specified in the “ECHO” section (see Section 7.7) then the solver information will be echoed to the results file.

```
SOLUTION
  solver=auto
  case eig
    eigen nmodes=50
    solver=gds
  case nlt
    NLtransient
    (other parameters)
  case frf
    directFRF
END
```

2.4.3.1. Limitation There is a limitation of restarts that the number of processors (i.e. MPI ranks) used to restart must be the same as used originally. There is one exception however. For the eigen solution method, it is possible to restart on a single processor. Beforehand, the distributed output file must be joined into a single file. Similarly for a transient solution, it is possible to restart in serial. In this case both the distributed output files and the distributed restart file must be joined into single files.

2.4.4. Lumped – option

To use a lumped mass matrix for all solution cases add option **lumped** anywhere in the **Solution** section. Most off diagonal terms are set to zero. The diagonal terms are increased to conserve mass.

The drilling degrees of freedom associated with beams and shells can generate spurious modes when they are lumped. As a consequence, **Sierra/SD** does not fully lump these degrees of freedom. They are lumped in the element coordinate frame, but transforming the mass matrix to the physical coordinates results in a 3×3 matrix.

2.4.4.1. Computational Acoustics In acoustics simulations with hexahedrons, in theory a special mass matrix minimizes dispersion error.²⁰ Adding `lumped_consistent` to the **Solution** selects the special mass matrix. Many have conjectured that the special mass matrix enhances simulations of linear elasticity. The CFL number $c \, dt/dx$ depends on the sound speed, c , the element size, dx , and the time step size, dt . A study⁵¹ of the `lumped_consistent` mass matrix reached the following conclusions.

- In the most accurate simulations, time step size is scaled with mesh size to maintain a constant CFL number.
- Simulation accuracy depends on the temporal (e.g. Newmark) as well as the spatial (finite element) discretization.
- A `lumped_consistent` mass matrix can enhance accuracy.
- With Newmark Beta time integration, the most accurate simulations used a consistent mass matrix and a CFL near 7/10.

2.4.5. Constraintmethod – option

The **ConstraintMethod** option is defined in the **Solution** section to indicate how multipoint constraints (MPC) will be applied. MPCs applied using **Lagrange** multipliers. Inverse problems sometimes use an experimental penalty method.

2.4.6. Scattering – option

For some acoustics and structural acoustics problems, it is advantageous to define the loads in terms of an incident pressure instead of a total pressure. The solutions for the scattered pressures follow the same differential equations as those of the total pressures. It may be necessary to combine the incident and scattered terms to compute a total pressure. A review is presented in subsubsection Acoustic Scattering subsection Acoustic and Structural Acoustic Boundary Conditions section Acoustics and Structural Acoustics of the Theory Manual Note that the **scattering** keyword applies to *all* loads in the solution case. It is nonsensical to mix scattering pressure inputs with total pressure inputs.

Scattering solutions require this keyword in the solution block. In addition, loads should be applied properly in the LOADS block. The user must apply a load to *both* the structural

and the acoustic side of a wet surface. This is typically done using a function tailored for that purpose.¹

2.4.6.1. symmetrize_struc_acous – option By default, coupled structural acoustic discretizations are symmetric. This is accomplished⁵² by scaling the acoustic equation by a -1 . In some cases (e.g. infinite elements) scaling is impossible, and cases the code internally reverts to the nonsymmetric formulation. The nonsymmetric formulation is always available by setting **symmetrize_struc_acous** to false. In a multcase solution **symmetrize_struc_acous** can vary from case to case. The pros and cons of two approaches have never been studied.

2.5. GDSW

GDSW is the workhorse for parallel solutions. It is the default linear solver on two or more processors. Many **Sierra/SD** features require that GDSW be the linear solver. In this manual, “the solver” is the linear solver. This section describes the GDSW parameters. Table 2-13 describes the basic solver parameters. Parameters for advanced usage are given in Tables 2-14 and Table 2-15 Report problems using the GDSW solver with the default solver parameters to the Sierra help system at **sierra-help@sandia.gov**.

To use a non-default linear solver (i.e. GDSW in serial) in the solution section set **solver=GDSW**. Non-default parameters are set in an optional GDSW section. The GDSW section applied to each solution case. The first subsection of this Section 2.5.1 presents a way to specify different GDSW parameters for each solution case. Subsection 2.5.1 describes additional parameters. Subsection 2.5.2 describes GDSW output. The corresponding parameters are summarized in Table 2-16. Subsection 2.5.3 describes some techniques for improving performance and reliability. Subsection 2.5.4 reviews theory related to solution accuracy. The GDSW Helmholtz solver is a relatively new capability. Subsection 2.5.5 presented the parameters for the parallel frequency response linear solver. The parameters are summarized in Table 2-19. Report any problems using the new Helmholtz solver to **sierra-help@sandia.gov**.

solver_tol It is important to control the accuracy of the solution. For all our linear solvers, *solver_tol* is the requested accuracy of the computed solution as measured by the relative residual error. In other words, the 2-norm of the residual vector for the computed solution divided by the 2-norm of the right-hand side force vector should be no greater than *solver_tol*.

orthog Convergence may be accelerated by storing and recycling search directions from previous solves. This feature requires additional memory, but may significantly reduce iterations. It is possible for the application of these vectors to be unstable.

¹the “plane_wave”, “planar_step_wave” and “shock_wave” functions compute both appropriate pressures on the structure, and normal velocities on the acoustic medium. See sections 2.8.10, 2.8.12 and 2.8.14.

Table 2-13. – GDSW Section Options. (Basic)

Variable	Values	Default	Description
max_iter	<i>integer</i>	1000	maximum number of iterations
solver_tol	<i>real</i>	1e-6	relative residual convergence tolerance
overlap	<i>integer</i>	2	number of layers of overlapping nodes for preconditioner
orthog	<i>integer</i>	1000	number of stored search directions used to accelerate solver convergence (also see num_vectors_keep)
prt_summary	<i>integer</i>	3	output flag: 0 - no summary 1 - summary 3 - detailed summary

This rare event is avoided by trying orthog=0. The default value of orthog is 1000. For the Helmholtz linear solver the corresponding parameter is orthogH.

krylov_method A variety of Krylov iterative methods are available as shown in Table 2-14. The default should work fine in most instances. If convergence problems arise, switching to classic right preconditioned GMRES is recommended. (krylov_method = gmresClassic) without the use of any stored search directions (orthog = 0).

default_solver The subdomain sparse direct linear solver is specified either by name (Esmond, Pardiso, NoPivot) or by the corresponding integer (1,3,6).

num_rigid_mode note: see **parameters** , Section 2.3. This keyword must not appear in the GDSW solver section. It is read from the Parameters section.

constrain_rbms Tells GDSW to numerically constrain rigid body modes of a structure in order to do a static solution of a free floating structure. Must be used in conjunction with the **FilterRBMLoad** parameter to ensure no net rigid body load on the constrained structure. The rigid body modes associated with **x** , **y** , **z** , **rotx** , **roty** , **rotz** can be selectively constrained, as can **p** the pressure degree of freedom in acoustic analysis. See Section 6.3.19 for more details on this use case.

max_numterm_C1 Constraints for the GDSW solver are classified by two types:

Type 1: simple constraints like those applied by an Rbar, tied contact, or rigid surfaces.

Type 2: more complex, averaging constraints like those in an RBE3.

Type 1 constraints have few terms, and Type 2 constraints have many terms in each constraint equation. Solution of problems with Type 2 constraints using Type 1 methods is possible and desirable if they are small enough, but the memory requirements could be prohibitive if the number of terms N in any constraint equation is too large. Specifically, storage of a dense matrix with at least N^2 terms

Table 2-14. – GDSW Section Options 1. (Advanced)

Variable	Values	Default	Description
krylov_method	<i>integer</i>	1	0-pcg: preconditioned conjugate gradients 1-GMRES: right preconditioned GMRES (generalized conjugate residual version) 2-lgmres: left preconditioned GMRES 3-flexgmres: flexible right precond GMRES 4-flexgmres2: variant of flexgmres 5-gmresClassic: right preconditioned
default_solver	<i>integer</i>	1	1-direct: Esmond Ng’s sparse direct solver 3 - Pardiso for Pardiso sparse direct solver (Intel MKL only), 6-NoPivot: Clark’s sparse direct solver note: see parameters , Section 2.3.
num_rigid_mode			Enables solution of static
constrain_rbms	<i>x y z rotx rotyz rotz p</i>		system with rigid body modes
max_numterm_C1	<i>integer</i>	250	max num terms for Type 1 constraints
coarse_option	<i>integer</i>	1	0 - additive coarse correction, 1 - multiplicative coarse correction
SC_option	<i>integer</i>	1	0-no/1-yes: eliminate subdomain interior unknowns using static condensation
weight_option	<i>integer</i>	2	1 - to not use weighted residuals for overlapping subdomain problems
coarse_size	<i>string</i>	auto	coarse space reduction option auto: automatic selection small: use reduced coarse space large: use larger coarse space
reorder_method	<i>string</i>	metis_edge minimum_degree, none	metis, metis_edge, rcm,
num_GS_steps	<i>integer</i>	1	number of Gram-Schmidt orthogonalization steps for stored search directions
con_tolerance	<i>real</i>	2.5e-9	singularity tolerance for processing constraints
con_row_tolerance	<i>real</i>	0.1	pivoting tolerance for processing constraints
scale_option		0	0 - no scaling in factorizations 1 - use scaling in factorizations
diag_scaling	<i>string</i>	none	none - no scaling of operator matrix diagonal - symmetric diagonal scaling
PTAP_solver	<i>integer</i>	1	solver for conjugate gradient matrix 0-diag: diagonal (in exact arithmetic) 1-full: full $\Phi^T A \Phi$ matrix
bailout	<i>keyword</i>		If keyword is found, ignore errors
coarsening_ratio	<i>integer</i>	1000	coarsening ratio for multilevel solver

Table 2-15. – GDSW Section Options 2. (Advanced)

Variable	Values	Default	Description
minCoarseLevels	<i>integer</i>	1	min number of coarse levels (for testing only)
maxCoarseLevels	<i>integer</i>	1	max number of coarse levels
maxCoarseSize	<i>integer</i>	3000	max size for coarsest problem
enforceActual...	<i>integer</i>	0	0-no/1-yes
...Residual			
graphPartitioner	<i>integer</i>	0	graph partitioner for multilevel solver 0-Parmetis, 1 PHG in Zoltan
num_vectors_keep	<i>integer</i>	900	related to orthog; see discussion below
stag_tol	<i>real</i>	0.01	Used to detect stagnation
dd_solver_output_file	<i>string</i>	dd_solver.dat	Output name for domain decomposition solver diagnostic file
krylov_solver_output_file	<i>string</i>	krylov_solver.dat	Output name for Krylov solver diagnostic file
useParallelDirectSolver	<i>integer</i>	0	0-no/1-yes
useParallelCoarseSolver	<i>integer</i>	0	0-no/1-yes
numCoarseProcs	<i>integer</i>	all available	number of processors for the coarse problem
useSuperLUDist	<i>integer</i>	0	0-no/1-yes: whether to use SuperLU-Dist
numProcRowSuperLUDist	<i>integer</i>		number of rows in SuperLU-Dist process grid
identify_low_quality_elements	bool	false	identify poorly shaped elements for special attention by solver
max_element_condition	double	Infinity	condition above which to consider element low quality

would likely be required. The parameter *max_numterm_C1* specifies the maximum number of terms that can appear in a Type 1 constraint following a constraint pre-processing step. Constraints with more than *max_numterm_C1* terms are then considered to be Type 2. The algorithm used to enforce Type 2 constraints in the preconditioner is generally not as efficient as the one for Type 1 constraints. If feasible, avoiding Type 2 constraints by increasing the *max_numterm_C1* in conjunction with a multiplicative coarse correction (*coarse_option=1*) will lead to a more efficient solution.

Avoiding Type 2 constraints generally reduces run times. To eliminate Type 2 constraints set the GDSW option *max_numterm_C1* to a sufficiently large number. The value that GDSW is using for *max_numterm_C1* is in the *dd_solver.dat* file. It is called *maxNumTermsForType1Constraints*. Also, in the *dd_solver.dat* file note the value of *maxNumNonZeros in Tran Matrix*. To eliminate Type 2 constraints *max_numterm_C1* must be larger than the value of *maxNumNonZeros in Tran Matrix*.

coarse_size Is used to specify a reduction strategy for the coarse problem size. There is no need to consider this parameter for problems run on fewer than a few hundred processors. However, as the number of processors (subdomains) becomes large, solving the coarse problem can become a bottleneck. The default (*auto*) automatically selects to use the *small* coarse space only if the number of processors exceeds 1000. Specifying a *small* rather than a *large* coarse space can often reduce the amount of memory needed by the solver.

reorder_method Allows one to specify a reordering method for a sparse direct solver. Currently, it is only available for *default_solver = direct* (see Table 2-14).

num_GS_steps is the number of orthogonalization steps of the stored search directions. Sometimes increasing it from its default to 2 is helpful. I know of no case of a value larger than 2 being necessary. It turns out that **num_GS_steps** does not apply to *gmresClassic*.

con_tolerance The GDSW solver uses a sparse LU decomposition algorithm to process the constraint equations. This involves choosing pivot rows for numerical stability (much like Gaussian elimination with partial pivoting). A constraint equation is deemed linearly dependent if the magnitude of its pivot is less than *con_tolerance*. The *con_tolerance* is the minimum acceptable ratio of the constraint pivot to the maximum matrix pivot. The *con_tolerance* can be viewed as a dimensionless parameter associated with numerical round off. The number of numerically redundant constraints in a model will typically be reduced as the **con_tolerance** is increased. Excessively high *con_tolerance* values could start to remove legitimate constraints.

Messages of the form,

```
min/max pivot for constraint factorization = some number
You may want to consider increasing the con_tolerance
parameter in the GDSW solver block.
```

are issued if the ratio of magnitudes of the smallest to largest pivots is less than 0.01. This provides a recommendation to carefully examine the constraints in the model for any potential problems. Additionally, a message will be written if there exists constraints on the cusp of being removed as redundant. Again presence of such nearly redundant constraints should be investigated closely for correctness.

Redundant or nearly redundant constraints can be generated in a variety of circumstances. For example a closed loop of Rbars, Rbars between nodes of a rigid set, or overlapping rigid sets. Although the SD constraint filtering algorithm is designed to detect and handle such redundancies it is strongly recommended that redundant constraint be avoided in model definition as much as possible. See Section 1.9 for more details.

scale_option There are presently two options for matrix scaling in the GDSW solver. Including **scale_option yes** or, equivalently, **scale_option 1** in the GDSW solver block will apply symmetric diagonal scaling to all matrices prior to them being passed to Esmond Ng's sparse direct solver. Notice for parallel runs that both the subdomain matrices and the coarse problem matrix will be scaled. In exact arithmetic, this option should have no effect on the number of iterations for each solve of a parallel run.

diag_scaling Including **diag_scaling diagonal** in the solver block will apply symmetric diagonal scaling to the original operator matrix and is not tied to a specific sparse direct solver. In contrast to the **scale_option** parameter, this parameter changes the number of iterations for each solve of a parallel run since GDSW is solving the scaled problem $DADy = Db$ to a specified relative residual tolerance rather than the original problem $Ax = b$ (note substitution of $x = Dy$, where D is a diagonal scaling matrix) for that same tolerance.

coarsening_ratio Is a target ratio between the number of subdomains prior to and after coarsening by the multilevel solver. For example, if there are originally 8000 subdomains (processors) and **coarsening_ratio** is chosen as 100, then the number of subdomains after coarsening will be 80.

maxCoarseLevels Is the maximum number of coarse levels allowed by the multilevel solver. For a standard 2-level method this parameter has a value of 1.

maxCoarseSize Is the largest size for the coarsest problem allowed before another level of coarsening is made. The solver parameter **maxCoarseLevels** takes precedence over **maxCoarseSize**.

enforceActualResidual This option is used to enable strict enforcement of the solver tolerance. If needed, iterative refinement steps are taken within the solver in an attempt to satisfy the convergence criterion. It should be noted that it may not be possible to drive the relative residual to values no greater than the solver tolerance. This can happen when the problem is poorly conditioned and/or the requested solver tolerance is too small. An error message is issued and the analysis stops if the specified solver tolerance cannot be achieved.

graphPartitioner Specifies which graph partitioning software to use when coarsening the subdomains.

num_vectors_keep applies to the classic version of GMRES (krylov_method gmresClassic). The parameter orthog (default 1000) controls the number of stored search directions. We store search directions in order to make the linear solver faster. Generally more is better but not always. The point to understand is which search directions are stored. The first 1000 search directions are stored. On later solves, the first 900 are saved and recycled. 100 search directions from the current solve are used. The number 900 is the default value of num_vectors_keep. Sometimes the solution has changed significantly and none of the old search help the solver. num_vectors_keep= 0 tells GDSW to never recycle any of the stored search directions between solves.

num_GS_step_gmres is the parameter for the Krylov method gmresClassic corresponding to the num_GS_steps for other Krylov methods.

useParallelDirectSolver Whether to use parallel sparse direct solver for the whole problem. Best suited to directfrf solution case.

useParallelCoarseSolver Whether to use a parallel sparse direct solver for the coarse problem. This can be helpful whenever the coarse problem is large enough that its single-processor solution becomes a bottleneck.

numCoarseProcs Number of processors to use if useParallelCoarseSolver = yes.

useSuperLUDist Whether to use SuperLU-Dist parallel sparse direct solver even if Cluster Pardiso is available. Applies to both the full problem (useParallelDirectSolver and coarse problem (useParallelCoarseSolver) cases.

numProcRowSuperLUDist SuperLU-Dist uses a rectangular grid of processes. This option allows the user to override the default choice, which is to make the number of rows about equal to the number of columns. The parameter is ignored unless it divides the total number of processors. It has been suggested that for large problems, the number of columns should be larger than the number of rows, but we have not yet tested cases where this choice proved to be beneficial.

The generalized conjugate residual (GCR) version of right preconditioned GMRES is the default Krylov method. GCR has the limitation that orthog must be at least as large as max_iter. On the other hand with gmresClassic max_iter and orthog may be specified independently.

2.5.1. Options

The GDSW section specifies GDSW linear solver parameters for all the solution cases. Different solution cases may call for different parameters. For example, an eigenvalue problem may use a shift parameter to eliminate rigid body modes, but a statics analysis cannot shift. Linear solver parameters for an individual solution case are set using the

`solver_options` keyword in the solution case, and adding the corresponding `solver_options` section.

The `solver_options` sections allow multiple definitions of the parameter. Each “solver_options” section may be called out from a separate case in a multicase solution block as illustrated in Input input 2.1. A “solver_options” section applies *only* to the GDSW solver.

Note that cases using a `solver_options` section will ignore any “global” options defined in the GDSW section, and will instead use the internal GDSW defaults for any options that aren’t explicitly defined.

```
Solution
  case preload
    statics
    solver_options for_preload
    load=10
  case eig
    eigen
    nmodes=100
    solver_options for_eig
End
Solver_options for_preload
  solver_tol = 1e-4
  constrain_RBMs="x y z"
End
Solver_options for_eig
  solver_tol = 1e-8
End
```

Input 2.1. Multiple Solver Options Example

2.5.2. Diagnostics

There is a variety of solver diagnostic information output to the files `dd_solver.dat` and `krylov_solver.dat`. Here we describe six different measurements of solution accuracy.

Note: the default names of the dd and Krylov solver diagnostic files (“dd_solver.dat” and “krylov_solver.dat”) can be overridden by the `dd_solver_output_file` and `krylov_solver_output_file` options in the **GDSW** and `solver_options` sections.

actual final residual Let’s say we want to solve the linear system $Ax = b$ for the vector of unknowns x given the right-hand side vector b . If x_a is the approximate solution from the solver, then the actual final residual is $\|b - Ax_a\|$, where $\|b\|$ denotes the 2-norm of b . The actual final residual is reported in `krylov_solver.dat` after each solve.

actual relative residual This is $\|b - Ax_a\|/\|b\|$ and shows up in column 6 of `dd_solver.dat`. In other words, this is the actual final residual divided by the norm of right-hand side vector b .

recursive final residual During GMRES or CG iterations, we can recursively calculate $b - Ax_a$ without having to directly calculate Ax_a (for efficiency reasons). In exact arithmetic the actual and recursive final residuals are identical, but in practice they can be different because of round off errors. If the effects of round off are not too big, then the actual and recursive final residuals should be close. The recursive final residual is reported after each solve in `krylov_solver.dat`.

recursive relative residual This is the recursive final residual divided by $\|b\|$ and shows up in column 5 of `dd_solver.dat`.

constraint error residual The constraint equations for a problem can be expressed as $Cx = 0$, where C is the constraint matrix. The constraint error residual is a normalized measure of Cx_a and should be small relative to 1 if the approximate solution x_a satisfies the constraints well. This residual is reported after each solve in `krylov_solver.dat`.

equilibrium error For problems with constraint equations, we solve the linear system $\begin{bmatrix} A & C^T \\ C & 0 \end{bmatrix} \begin{Bmatrix} x \\ \lambda \end{Bmatrix} = \begin{Bmatrix} b \\ 0 \end{Bmatrix}$, where x is the vector of unknowns, λ is the vector of Lagrange multipliers, and C^T denotes the transpose of the constraint matrix C . For the approximate solution vector x_a , the equilibrium error is $\|b - Ax_a - C^T \lambda\|$ and is reported after each solve in `krylov_solver.dat`. In the absence of round off errors, the equilibrium error and the actual final residual should be identical.

All these residual measures may be more than is usually of interest, but they can provide valuable information for cases when solver convergence is an issue.

Picking a suitable solver tolerance for GDSW or any other iterative solvers requires close attention. If the solver tolerance is too high, then simulation results may not have sufficient accuracy. Likewise, if the solver tolerance is too low, then more analysis time may be spent obtaining a needlessly over-accurate solution. Solving a problem with two or more values for the solver tolerance can be useful to help avoid unnecessarily accurate solves and to also ensure that the solves are accurate enough. For example, let's say you do a modal analysis with a solver tolerance of $1e-4$ and $1e-5$. If you don't see a concerning change in the modal frequencies, then either choice for the solver tolerance is probably fine for the modal frequencies themselves. If, however, the modal solution is used as the basis for a subsequent analysis such as modal transient, then it is recommended that the effects of solver tolerance on the final results also be considered. Similar comments regarding the choice of a solver tolerance also hold for other solution cases such as direct transient analysis.

We could provide users with estimates for how small a solver tolerance is needed to guarantee a certain measure of accuracy, but these estimates would usually be way too pessimistic to be of any practical value.

Additional details and troubleshooting strategies for the GDSW solver can be found in §2.5.3 and documentation available on the compsim.sandia.gov website. Relevant documentation includes GDSW 101 and the GDSW Solver Tutorial. Solver strategies for dealing with poor mesh decompositions caused by the presence of constraints equations or multiple physics (i.e. structural-acoustics problems) are described in the GDSW Solver Tutorial. These include rebalancing algorithms internal to the solver that can be accessed using GDSW solver parameters. We hope this will provide a useful interim solution for challenging problems prior to the deployment of alternative decomposition tools that effectively address these issues prior to the solution phase. Input 2.2 provides recommended options for minimum memory use.

```
GDSW
  overlap=0
  max_iter=50    // set to minimum required for a solution
  krylov_method = gmresClassic
  orthog=0
  precision_option_0=single
  precision_option_coarse=single
END
```

Input 2.2. Minimum Memory Recommended Options. These options, while not usually optimal for speed, may use the lowest memory.

Table 2-16. – GDSW Section Options. (Supplemental Output)

Variable	Values	Default	Description
prt_coarse	<i>integer</i>	0	0-no/1-yes: print coarse matrix
prt_constraint	<i>integer</i>	0	0-no/1-yes: print constraint matrix
prt_memory	<i>integer</i>	0	0-no/1-yes: print memory information
prt_timing	<i>integer</i>	0	0-no/1-yes: print timing information
prt_interior	<i>integer</i>	0	0-no/1-yes: print interior matrices
prt_overlap	<i>integer</i>	0	0-no/1-yes: print overlap matrices
write_orthog_data	<i>integer</i>	0	0-no/1-yes: write orthogonalization data to file

2.5.3. Troubleshooting

Ideally, the linear solver should always return a solution that satisfies the requested accuracy in terms of the relative residual tolerance (10^{-6} by default). This subsection provides some troubleshooting guidelines for situations when this is not the case. Additional information is available at compsim.sandia.gov in the Sierra/SD online documentation *GDSW Solver Tutorial* and *GDSW 101*. If problems persist, please submit

a Sierra help ticket or reach out to a member of the Sierra/SD development team for assistance.

- **Default Solver Parameters:** In case of solver convergence problems, it is recommended that one first verify that the default solver parameters do not work. A notable exception is for problem types like structural acoustics where a smaller solver tolerance than the default may be needed to obtain accurate structural and acoustic responses.
- **Negative Shift for Modal Analyses:** A common mistake is to not specify a negative shift for modal analyses of structures with no essential boundary conditions. If one or more rigid body modes are present, then the stiffness matrix will be singular and the solver will likely have problems converging except in special cases (see Singular Solves bullet below). The recommended shift is $-(2\pi f)^2$, where f is an estimate of the natural frequency (in Hz) of the first flexible mode. Of course the modes are not known in advance. The point is that the shift does not need to be an accurate estimate. For modeling an ordinary (larger than a paper clip, smaller than a house) steel component in Imperial units -10^6 almost always works. Caution: specifying a negative shift that is too large in magnitude may help the linear solver, but it can cause the algorithm that solves the eigenvalue problem to either require too many linear solves or not converge at all.
- **Memory Considerations:** The linear solver requires that enough memory be available to store the factorizations of subdomain and coarse space matrices. If the subdomain or coarse matrices are too large, then the memory capacity of the computing resource will be exceeded and a run will fail. There are a number of different ways to address such problems. First, one may request a smaller number of processors per compute node. For example, the CTS-1 machine *eclipse* has 36 cores per compute node, but one may request that only half or even a smaller number of cores be used per node. This has the effect of providing more memory for each subdomain (MPI rank). In the same spirit, another option is to run the problem on a larger number of processors. This will result in smaller subdomain matrices, but the size of the coarse matrix will increase. For large problems, it may be necessary to use a second coarse level to limit the coarse matrix size. Memory resources might also be exceeded by the storage of search directions used by the Krylov method (e.g. GMRES). The default maximum number of iterations is 1000, and reducing this number will result in memory savings. Additional information on memory usage is available in the online documentation *GDSW Memory Use Tutorial*.
- **Convergence Issues:** Due to the limitations of finite precision arithmetic, it may not always be possible to provide a solution which satisfies the specified relative residual tolerance. This is especially the case for poorly conditioned linear systems or unrealistically small solver tolerances. Discussion of this topic continues in the online document *Solver Accuracy Notes*

Increasing the maximum number of iterations is one simple option that can result in successful solves. Other options for improving convergence are described below in

separate bullets. If convergence to the specified tolerance still cannot be achieved, then one option is to increase the relative residual tolerance until the solves are successful. To confirm that the quantities of interest have converged, one or more additional runs with even smaller tolerance are recommended.

- **Subdomain Overlap:** The default value of the `overlap` solver parameter is two. This means that the original subdomains are extended by two layers of elements when determining their overlap. Increasing the overlap often reduces the number of iterations needed for convergence, but the memory requirements increase and each iteration will require more time. Try a value for overlap of 3, 4, or more and see what effects it has. Even larger values of overlap can be used for models with only shell elements due to the reduced demands on the subdomain solver.
- **Static Condensation:** For non direct frequency response problems, the linear solver eliminates subdomain interior residuals at each iteration by default and solves the resulting Schur complement system of equations. For models with higher-order elements (i.e. polynomial degree greater than 2), it is recommended that the `SC_option` parameter be set to `none` since elimination of residuals interior to the subdomain can be problematic (i.e. effects of round off errors can be more pronounced). Similar concerns are present for direct frequency response problems, but the default option for such problems is `none`.
- **Frequency Response Analysis:** The iterative solution of direct frequency response problems can be challenging. The coefficient matrix $K + i\omega C - \omega^2 M$ can be both indefinite and complex depending on the input circular frequency ω and the damping matrix C . A more thorough discussion of trouble shooting the Helmholtz linear solver⁵³ is in the **Sierra/SD** How To manual in section Frequency response linear solver.

A distributed memory sparse direct solver is recommended if the problem is not too large (to fit in memory). A direct solver works well for structures that are not blocky and shell element models. In the GDSW solver block set `useParallelDirectSolver = yes`. If available, this uses the Intel code Cluster Pardiso. Otherwise, SuperLU-Dist is used. If it is desired to use SuperLU-Dist on a platform where Cluster Pardiso is available, the option `useSuperLUDist = yes` must be specified.

If the distributed memory sparse direct solver option is not viable, then convergence may be improved as the number of processors increases. This leads to smaller subdomains and better performance of the coarse part of the GDSW preconditioner. Some users have also seen improved convergence by using the non-default solver parameter settings `orthogH = 0` and `precondUpdateFreq = 1`. If the problem already has a fair amount of damping, it may also help to set `structural_damping = 0`. This manual documents the interface [2.5.5](#).

Fiscal year 2021 planned work includes the capability to solve for multiple subdomains per MPI process, which should improve performance without the need to run on more processors.

- **Constraint Equations:** Certain types of constraint equations like those introduced by RBE3 elements can cause challenges for the linear solver (see earlier discussion of the `max_numterm_C1` solver parameter, Type 1 constraints, and Type 2 constraints). Models with Type 2 constraints are generally harder to solve than those without them. Two lines in the solver output file `dd_solver.dat` relevant to avoiding Type 2 constraints are `maxNumNonZeros in Tran Matrix` and `maxNumTermsForType1Constraints`. The current value of the `max_numterm_C1` solver parameter is reported by `maxNumTermsForType1Constraints`. Type 2 constraints can be avoided by setting `max_numterm_C1` to be no less than `maxNumNonZeros in Tran Matrix`, but in some cases this can lead to excessive memory requirements for the solver.

Another source of potential difficulty is the presence of dependent or nearly dependent constraints. This means that the coefficient matrix $\hat{C}$ in the constraint equations $\hat{C}x = 0$ has either an infinite or large condition number. There are filters inside of Sierra/SD to eliminate linearly dependent constraints, but some may still be passed to the solver. The solver also identifies and eliminates dependent or nearly dependent constraints, but this is not always foolproof. It is advised that an analyst carefully check their model's contact definitions to avoid potential problems in terms of model correctness or for the linear solver. Additional information on constraint equations and their interactions with the linear solver is available in the online documentation *Constraints*.

- **Singular Solves:** The stiffness matrix for a structure with no essential boundary conditions is singular. Thus, static analyses or modal analyses with a zero shift require special considerations. First, any rigid body mode component of the applied loads must be removed (see `FilterRbmLoad` and `num_rigid_mode` in the **Parameters** section). The GDSW solver is able to solve singular systems, but it is important that the rigid body modes actually exist. For example, consider a model with curved surfaces that are connected using node-to-face tied surfaces. If a dependent node is not located on the independent surface, then one or more of the rotational rigid body modes will be lost. If the solver expects there to be six rigid body modes but there are say only three, then solver convergence problems are likely to occur. To avoid this problem, ensure that your model does indeed have the specified rigid body modes.
- **Very Small Solver Tolerances:** As noted earlier, it may not be possible for the solver to provide a solution which satisfies the requested relative residual tolerance. Nevertheless, there are some situations where it may be possible to reduce the actual relative residual below what is possible using default solver parameters. Sierra/SD has a guardrail to error out if the actual relative residual is more than 100 times of that requested (see `linear_solver_bailout_factor` in the **Parameters** section). By increasing this parameter while reducing the solver tolerance, it may be possible to reduce the actual relative residuals for challenging problems like structural acoustics.
- **Scalability:** If you add `prt_debug=1` or `prt_debug=yes` in the GDSW section, then the information needed to measure the load balance is written to a file. The filename is `subdomainData.dat`. A row of data in the `subdomainData` file has 8 columns of

integers. There is one row per subdomain (or MPI rank or core). Typically, the integers in each column are defined as follows.

1. Interior unknowns (i.e. number of unknowns for static condensation).
2. Non-zeros in the factors of the interior matrix.
3. Unknowns on subdomain boundary.
4. Owned unknowns on subdomain boundary.
5. Unknowns in subdomain.
6. Owned unknowns in subdomain.
7. Unknowns in overlapping subdomain.
8. Non-zeros in the factors of the overlap matrix.

In the special case of a diagonal preconditioner (`preconditioner_type=DIAG`), only 2 integers are printed.

1. Unknowns in subdomain.
2. Owned unknowns in subdomain.

Each iteration sometimes accesses each of these matrices a constant number of times. Often times the non-zeros in the factors of the overlap matrix predict the load balance. Larger numbers are generally more important. If the interior unknowns are eliminated, then non-zeros in the factors of the interior matrix will sometimes predict the load balance. The unknowns in a subdomain is governs the load balance of matters related to `orthog`. As `orthog` or `orthogH` increases, the influence of the unknowns per subdomain increases.

2.5.4. Mathematical Conditioning Issues

The performance of the solver is closely tied to the mathematical conditioning of the equation system it is solving. Conditioning can be thought of as the ratio of the largest eigenvalue of the system to the smallest eigenvalue. The larger this ratio becomes the more difficult it will be to solve the system and the larger potential error there will be in the solution.

Many model features can increase (or worsen) model conditioning.

- Rigid body modes (which have zero eigenvalues) are a special case which is handled as carefully as is theoretically possible, but which nonetheless have robustness problems.
- Mesh refinement increases gradients.
- Poorly shaped elements can lead to ill conditioned linear systems

- Understanding Type 1 versus Type 2 constraints as discussed in the definition of the GDSW option *max_numterm_C1* can lead to substantial performance improvements.
- Models in which a stiff material abuts a soft material are poorly conditioned.
- Mixtures of high density and low density components.
- Meshes with mixtures of solid and structural elements.

A poorly conditioned system can cause difficulties for the solver, both from an accuracy and robustness standpoint.

- Long solve times or many iterations required for convergence
- Difficulty obtaining a low target residual during the solve
- Large difference between the 'recursive' and 'actual' residuals reported by the solver. This can also indicate the solver result has lower accuracy than intended.
- Significant localized errors in the result obtained at some nodes

2.5.5. Frequency Response Functions

Several additional SD training documents⁵³ and presentations describe the solver behavior, debugging, and usage guidelines in more detail. See the **Sierra/SD** training documents on 'GDSW Solver Accuracy Notes', 'GDSW 101', 'GDSW Memory Use Tutorial'.

2.5.6. Parameters

SubdomainData is described in Section 2.5.3.

The *custom* option for *Hprecond* in Table 2-19 preconditions the matrix

$$-\omega^2(\alpha_M + i\beta_M)M + i\omega C + (\alpha_K + i\beta_K)K,$$

where ω is the circular frequency of excitation, i is the imaginary unit, and M , C and K are the mass, damping, and stiffness matrices, respectively. With $\gamma = \text{structural\_damping}$ and $\beta = \text{viscous\_damping}$ the non-zero parameters for the other preconditioning options are $\alpha_K = 1$ for *stiffness*, $\alpha_K = 1$, $\alpha_M = -1$ for *Laird-Giles*, $\alpha_K = 1$, $\alpha_M = 1$, $\beta_M = -\gamma$ for *shifted Laplacian*, and $\alpha_K = 1$, $\beta_K = \gamma + \beta\omega$, $\alpha_M = 1$ for *operator*. Notice one should not use the *stiffness* preconditioning option for ω near zero for structures with rigid body modes since K is singular or near singular in this case.

Table 2-17. – GDSW Section Options (Advanced).

Variable	Values	Default	Description
version	<i>integer</i>	2	GDSW version
num_vectors_keep	<i>integer</i>	-1	number of stored search directions to keep after culling
max_previous_sols	<i>integer</i>	0	maximum number of previous solutions used to accelerate convergence
num_sub_per_proc	<i>integer</i>	1	number of subdomains per processor
num_iter_improve_I	<i>integer</i>	0	number of iterative improvement steps for I_solver = LDM and precision_option_I = single
num_iter_improve_O	<i>integer</i>	0	number of iterative improvement steps for O_solver = LDM and precision_option_I = single
num_iter_improve_coarse	<i>integer</i>	0	number of iterative improvement steps for coarse_solver = LDM and and precision_option_coarse = single
overlap_method	<i>integer</i>	0	0 - standard graph-based 1 - element-based
cull_method	<i>integer</i>	1	0 - none: reach maximum and then stop 1 - simple: reach maximum and then remove most recent ones 2 - eigen: cull search directions based on solution to eigenvalue problem
num_vectors_keep	<i>integer</i>	-1	number of stored search directions to keep after culling (-1 automatic selection)
reduced_option_coarse	<i>integer</i>	3	same as reduced_option but for subregions
coarse_size	<i>integer</i>	0	0 - automatic selection of coarse space size 1 - select small coarse space 2 - select large coarse space
preconditioner_type	<i>integer</i>	2	1 - BDDC 2 - GDSW 3 - DIAG 4 - NODAL
interface_precond	<i>integer</i>	0	0 - do not require interface preconditioner 1 - require interface preconditioner
coarse_connectivity_option	<i>integer</i>	1	algorithm number for coarse elem connect
viscous_damping	<i>real</i>	0	viscous damping coefficient for preconditioner operator
structural_damping	<i>real</i>	0.12	structural damping coefficient for preconditioner operator
rbm_tolerance	<i>real</i>	1e-12	tolerance used to flag rigid body modes in dd_solver.dat. Recommend usage of the rbmtolerance in the parameters section.
con_tolerance	<i>real</i>	2.5e-9	singularity tolerance for processing constraints
con_row_tolerance	<i>real</i>	1e-1	pivoting tolerance for processing constraints
ML_max_level	<i>integer</i>	7	maximum number of levels for multilevel local solver
ML_max_coarse	<i>integer</i>	1000	maximum number of unknowns for coarsest level
use_epetra_coarse	<i>integer</i>	0	0 - do not use Epetra matrices for coarse correction 1 - use Epetra matrices for coarse correction
parmetis_option	<i>integer</i>	0	Parmetis option for coarse problem partitioning: 0 - PartKway 1 - PartGeomKway

Table 2-18. – GDSW Section Print Options.

Variable	Values	Default	Description
ML_print_coarse	<i>integer</i>	0	0 - no output 1 - print coarse stiffness matrix
ML_print_Phi	<i>integer</i>	0	0 - no output 1 - print interpolation matrix
pardiso_message_level	<i>integer</i>	0	0 - no messages 1 - print messages
prt_matrix	<i>integer</i>	0	0 - no output 1 - print out matrix in 3-column format 2 - print out matrix in CSR format
prt_subdomain_coarse	<i>integer</i>	0	0 - do not print subdomain coarse matrix 1 - print subdomain coarse matrix
prt_subdomain_PU	<i>integer</i>	0	0 - do not print sub partition of unity 1 - print sub partition of unity
prt_stabilization	<i>integer</i>	0	0 - do not print coarse stabilization matrices 1 - print coarse stabilization matrices
prt_interior	<i>integer</i>	0	0 - do not print interior matrices 1 - print interior matrices
prt_overlap	<i>integer</i>	0	0 - do not print overlap matrices 1 - print overlap matrices
prt_memory	<i>integer</i>	0	0 - do not print gdsw memory diagnostics 1 - print gdsw memory diagnostics
prt_debug	<i>integer</i>	0	0 - do not print subdomainData.dat 1 - print subdomainData.dat
write_orthog_data	<i>integer</i>	0	0 - do not write orthogonalization data to file 1 - write orthogonalization data to file
ML_print_coarse	<i>integer</i>	0	1 - ignore memory reporting 0 - do not print coarse stiffness matrix 1 - print coarse stiffness matrix
ML_print_Phi	<i>integer</i>	0	0 - do not print interpolation matrix 1 - print interpolation matrix

Table 2-19. – GDSW Section Options (Helmholtz).

Variable	Values	Default	Description
Hprecond	<i>integer</i>	5	Helmholtz preconditioner: 0-stiffness: Stiffness based 1-LG: Laird-Giles 2-custom: Custom 3-SL: shifted Laplacian 5-operator: Operator with damping
orthogH	<i>integer</i>	20	maximum number of stored search directions for Helmholtz problems
max_previous_sols	<i>integer</i>	0	maximum number of previous solutions used to accelerate convergence
precondUpdateFreq	<i>integer</i>	10	frequency to update preconditioner as operator changes
viscous_damping	<i>real</i>	0	viscous damping coefficient (see text)
structural_damping	<i>real</i>	0.12	structural damping coefficient (see text)
alphaK	<i>real</i>	0	custom precond stiffness coefficient (see text)
betaK	<i>real</i>	0	custom precond stiffness coefficient (see text)
alphaM	<i>real</i>	0	custom precond mass coefficient (see text)
betaM	<i>real</i>	0	custom precond mass coefficient (see text)
krylov_methodH	<i>integer</i>	5	same as krylov_method but for Helmholtz problems
SC_optionH	<i>integer</i>	0	same as SC_option but for Helmholtz problems

Table 2-20. – GDSW Section Options (Solvers).

Variable	Values	Default	Description
default_solver	<i>integer</i>	1	1-direct: Esmond Ng's sparse direct solver 2-LDM: Clark's LDM' sparse direct solver 6-NoPivot: Clark's sparse direct solver
PTAP_solver	<i>integer</i>	1	solver for conjugate gradient matrix 0-diag: diagonal (holds in exact arithmetic) 1-full: full $\Phi^T A \Phi$ matrix

2.6. Sensitivity

Sensitivity to parameters is available for modal analysis,⁵² Craig-Bampton reduction (CBR), static solutions and some transient solutions. An example input file for modal analysis is given in the Section 9. In the case of CBR analysis, we refer to sections 3.5 and 3.5.1 for a detailed discussion of how to perform sensitivity analysis. The **sensitivity** section controls global parameters related to sensitivity analysis. Sensitivity analysis is not performed in **Sierra/SD** unless this section is present in the input file. The following example illustrates the legal keywords. Valid keywords are identified in Table 2-21. Lists of numbers should follow the rules for integer lists detailed in Section 2.1.

```
SENSITIVITY
  values all
  vectors 1,3,5,7:9
  iterations 8
  tolerance 1e-7
  Attune
  AttuneNodeset sensNS
END
```

Keyword	argument	Description
values	“all”/“none” or list	eigenvalue selection
vectors	“all”/“none” or list	eigenvector selection
iterations	integer	number of eigen vector iterations
tolerance	float	convergence tolerance for eigen vectors
attune	n/a	enable attune output
AttuneNodeset	string	nodeset for reduced model

Table 2-21. – Sensitivity Analysis Keywords.

The keywords **values** and **vectors** are used to control what types of sensitivities are computed for which cases in the analysis. In modal analysis, these refer to the eigenvalues and eigenvectors, respectively, and the case numbers represent the mode numbers. In static analysis, **vectors** refers to the displacement vector results, and **values** has no meaning. Also, in modal analysis, eigenvalue sensitivities are always computed when eigenvector sensitivities are requested for a mode. Allowable values are:

```
vectors all           // compute for all cases/modes
vectors none          // compute for no cases/modes
vectors 1:3,5         // range of cases/modes
```

Omitting the keyword **vectors** (or **values**) is equivalent to not requesting those sensitivities; in other words, it is equivalent to **vectors none**. The keywords **iterations** and **tolerance** are used in computing eigenvector derivatives. The default values are 10

and $1.0\text{e-}06$, respectively. If the eigenvector sensitivity approximation fails to converge to the desired tolerance in the specified number of iterations, a fatal error occurs.

2.6.1. Attune

An interface is provided to the **Attune** test/analysis correlation code supplied by ATA engineering. The data is written to an external text file, with a file name based on the input (*.inp) file name. A surrogate of the finite element model is determined using eigenvectors. Attune applies *only* to eigen sensitivity analysis, and the eigen modes must be selected using **values**. For output through this interface, the following two parameters must be defined.

attune: request interface output.

AttuneNodeset: identification of a nodeset to be associated with the test degrees of freedom. Note that even if test mode shapes are not available, Attune requires the definition of a reduced space model (using this nodeset). It is required for mode tracking.

To use **Attune** , please refer to the ATA website and on line documentation.⁵

2.6.2. Sensitivity Output

Sensitivity results are output to the same file as the nominal results. The arrangement of the output depends on the analysis. The statics nominal result is output, followed by the sensitivity result for each parameter. In eigenvalue problems, the nominal frequencies and eigenvectors are output, followed by the eigenvalue and eigenvector sensitivities with respect to the first parameter, the second parameter, etc. The eigenvalue sensitivities are placed in the time field of each output record, like the frequencies are for the nominal modal parameters. For transient analysis, the nominal response for each time step is output, followed by the sensitivities for that time step. Then the nominal results for the next time step are output. See Figure 2-5 for an example of eigen sensitivity.

The change of parameter (or tolerance) may be specified in any of three ways.

1. Specify an absolute tolerance by entering “+/-” followed by the number, e.g. “+/- 1.05e-4”.
2. Specify a relative tolerance by entering “+/-” followed by a number and the keyword “percent”. Each field should be separated by a space. For example,
56 +/- 2.0 percent
3. Use the default tolerance by entering only the “+/-” by itself. The default tolerance is 2 percent.

```
EXPLORE> select step 4
```

```
Time = 3.504E+3 (time step 4 of 30)
```

```
EXPLORE> gvar
```

```
Time = 3.504E+3 (time step 4 of 30)
```

```
Global Time Step Variables
```

```
ModeNumber = 4.0000E+00
```

```
EigenFrequency = 3.5042E+03
```

```
EigenVectScale = 1.0000E+00
```

```
deriv_Block_1000_area = 3.9362E+03
```

```
deriv_Block_101_thickness = 1.4426E+07
```

The order of parameters can be determined from the global variables. It is also available in the results file. The sensitivities may be extracted using the global variables.

```
EXPLORE> times
```

```
Number of time steps = 18
```

```
Step 1) 725.3E+0
```

```
Step 2) 725.3E+0
```

```
Step 3) 3.005E+3
```

```
Step 4) 3.504E+3
```

```
Step 5) 3.504E+3
```

```
Step 6) 4.929E+3
```

```
Step 7) 602.1E+0
```

```
Step 8) 602.1E+0
```

```
Step 9) 6.512E+0
```

```
Step 10) 3.936E+3
```

```
...
```

The first nmodes (6 in this example) eigenvalues and vectors are associated with the nominal structure. The next nmodes values are the $d\lambda/dP_1$ values associated with the first parameter, P_1 . The corresponding vectors are $d\phi_i/dP_1$.

Figure 2-5. – Eigen Sensitivity Example Data. In this example, both eigenvalue and eigen-vector sensitivities are computed. The data is probed using “explore”. Global variable include sensitivities to area in block 1000 and thickness in block 101.

The selection of parameters is controlled by the inclusion of a +/- symbol following a parameter in the input deck. Examples of valid sensitivity parameter definitions are:

```

MATERIAL 1
  E 10e6 +/- 1e6           // absolute tolerance specified
  density 2.59e-4 +/-      // no tolerance, use default
END

BLOCK 1
  area 0.10 +/- 5 percent  // relative tolerance specified
END

BLOCK 2
  thickness +/- 1 percent  // relative to Exodus attr
END

Loads
  nodeset 1
  force 0. 0. 1000 +/- 0 0 10 // tolerance for vector param
end

```

Note that the tolerances are specified on the parameters where they normally appear in the input file. That is, these definitions do not appear in the **sensitivity** section.

The sensitivity quantities output to the exodus file are derivatives, and can be used to compute a first-order approximation to the change in an output quantity with respect to a parameter. For example, the change in an eigenvalue λ that depends on a parameter p can be approximated

$$\lambda(p + \Delta p) = \lambda(p) + \frac{d\lambda}{dp} \Delta p + o(\Delta p).$$

The quantity written to the exodus output is $d\lambda/dp$. The tables printed in the results and standard output also include the linear approximation to the change in λ ,

$$\frac{d\lambda}{dp} \Delta p.$$

This approximation may be inaccurate if Δp is too large. The size of Δp is specified by the user in the input deck for each parameter included in the sensitivity study.

Derived Output Quantities The sensitivity analysis methods use a semi-analytic method. Primary variables (usually displacement) are computed in terms of changes in stiffness and mass matrices, and resultant displacements are then computed analytically. See Figure 2-6. Many other output quantities are computed in terms of these primary variables using the standard output routines. Most of these output quantities depend linearly on the primary variables. For example, computation of the derivatives of strain can be readily computed using the chain rule, and may also employ the same procedures for

strain computation. Let $\varepsilon = \kappa u$ define the strain/displacement relationship at a given point in the model. Here κ is a constant.

$$\begin{aligned}\frac{d\varepsilon}{dp} &= \frac{d\varepsilon}{du} \frac{du}{dp} \\ &= \kappa \frac{du}{dp}\end{aligned}$$

Code that computes ε from u may be used to compute $\frac{d\varepsilon}{dp}$ from $\frac{du}{dp}$.

Other variables are not linear with respect to the primary variables. For example, the strain energy or the von Mises stress include quadratic terms in displacement.

$$\begin{aligned}E_s &= u^T K u \\ \frac{dE_s}{dp} &= \frac{dE_s}{du} \frac{du}{dp} \\ &= u^T K \frac{du}{dp}\end{aligned}$$

These variables may not use the same code paths. Data is written, but is not correct for these variables.

Statics	$Ku = f$ $\Delta u = K^{-1}(\Delta f - \Delta K u)$
Eigen	$(K - \lambda M)\phi = 0$ $\Delta \lambda_i = \phi^T \Delta K \phi - \lambda \phi^T \Delta M \phi$

Figure 2-6. – Semi-Analytic Methods for Sensitivity Analysis.

2.6.2.1. Solution Types Sensitivity analysis is available only for the solution types shown in Table 2-22. The primary application is in eigenvalue problems where the semi-analytic solutions can provide significant computation and accuracy benefit over a finite difference approach.⁴

Table 2-22. – Sensitivity Analysis Solution Type Availability.

Name	Section	Description
eigen	3.10	Normal Modes
statics	3.28	Linear Statics

Sensitivity Limitations

1. Subsection Eigen Sensitivity Analysis section Solution Procedures of the Theory Manual explains how sensitivity analysis may be performed using most solvers.
2. Outputs are limited to variables linear in displacement (2.6.2). In particular, output of von Mises stress is output for these solutions, but is not correct. That is the derivative of the von Mises stress is not the von Mises stress of the derivative of the stress.

2.7. *Coordinate*

Local coordinate systems may be defined to orient directional materials (section 4.7.2.3), define constraints, boundary conditions, or loads in local orientations (sections 2.9, 6.1 and 6.3), or transform outputs to a local coordinate system (sections 7.3 and 7.4). The basic/default coordinate frame can also be explicitly referenced as coordinate “0”, if desired.

The ability to visualize and verify correctness of coordinate system definitions is a key element of the overall workflow. At present, support is limited to the material direction output for anisotropic isosolid elements; see section 4.1.3.

```
BEGIN RECTANGULAR COORDINATE SYSTEM <string>name
  ORIGIN = <real> <real> <real>
  Z POINT = <real> <real> <real> # Point on  $\tilde{Z}$  axis
  XZ POINT = <real> <real> <real> # Point on  $\tilde{X}\tilde{Z}$  plane
  ROTATE <real> ABOUT AXIS X|Y|Z # angle in radians
END

BEGIN RECTANGULAR COORDINATE SYSTEM <string>name
  ORIGIN NODESET = <nodelist> # Single-node nodesets
  Z POINT NODESET = <nodelist> # Node on  $\tilde{Z}$  axis
  XZ POINT NODESET = <nodelist> # Node on  $\tilde{X}\tilde{Z}$  plane
  ROTATE <real> ABOUT AXIS X|Y|Z # angle in radians
END
```

Syntax 2.2. Rectangular Coordinate System Syntax

```
BEGIN CYLINDRICAL COORDINATE SYSTEM <string>name
  ORIGIN = <real> <real> <real>
  Z POINT = <real> <real> <real> # Point on  $\tilde{Z}$  axis
  XZ POINT = <real> <real> <real> # Optional: point on  $\tilde{X}\tilde{Z}$  plane; sets
                                # location of azimuthal angle theta=0
  ROTATE <real> ABOUT AXIS X|Y|Z # angle in radians
END
```



```

BEGIN CYLINDRICAL COORDINATE SYSTEM <string>name
  ORIGIN NODESET = <nodelist>      # Single-node nodesets
  Z POINT NODESET = <nodelist>      # Node on  $\tilde{Z}$  axis
  XZ POINT NODESET = <nodelist>      # Optional: node on  $\tilde{X}\tilde{Z}$  plane; sets
                                      # location of azimuthal angle theta=0
  ROTATE <real> ABOUT AXIS X|Y|Z    # angle in radians
END

```

Syntax 2.3. Cylindrical Coordinate System Syntax

```

BEGIN SPHERICAL COORDINATE SYSTEM <string>name
  ORIGIN = <real> <real> <real>
  Z POINT = <real> <real> <real>    # Point on  $\tilde{Z}$  axis
  XZ POINT = <real> <real> <real>    # Optional: point on  $\tilde{X}\tilde{Z}$  plane; sets
                                      # location of azimuthal angle theta=0
  ROTATE <real> ABOUT AXIS X|Y|Z    # angle in radians
END

BEGIN SPHERICAL COORDINATE SYSTEM <string>name
  ORIGIN NODESET = <nodelist>      # Single-node nodesets
  Z POINT NODESET = <nodelist>      # Node on  $\tilde{Z}$  axis
  XZ POINT NODESET = <nodelist>      # Optional: node on  $\tilde{X}\tilde{Z}$  plane; sets
                                      # location of azimuthal angle theta=0
  ROTATE <real> ABOUT AXIS X|Y|Z    # angle in radians
END

```

Syntax 2.4. Spherical Coordinate System Syntax

```

BEGIN CONICAL COORDINATE SYSTEM <string>name
  ORIGIN = <real> <real> <real>
  Z POINT = <real> <real> <real>    # Point on  $\tilde{Z}$  axis
  XZ POINT = <real> <real> <real>    # Optional: point on  $\tilde{X}\tilde{Z}$  plane
  ROTATE <real> ABOUT AXIS X|Y|Z    # angle in radians
  ANGLE = <real>
END

BEGIN CONICAL COORDINATE SYSTEM <string>name
  ORIGIN NODESET = <nodelist>      # Single-node nodesets
  Z POINT NODESET = <nodelist>      # Node on  $\tilde{Z}$  axis
  XZ POINT NODESET = <nodelist>      # Optional: Node on  $\tilde{X}\tilde{Z}$  plane
  ROTATE <real> ABOUT AXIS X|Y|Z    # angle in radians
  ANGLE = <real>
END

```

Syntax 2.5. Conical Coordinate System Syntax

```
BEGIN ELLIPSOIDAL COORDINATE SYSTEM <string>name
  ORIGIN = <real> <real> <real>
  Z POINT = <real> <real> <real> # Point on  $\tilde{Z}$  axis
  XZ POINT = <real> <real> <real> # Point on  $\tilde{X}\tilde{Z}$  plane
  ROTATE <real> ABOUT AXIS X|Y|Z # angle in radians
  AXIS STRETCHING = <real> <real> <real> #angle in radians
END

BEGIN ELLIPSOIDAL COORDINATE SYSTEM <string>name
  ORIGIN NODESET = <nodelist> # Single node nodeset
  Z POINT NODESET = <nodelist>
  XZ POINT NODESET = <nodelist>
  ROTATE <real> ABOUT AXIS X|Y|Z # angle in radians
  AXIS STRETCHING = <real> <real> <real> #angle in radians
END
```

Syntax 2.6. Ellipsoidal Coordinate System Syntax

The prefix to `COORDINATE SYSTEM` specifies the type of local coordinate system. Supported types are: `RECTANGULAR` (Cartesian), `CYLINDRICAL` (Polar), `SPHERICAL`, `CONICAL` (a cylindrical system with an aperture angle), and `ELLIPSOIDAL` (a spherical system stretched by `AXIS STRETCHING`).

A local element coordinate system is defined by a set of three control points. These points may be defined using absolute locations (`ORIGIN`, `Z POINT`, `XZ POINT`) or using nodelists in the mesh file (`ORIGIN NODESET`, `Z POINT NODESET`, `XZ POINT NODESET`) that *contain exactly one node each*. The control points define a local $\tilde{X}\tilde{Y}\tilde{Z}$ Cartesian coordinate system centered at the coordinate system origin. For cylindrical, spherical, conical, and ellipsoidal systems the orientation of the basis vectors at individual nodes and elements is spatially dependent. In this documentation, the local basis vectors at each node or element will be denoted as $\hat{r}$, $\hat{s}$, and $\hat{t}$.

Note: coordinate systems defined by nodesets use only the initial location of the nodes and do not update with model displacement.

- **ORIGIN:** Required. Specifies the origin location for the new coordinate system.
- **Z POINT:** Required. Specifies a point on the $\tilde{Z}$ axis for the new coordinate system. For `SPHERICAL`, this defines where the zenith (polar) angle ϕ is 0.
- **XZ POINT:** Required only for `RECTANGULAR` and `ELLIPSOIDAL`. Third point to specify the $\tilde{X}\tilde{Z}$ -plane of the new system. Must not lie on the $\tilde{Z}$ axis. If not specified, an arbitrary axis will be chosen based on the global coordinate system. For `CYLINDRICAL` and `SPHERICAL`, this defines where the azimuthal angle θ is 0. `XZ POINT` also affects

behavior at ill-defined points such as the origin of SPHERICAL systems, or along the $\tilde{Z}$ axis of CYLINDRICAL systems (see figure 2-11).

- **ANGLE** (CONICAL COORDINATE SYSTEM only): Required. Defines the angle in radians to rotate the axial direction away from the $\tilde{Z}$ axis (see figure 2-9).
- **AXIS STRETCHING** (ELLIPSOIDAL COORDINATE SYSTEM only): Required. Defines the aspect ratio to stretch a spherical system (see figure 2-10).
- **ROTATE**: Processed last. Specifies a rotation (in radians) of the coordinate system about one of its own axes. The conical system is effectively a cylindrical system with a local rotation about the $\tilde{Y}$ axis. Only one rotation may be specified.

The three coordinate control points are illustrated in figure 2-7. The coordinate system types are depicted in figures 2-8 to 2-10.




Figure 2-7. – Coordinate system definition vectors: note that XZ POINT determines the $\tilde{X}$ axis, but need not lie along it, which is the case depicted in this figure.




Figure 2-8. – Three examples (rectangular, spherical, and cylindrical) of transformed coordinate systems are given. The center of the rectangular block ($6 \times 2 \times 3$) is located at $(3, -1, 5)$ and rotated 30 degrees about the X axis. The center of the sphere (radius of 2) is located at $(3, -1, 5)$. The center of the cylinder (radius of 1, height of 2, and rotated 90 degrees about the X axis) is located at $(-5, 3, 0)$. The new coordinate systems are defined respectively of each of the geometries and are indicated by $\tilde{X} \tilde{Y} \tilde{Z}$ and the subscripts (r , s , and c).




Figure 2-9. – Conical Coordinate System Definition at $\tilde{X} \tilde{Z}$ plane




Figure 2-10. – Ellipsoidal Coordinate System Using `axis_stretching=(2.0, 1.1, 1.0)`. $\hat{r}$ is red, $\hat{s}$ is green, and $\hat{t}$ is blue.

Examples of rectangular, spherical, and cylindrical coordinate systems are given in figure 2-8. In those examples, we wish to define coordinate systems for three geometries: a brick, a sphere, and cylinder. The corresponding input deck is shown below.

```
begin rectangular coordinate system rc
  origin 3 -1 5
  z point 5 -1 5
  xz point 3 0.7321 6
end
coordinate spherical coordinate system ballLike
  origin 5 4 -2
  z point 3 4 -2
  xz point 5 2 -2
end
begin cylindrical coordinate pinSystem
  origin -5 2 0
  z point -5 4 0
  xz point -5 2 2
end
```

Input 2.3. Example coordinate system input

For cylindrical, conical, spherical, and ellipsoidal systems, the local basis vectors $\hat{r}$, $\hat{s}$, and $\hat{t}$ are not all well-defined when sampled along the $\tilde{Z}$ axis. At the origin, all coordinate

systems fall back to a local Cartesian system with $\hat{r}$ along the $\tilde{X}$ axis, $\hat{s}$ along the $\tilde{Y}$ axis, and $\hat{t}$ along the $\tilde{Z}$ axis. This is also true for the cylindrical and conical systems on the $\tilde{Z}$ axis. Spherical and ellipsoidal systems on the $\tilde{Z}$ axis will instead maintain the first vector ($\hat{r}$) pointing radially outward, the third vector ($\hat{t}$) pointing along the $\tilde{X}$ axis, and the second vector ($\hat{s}$) as either the positive or negative $\tilde{Y}$ axis (generating a right handed system). This maintains the normal and tangential properties of the three vectors for spherical coordinates. See figure 2-11 for a visual representation of this behavior.




Figure 2-11. – Coordinate system behavior near the $\tilde{Z}$ axis. $\hat{r}$ is red, $\hat{s}$ is green, and $\hat{t}$ is blue.

The **Exodus** mesh input file also has the ability to define local coordinate frames. These frames will be read and available during an analysis. In the case of a coordinate frame being defined in both the **Exodus** file and the input file, the input file definition will be used, and a warning will be logged.

In spherical coordinates, it may help to consider the Cartesian frame $(\tilde{X}, \tilde{Y}, \tilde{Z})$ with the same orientation as (r, θ, ϕ) :

Coordinate System	History Variable	Corresponding Coordinate
Cylindrical	X	r
	Y	θ
	Z	z
Spherical	X	r
	Y	θ
	Z	ϕ

Table 2-23. – Coordinate Names for history files.

$$\begin{aligned}
\tilde{X} &= r \sin(\phi) \cos(\theta) \\
\tilde{Y} &= r \sin(\phi) \sin(\theta) \\
\tilde{Z} &= r \cos(\phi),
\end{aligned} \tag{2.1}$$

$$0 \leq \phi < \pi, \quad 0 \leq \theta \leq 2\pi.$$

If the user specifies a coordinate system in the **History** section, notice that its applicability may be somewhat limited, though convenient. In particular, only a single history file is written in each analysis, and only one coordinate frame may be outputted per node (see section 7.3). The history file will display variables as Cartesian regardless of coordinate choice. Table 2-23 shows the corresponding values for cylindrical and spherical coordinates.

Finally, we provide a full example of how we can use the coordinate definitions. We use the same coordinate definitions from figure 2-8 and input 2.3, and we define the brick as block 1, the sphere as block 2, and the cylinder as block 3, respectively. We also specify a sideset 100 on the brick, a nodeset 200 on the sphere, and a nodeset 300 on the cylinder as shown in figure 2-12. We wish to apply a radial force emanating from the middle of the cylinder, and we want to record the output accelerations on the sphere. The corresponding input deck is attached, with associated coordinate systems defined in input 2.3.

```

Loads
  nodeset 300
  coordinate ballLike
  force = 1 0 0
end
History
  nodeset 200
  coordinate pinSystem
  accel
end

```

Tractions can also be specified on a particular coordinate frame, but special care is required. We refer the reader to section 6.3.3 for using arbitrary coordinates with tractions.




Figure 2-12. – An example of how coordinate systems are used. Accelerations are measured on the sphere (nodeset 200) due to a radial force applied on the cylinder (nodeset 300).

2.8. *Function*

Time, frequency and/or spatially dependent functions for transient and frequency response analysis can be defined using the **function** section. The following are simple examples of the use of a **function** .

```
FUNCTION test_func1
  type LINEAR
  data 0.0 0.0
  data 0.0150 0.0
  data 0.0152 1.0
  data 0.030 0.0
END

FUNCTION poly_fun
// a pulse of duration .05
// peak value 6/7 at 1/49 sec
// pulse(t)=8*sqrt(t)-686*t^2
  name "Smooth Pulse with Duration 0.05"
  type POLYNOMIAL
```



```
data 0. 0.  
data .5 8.  
data 2. -686.  
END
```

The function identifier string is given after function keyword. This identifier is used to reference the function in other parts of the input deck. It can be any string, but once this identifier could only be an integer. The keywords for these function definitions are:

1. TYPE to define the functional form,
2. NAME Additional reference information for the function. This may be printed to error messages or other informational output regarding the function, however it is the identifier that is always used to call the function from other parts of the input deck.
3. DATA for the functional parameters.

Other function definitions may require more parameters.

2.8.1. Function Offset/Shifts

Function input/output values can be offset and scaled. The syntax for defining these transformations is below.

```
X|abscissa offset = <real, default=0>  
Y|ordinate offset = <real, default=0>  
X|abscissa scale = <real, default=1>  
Y|ordinate scale = <real, default=1>
```

Given a function $y = f(x)$, the transformed output would be:

$$y = y_{scale} \cdot [f(x_{scale} \cdot [x + x_{offset}]) + y_{offset}].$$

Note that function offsets/shifts are currently only enabled for single-input, single-output functions. Otherwise, a warning will be issued and the offsets/shifts will be ignored.

2.8.2. Linear Functions

For linear functions, the data elements are pairs specifying the independent variable (e.g., time) and the corresponding function value. Evaluation of the function at unspecified points is performed via linear interpolation. In order to enforce uniqueness of the interpolant, **Sierra/SD** ignores points with an independent variable value less than that of the previous point. For example, the last point in the `ignored_point` function is ignored.

```
FUNCTION ignored_point  
type linear  
data 0.00 0.
```

```

data 0.01 1.
data 0.05 1.
data 0.04 0. // ignored: column 1 may not decrease
END

```




Figure 2-13. – Linear function "ignored_point".

Figure 2-13 shows the graph of function 3.

Extrapolation refers to evaluating the function before the first given time or after the last given time. Linear functions use the value of the nearest data point to extrapolate. For example, in the following function, $f(3/100) = 1/2$.

```

FUNCTION extrap_fun
  type linear
  data 0.00 0.
  data 0.01 1.
  data 0.02 0.5
END

```




Figure 2-14. – Linear function "extrap_fun".

If an independent variable value is listed multiple times, then the function value at that point is the average of the specified function values. For example, in the following function, $f(3/100) = 3/4$.

```

FUNCTION 5
  name "multiple_fun"
  type linear
  data 0.00 0.
  data 0.01 1.
  data 0.03 1.
  data 0.03 0.5
END

```

Function 5 is pictured in Figure 2-15.




Figure 2-15. – Linear function #5. "multiple_fun".

2.8.3. Sierra SM Piecewise Linear Functions

Some limited Sierra SM syntax is supported for specifying a piecewise linear function in **Sierra/SD**. The data entered through this syntax is interpreted as a linear function, and follows all the rules specified above. This comes with limitations. Only the **piecewise linear** function type is supported, and a number of features are missing. Most importantly, ending the function with "END FUNCTION <string>" is not supported. This comes from an incompatibility with existing **Sierra/SD** syntax; "FUNCTION <string>" happens to be valid **Sierra/SD** syntax in this scope and context.

```

Begin Function function_name
  Type is Piecewise Linear
  Begin Values
    0.00, 0.
    0.01, 1.
    0.05, 1.
  End Values
End

```

2.8.4. Functions using Tables

Functions may be specified by reference to a linearly interpolated **table** (as discussed in section 2.8.19). The table must be of dimension=1. 1-dimensional tables behave identically to the linear functions described above. However, they will typically be much faster and more memory efficient than linear functions, especially as the data size grows.

The function in the following example is a tabular representation of the data of Figure 2-14 and Function `extrap_fun` above.

```
FUNCTION 7
  type table
  tablename=example7
END

Table example7
  dimension=1
  size=5
  datafile='example7.txt'
  origin 0.0
  delta .01
END
```

Within the datafile, “example7.txt”, the following data would be represented.

```
0.0
1.0
0.5
0.5
0.5
```

The linear function can be evaluated for any time, and the table is limited to the range 0-0.04. Table type functions require the **tablename** keyword.

2.8.5. Polynomials

For polynomials, the data points given are the exponent of the independent variable and the coefficient. The independent variable taken to any real power will always be evaluated as positive. If powers are repeated, their coefficients will sum. For example,

```
FUNCTION 6
  name "poly_fun"
  type polynomial
  data 0.0 0.
  data 1.0 1.
```

```
data 2.0 0.1
data 1.0 0.5
END
```

is equivalent to

```
FUNCTION 6
name "poly_fun"
type polynomial
data 0.0 0.
data 1.0 1.5
data 2.0 0.1
END
```

The function value as a function of the independent variable t is,

$$f(t) = 1.5t + 0.1t^2.$$

2.8.6. LogLog Functions

In frequency domain analysis, log/log functions are commonly used for application of loads. This is particularly true for random vibration inputs which are commonly specified on log/log plots. The LogLog option allows linear interpolation on a log/log plot so that only the corner frequencies need be specified. An example follows.

```
FUNCTION 1
name "my_LogLog"
type LogLog
data 1.0 1e-8
data 299 1e-8
data 300 0.01
data 2000 0.03
data 8000 0.03
data 10000 0.01
data 10001 1e-8
END
```

2.8.7. SamplingRandom

The random pressure loading defined in subsection Random Pressure Loading section Loads and Materials of the Theory Manual provides a means of applying a pressure load with a specified spatial and temporal correlation. In many cases, the desired random

function is a function of time only. In those cases the `SamplingRandom` function provides a mechanism for applying the load.

Keyword	Values	Default	Description
type	string	required	must be “SamplingRandom”
cutoff_freq	<i>Real</i>	required [†]	cutoff frequency (in Hz)
omega_c (ω_c)	<i>Real</i>	required [†]	cutoff frequency (in rad/s)
deltaT	<i>Real</i>	π/ω_c	coarse time step
ntimes	<i>integer</i>	5	# of terms in time interpolation
correlation_function	<i>string</i>	defaults to $\sin(x)/x$	function for time interpolation
scale_function	string	defaults to $\sigma(z) = 1$	

[†]Specify either `cutoff_freq` or `omega_c`, but not both.

Table 2-24. – SamplingRandom function parameters.

Note that

$$\omega_c = 2\pi \times \text{cutoff\_freq},$$

and that only one of the two parameters `omega_c` and `cutoff_freq` can be specified. More detailed descriptions of these parameters are given in Section 6.3.15. Random time functions can be used to specify any type of random load, including pressure loads, force loads, acoustic loads, etc. Below we give an example for the case of an acoustic load.

```
LOAD
  sideset 1
    function = 1
    acoustic_vel = 1.0
END

FUNCTION 1
  type SamplingRandom
  cutoff_freq 1000
  deltaT 8.0e-4
  ntimes 5
END
```

The `SamplingRandom` function is a special case of a zero mean, unit variance Gaussian function. Sampling methods allow a reduced memory method of computing the time realization. In a transient analysis, the time integration step should be less than the coarse time step, “deltaT”. Statistics for the functions may be output by specifying “input_summary” in the “ECHO” section of the input file (see Section 7.7).

2.8.8. RandomLib Functions

There are two tests of the `RandomLib` function. One is a verification test.

In many cases, a random load on a structure may need a spatial correlation with other loads on the structure. The **RandomLib** function was created to address this need. ²

Run time parameters for “RandomLib” functions are listed in Table 2-25, and an example is provided in input 2.4. Each parameter is described in more detail below.

The RandomLib function operates only by reading data from an external **Exodus** data file. The data file is an **Exodus** file that contains nodal scalar boundary conditions on a nodeset that covers the same nodes as the sideset. This nodeset is assumed to have the name “surface_1_nodes” where the “1” in this case corresponds to sideset 1. These nodal loads are typically generated within MATLAB code and merged with the **Exodus** file definition.

Currently, this function has been applied only to apply a scalar function on the nodal locations of a single sideset in the model. Such functions can be used to apply pressures (which are applied as piecewise linear functions within the elements). It can also be used to apply prescribed accelerations at the nodal locations.

Keyword	Values	Description
type	randomlib	required to specify function
interp		temporal interpolation scheme none=nearest linear=linear interpolation
sideset	<i>int/string</i>	<i>sideset id/name where pressures are applied</i>

Table 2-25. – RandomLib function parameters.

```
function 55
  type=randomlib
  interp=none
  sideset=1
  exo_var scalar pressure
end
```

Input 2.4. Example RandomLib Function Specification

type The specification “type=randomlib” is *required* to reference the randomlib function and its capabilities.

interp A restart file contains time samples of a random function. **Sierra/SD** references these values at each time step to properly load the function. Figure 2-16 shows how interpolation influences the actual value returned.

²The RandomLib function is an external library interface to **Sierra/SD**. Additional functionality and the interface to other applications are described in separate documentation.

sideset Pressure is applied over a single sideset of the model. This sideset must match the definition in the load section. ¹

exo_var Specification of the name of the nodeset variable that represents the nodal loading keyword **exo_var** is required. The format should be “**exo_var scalar pressure,**” where ‘pressure’ is the name of the variable used in the **Exodus** file.




Figure 2-16. – RandomLib Temporal Interpolation. Because of different time steps in the RandomLib data library and the **Sierra/SD** time step algorithm, the function value returned depends on the time interpolation algorithm. With **interp=none**, the first value returned to **Sierra/SD** is about -1.0, as that is the nearest time sample in the data. With **interp=linear**, the value returned is about -0.6. Note that round off can cause odd behavior with **interp=none**, even if the two data sets have the same fundamental time step.

2.8.9. Analytic Functions

Analytic functions can be used define complex mathematical functions directly in the input deck. An example of the input for this type of function is:

¹Data for the **Exodus** file is usually provided using specialized tools such as **mkrandloadrst**. A sideset provides information about the extent of the load, and for pressure loads, it is required to identify the faces upon which the load is applied. Actual time history data is associated with a nodeset which includes the same nodes as the sideset.


```

FUNCTION sine
  type analytic
  evaluate expression = "sin(2 * pi * t)"
END

```

Sierra/SD uses the STK⁷⁰ expression parser. Analytic functions cannot be used for acceleration boundary conditions. Integration of analytic functions is not implemented.

Rules and options for composing algebraic expressions. If you choose to use the EVALUATE EXPRESSION command line, you will need to write the algebraic expressions. The algebraic expressions are written using a C-like format. Each algebraic expression is terminated by a semicolon(;). The entire set of algebraic expressions, whether a single expression or several, is enclosed in a single set of double quotes(" ").

Note: analytic function variables are case *insensitive*, including any local variable definitions, and there is currently no warning for multiply-defined variables (the last variable name “wins”), so caution should be taken to ensure that duplicate variable names aren’t used.

Expressions compute an output value as a function of input values. The input value of an expression is called an independent variable of the expression. An expression can use any name for the expression independent variable. The examples shown here use **x** for the independent variable.

Example: Return **sin(x)** as the value of the function.

```

function sinx
  type is analytic
  evaluate expression is "sin(x)"
end

```

Example: Return a piecewise linear interpolated ramp function. For x values less than 0.0 the function will return 0.0. For x values between 0.0 and 0.5 the function will return y values linearly interpolated between 0 and 100. For x values greater than 0.5 the function will return 100:

```

function pressure
  type is analytic
  evaluate expression is " \#
    (x <= 0.0) ? ( \#
      0.0 \#
    ) : ( \#
      (x < 0.5) ? ( \#
        x*200.0 \#
      ) : ( \#
        100.0 \#
      )
    )
  "
end

```

```

    )
    )
    "
end
\#
\#

```

Operators and Functions Available within Expressions The following functionality is currently implemented for the expressions:

Operators Valid arithmetic and Boolean operations are as follows:

Symbol	Operation
+, -	plus, minus
*, /	times, divide
^	power
==, !=	equivalent, not equivalent
>, <	greater/less than
>=, <=	greater/less than or equal to
!	not
&, , &&,	and, or
?, :	ternary if, then, else

Parentheses (), represents standard mathematical meaning of operation ordering.

Component indexing `value[i]`, accesses the i th component of a multi-component variable such as a vector, tensor, etc. This index is one-based. For a vector `vec`,

```

vec_x = vec [1]
vec_y = vec [2]
vec_z = vec [3]

```

for a symmetric tensor `sym`,

```

sym_xx = sym [1]
sym_yy = sym [2]
sym_zz = sym [3]
sym_xy = sym [4]
sym_yz = sym [5]
sym_zx = sym [6]

```

and for a full tensor `full`,

```

full_xx = full [1]
full_yy = full [2]
full_zz = full [3]
full_xy = full [4]
full_yz = full [5]
full_zx = full [6]

```

```

full_yx = full [7]
full_zy = full [8]
full_xz = full [9]

```

Math functions Valid mathematical operations are as follows:

Function	Operation
<code>abs(x)</code>	absolute value of x
<code>mod(x,y)</code>	modulus of $x y$
<code>min(x,y)</code>	minimum value of x, y
<code>max(x,y)</code>	maximum value of x, y
<code>sign(x)</code>	sign operator: -1 if x is negative, 1 if positive
<code>ipart(x)</code>	integer part of x
<code>fpart(x)</code>	fractional part of x

Power functions Valid functions for expressions containing exponents are as follows:

Function	Operation
<code>pow(x,y)</code>	x^y
<code>pow10(x)</code>	10^x
<code>sqrt(x)</code>	$\sqrt{x}$

Trigonometric functions Valid trigonometric operations are as follows:

Function	Operation
<code>acos(x)</code>	arccosine of x
<code>asin(x)</code>	arcsine of x
<code>asinh(x)</code>	inverse hyperbolic sin of x
<code>atan(x)</code>	arctangent of x
<code>atan2(y,x)</code>	arctangent of y/x , signs of x and y
<code>cos(x)</code>	cosine of x
<code>cosh(x)</code>	hyperbolic cosine of x
<code>sin(x)</code>	sine of x
<code>sinh(x)</code>	hyperbolic sine of x
<code>tan(x)</code>	tangent of x
<code>tanh(x)</code>	hyperbolic tangent of x

Logarithm functions Valid logarithmic operations are as follows:

Function	Operation
<code>log(x)</code> or <code>ln(x)</code>	natural logarithm of x
<code>log10(x)</code>	base-10 logarithm of x
<code>exp(x)</code>	e^x

Rounding functions Valid mathematical rounding operations are as follows:

Function	Operation
<code>ceil(x)</code>	smallest integer value not less than x
<code>floor(x)</code>	largest integer value not greater than x

Random functions Valid functions for generating random data are as follows:

Function	Operation
<code>random()</code>	random real number $0.0 \leq x < 1.0$
<code>random(x)</code>	seeds the random number generator
<code>time()</code>	elapsed time in seconds since January 1, 1970

Coordinate conversion functions Valid functions for converting to and from polar and Cartesian coordinate systems are as follows:

Function	Operation
<code>deg(x)</code>	converts x from radians to degrees
<code>rad(x)</code>	converts x from degrees to radians
<code>recttopolr(x,y)</code>	magnitude of vector (x, y)
<code>recttopola(x,y)</code>	angle of vector (x, y)
<code>poltorectx(r,th)</code>	x -coordinate of angle th at distance r
<code>poltorecty(r,th)</code>	y -coordinate of angle th at distance r

If, then, else ternary syntax $A ? B : C$, where A is a Boolean statement such as $(x < 2.0)$, B is a statement (or set of statements) to be evaluated if A is true, and C is a statement (or set of statements) to be evaluated if A is false. **Note:** these statements can *only* be used for return values, i.e. you cannot define local variables inside a ternary statement.

Other utilities Miscellaneous other utility functions are as follows:

- `cos_ramp(x, xstart, xend)` defines a cosine ramp loading function. For $x < xstart$, it returns 0.0. For $x > xend$, it returns 1.0. If $xstart \leq x \leq xend$, the function has the value $(1.0 - \cos(\text{Pi} * (x - xstart) / (xend - xstart))) / 2.0$. The primary purpose for the cosine ramp is to provide smooth loading. If attempting to run a nearly quasistatic problem with the dynamic solver, a displacement boundary condition applied via the `cos_ramp` will generally give a smooth response for a given loading time.
- `cycloidal_ramp(x, xstart, xend)` defines a cycloidal ramp function. For $x < xstart$ it returns 0.0. For $x > xend$, it returns 1.0. If $xstart \leq x \leq xend$, the function has the value $(x - xstart) / (xend - xstart) - 1 / (2 * \text{pi}) * \sin(2 * \text{pi} / (xend - xstart) * (x - xstart))$. The primary purpose for the cycloidal front ramp is to provide smooth loading for prescribed displacements and velocities. This function has continuous acceleration derivatives at `xstart` and `xend`, providing a smoother response than the `cos_ramp` for prescribed displacement boundary conditions.

- `haversine_pulse(x, xstart, xend)` defines a haversine pulse function. For $x < xstart$, it returns 0.0. For $x > xend$, it returns 0.0. If $xstart \leq x \leq xend$, the function has the value $\text{pow}(\sin(\text{Pi} \cdot (x - xstart) / (xend - xstart)), 2)$. Design shock loads for components are often specified as haversine acceleration pulses; this function is included to make it easier to apply this loading.

Constants. There are three predefined constants that may be used in an expression. These three constants are `e`, `pi`, and `two_pi`. Note that these constants are reserved variable names and thus cannot be redefined in an expression.

Math symbol	Expression	Value
e	<code>e</code>	2.7182818284...
π	<code>pi</code>	3.1415926535...
2π	<code>two_pi</code>	$2 * 3.1415926535...$

Also, there are two predefined constant functions that can be used, `SIERRA_CONSTANT_FUNCTION_ZERO` and `SIERRA_CONSTANT_FUNCTION_ONE`. These two functions are equivalent to defining functions with constant expressions “0.0” or “1.0”.

Multiple inputs By default, we assume that analytic functions have a single input value.

Typically, the input would represent time, but the variable name can be whatever you choose (e.g. “x”, “t”, or “time”). Sometimes, however it is beneficial to define functions that depend on more than one input, for example both time *and* the current displacement. In that case, the user must supply additional information to define the meaning of each input variable:

```
expression variable <string> = time|coord|disp|velocity|acceleration|nodeId
expression variable <string> = nodal <string>
```

Note that an “expression variable” line is required for *each* independent variable, even if the name and meaning happen to be the same (e.g. `expression variable time = time`). Also, note that analytic functions are evaluated at each application node on the structure (i.e. each node in a nodeset). The list of supported input variables is given in table 2-26.

Nodal variables may be used in the “expression variable” line to define independent analytic function variables. Current testing only guarantees accuracy for variables read from the input mesh, as they do not change with time. Output variables change over time, and current testing does not guarantee proper evaluation order of functions and outputs. Nonetheless, output variables are accessible through this interface.

Analytic functions are case-insensitive. `PI`, `Pi`, `pI` and `pi` are equivalent. If it isn’t, then that is a bug. Note however that if you define both `pi` and `Pi` as different variables, the expression variable parser treat them as the same variable, and only use the last definition.

Variable	Description
time	Current time for the function evaluation.
coord	Undeformed coordinates at each node.
disp	Deformation vector at each node.
velocity	Velocity vector at each node.
acceleration	Acceleration vector at each node.
nodeid	The global node ID at each node. Note: <i>unmapped Exodus node id</i> (1:N)

Table 2-26. – Predefined Analytic Input Variables.

This section combines documentation from Sierra SM and Sierra SD. Having a shared document has pitfalls. **Sierra/SM** and **Sierra/SD** have different syntax. For example this section uses the “is” and “are” of Sierra SM. In SM usually ‘=/is/are’ can be used interchangeably. Readers should be aware that “are” ain’t defined in Sierra SD. It is not obvious what will happen if someone uses the “is” Sierra SM syntax in Sierra SD. One way to translate the Sierra SM syntax to Sierra SD syntax is to replace “is” and “are” by =.

Some parts of the documentation recommend using time as the independent variable. Sometimes the independent variable is time. Other parts of the documentation states that time() is a built in function that measures time is a surprising way. This part of the documentation discourages using time as the independent variable.

The documentation says that each algebraic expression is terminated by a semicolon. There are no examples. No one has time to add examples. There are examples in the regression test suite.

A function has a name parameter. The name is not documented. No one has time to document it.

Scientific notation is not mentioned in the documentation. If it is possible to rewrite 100000 as 1.0e5, then this would be worth documenting. What about E, d or D? The scientific notation "1.0e-5" works. No one knows about other notation.

```
function left
  type analytic
  name "otherleft"
  evaluate expression = "1500*pow(sin(1.0e5*t*pi/9),2)"
end
```

2.8.10. Plane Wave (Time Domain)

Plane wave functions are applied primarily in scattering problems where a load on a surface is analytically described as an incident plane wave. We define the wave in terms of the following parameters.

Keyword	Values	Description
type	plane_wave	identifier keyword
Direction	3 <i>reals</i>	direction of the wave, $\mathbf{e}_k = \mathbf{k}_0/ \mathbf{k}_0 $
material	<i>string</i>	acoustic material
K0	<i>real</i>	wavenumber, $k_0 = \mathbf{k}_0 $
origin	3 <i>reals</i>	wave origin, $\mathbf{x}_0$

The material specification provides the wave speed c_0 and the fluid density ρ_0 , and the pressure amplitude p_0 is specified in the **loads** section. The angular frequency is $\omega = 2\pi f$ and $k_0 = \omega/c_0$. A time-harmonic plane wave can then be written as

$$p = p_0 \cos[k_0 c_0 t - \mathbf{k}_0 \cdot (\mathbf{x} - \mathbf{x}_0)], \quad (2.2)$$

The velocity is also computed for scattering problems,

$$\mathbf{v} = \mathbf{e}_k \frac{p}{\rho_0 c_0}. \quad (2.3)$$

```

loads
  sideset 1 // acoustic
    acoustic_vel = 1.0
    function = 65
  sideset 2 // structure
    pressure = 1.0
    function = 65
end

function 65
  type = plane_wave
  Direction = 1 0 0
  origin = 0 0 0
  K0 = 1000
  material = air
end

material air
  acoustic
  c0=332.0
  density=1.29
end

tied data
  surface 1,2
end

```

Input 2.5. Example Planewave Function Specification

An example is shown in input 2.5. Note that if the `tied data` block is omitted, there is no interaction between the acoustic domain and the structure. Instead, the boundary of the acoustic domain is rigid, and the scattered pressure field is from a rigid boundary instead of from a structure with the specified material properties.

2.8.11. Plane Wave (Frequency Domain)

An incident plane wave can also be applied in the frequency domain for scattering problems. In the frequency domain, a plane wave is defined in terms of the following parameters.

Keyword	Values	Description
type	plane_wave_freq (or iplane_wave_freq)	identifier keyword
Direction	3 reals	direction of the wave, $\mathbf{e}_k = \mathbf{k}_0/ \mathbf{k}_0 $
material	string	acoustic material
origin	3 reals	wave origin, $\mathbf{x}_0$

The wave speed c_0 and the density ρ_0 are specified by the choice of material, the frequency f is specified in the `frequency` block, and the particle velocity amplitude $u_0 = p_0/\rho_0 c_0$ is specified in the `loads` block. A time-harmonic plane wave (with the time-dependence dropped) can then be written as

$$p(\mathbf{x}) = p_0 e^{-i\mathbf{k}_0 \cdot (\mathbf{x} - \mathbf{x}_0)}. \quad (2.4)$$

```
loads
  sideset 1 // acoustic
    acoustic_vel = 1.0
    function = 66
  sideset 1
    iacoustic_vel = 1.0
    function = 67
  sideset 2 // structure
    pressure = 1.0
    function = 66
  sideset 2
    ipressure = 1.0
    function = 67
end

function 66
  type = plane_wave_freq
  Direction = 1 0 0
  origin = 0 0 0
```



```

    material = air
end
function 67
    type = iplane_wave_freq
    Direction = 1 0 0
    origin = 0 0 0
    material = air
end

material air
    acoustic
    c0=332.0
    density=1.29
end

tied data
    surface 1,2
end

```

Input 2.6. Example PlanewaveFreq Function Specification

An example is shown in input 2.6. The syntax for the plane wave frequency function does not require a wavenumber, K_0 , but is otherwise the same as **plane_wave** function from the previous section. Note that both a real and an imaginary plane wave frequency function are applied to their corresponding real and imaginary loads—which is currently necessary in order to apply the correct phase shift between real and imaginary parts of propagating waves—and that real and imaginary loads should be given identical amplitude, direction, and origin.

2.8.12. Planar Step Wave

The planar step wave, keyword=“planar_step_wave” provides a means of applying a traveling exponential step wave to an acoustic scattering problem. The function provides both a pressure on a structure and a velocity load on an acoustic model. Parameters are listed in Table 2-27. The exponential step wave is useful for verification problems in scattering, but is not realizable physically. The pressure definition is similar to the plane wave, but employs a Heaviside step function, $H(t - t')$, where $t' = \frac{\hat{d} \cdot [\vec{x} - \vec{x}_o]}{c_o}$.

$$P = P_o \cdot e^{-\beta \cdot (t - t')} H(t - t') \quad (2.5)$$

A standard planar step wave function can be defined by using $\beta = 0$. This is the default behavior if no beta parameter is specified.

Keyword	Values	Description
type	planar_step_wave	identifier keyword
Direction	3 reals	direction of the wave, $\vec{d}$
material	string	acoustic material
origin	3 reals	wave origin, $\vec{x}_o$
beta	real	exponential decay factor, β

Table 2-27. – Planar Step Wave Parameters.

2.8.13. Spherically Spreading Wave

A spherically spreading wave, keyword=**spherical_wave** , computes the response of a point source excitation in an acoustic medium. The function applies both a pressure on the structure and a velocity load on an acoustic model. Parameters are listed in Table 2-28. Figure 2-17 illustrates the geometry.

Keyword	Values	Description
Type	spherical_wave	identifier keyword
origin	3 reals	wave origin, $\vec{x}_o$
reference_location	3 reals	reference location $\vec{R}$
material	string	acoustic material (alternate to C0)
pressure_function	string	new function for user supplied pressures

Table 2-28. – Spherical Wave Parameters.

A spherical wave is used only in transient dynamics analyses. An example input is described in input 2.7. Function 1 in the example defines the spherical wave function, which describes the geometry of the loading. The time history of the loading is referenced in the function, function 11 in the example, must be a simple function of time. It could be a linear function, a runtime compiled function or a table. It cannot be a function of space and time.

```

LOAD 10
  sideset 1001
    acoustic_vel 1.0
    function = 1
  sideset 50000000
    pressure 1.0
    function = 1
END

FUNCTION 1
  type = spherical_wave
  origin = 0 1000 0
  pressure function = 11

```




Figure 2-17. – Spherical Wave Geometry.

```

material = 1000    //material for acoustic medium
END

FUNCTION 11
  Data 0.0      0.00000
  Data 1e-6     0.00001
  Data 2e-6     0.00002
END

```

Input 2.7. Spherical Wave Example

2.8.14. Undex Structural Acoustic Loads

For Navy scattering applications, the “multicycle_bubble” and “Undex_shock_wave” functions provide a numerical function for analysis of exterior shock loading. The parameters of the loading are listed in Table 2-29. Details of the theory and implementation are available from the Navy Surface Warfare Center, Carderock Division (NSWC/CD). An example input is shown in input 2.8.

```

function 67
  type = hicks_bubble

```

Keyword	Values	Default	Description
type	<i>single_decay</i> <i>double_decay</i> <i>hicks_bubble</i> <i>Undex_shockwave</i>	required	identifier keyword
charge_weight	<i>real</i>	required	in pounds of TNT
charge_location	<i>3 reals</i>	required	explosive location
waterline_depth	<i>real</i>	0	
free_surface_flag	<i>integer</i>	1	
material	<i>string</i>	required	acoustic material

Table 2-29. – Undex Load Parameters. input coordinates are in inches.

```

charge_weight = 10
charge_location = 100.0 0. 50
waterline_depth = 10
free_surface_flag = 1
material = water
end

material water
  acoustic
  c0=4872
  density=62.4
end

```

Input 2.8. Example Hicks Bubble Function Specification

A “free_surface_flag” of one indicates generation of an applicable image source above the surface of the water, where a “free_surface_flag” of 0 indicates a load without a free surface. In this routine, “z” is upwards and normal to the water surface. The depth is the distance below the water surface, i.e.

$$\text{waterline\_depth} = z_{\text{waterline}} - z_{\text{charge}}$$

where z_{charge} is the z component of the charge location. If the free surface flag is not specified, no effects of the surface are included.

One special type of shock wave (contained within the general “Undex” shock wave function definition), is that of a single decay shock wave. The single decay function has the

following simple analytical solution:

$$P(w, r, t) = \begin{cases} 0.0 & t < \text{toa}(r) \\ \frac{P_{\max}(w, r)}{\exp\{(t - \text{toa}(r))/\theta(w, r)\}} & t > \text{toa}(r) \end{cases}$$

Where:

w = charge weight (lb. TNT)

r = standoff distance (ft)

t = time (seconds)

c = sound speed in water (ft/s)

$$P_{\max}(w, r) = k_1 \left(\frac{\sqrt[3]{w}}{r} \right)^{a_1}$$

k_1, a_1 = similitude constants

$$\theta(w, r) = k_2 \sqrt[3]{w} \left(\frac{\sqrt[3]{w}}{r} \right)^{a_2}$$

k_2, a_2 = similitude constants

$\text{toa}(r) = r/c$ = time of arrival

2.8.15. Fluid Structure Interaction

For fluid-structure interaction (FSI) applications, the **FSI** keyword provides a means of applying a prescribed nodal pressure load along the wetted surface. The FSI function is referenced in the **Sierra/SD** input as follows,

```
Loads
  sideset 1
    pressure 1
    scale 1
    function 1
End

Function 1
  type = FSI
End
```

The above input file assumes sideset 1 is the wetted surface. **Sierra/SD** will communicate nodal locations of the sideset to the fluid. These are the locations at which pressures are sent to **Sierra/SD**. Then, **Sierra/SD** calculates a consistent load based on the values at the nodes. Finally, if restarts are needed, “restart=auto” is required in the solution section.

Sierra/SD also supports two-way coupling for Fluid-Structure interaction. Interpolation from structural nodes to fluid nodes and from fluid nodes to structural nodes is

implemented and unit tested. Figure 2-18 shows the infrastructure for FSI. There are many details to use of this coupling, such coupling is turned with the “acoustic_coupling” transient solution option.

Coupling to the Sierra code Fuego through Sierra Toolkit Transfers is currently being implemented. Users of current sprint releases must be careful to not use the “Fuego_coupling” transient solution option, as this is only partially implemented.




Figure 2-18. – Fluid-Structure Interaction (FSI) Infrastructure.

2.8.16. Blending

In some applications, a combination of functions is necessary to accurately describe the loading. For example, a blended shock/bubble function is applied to a ship surface. The shock wave reaches the surface first, and is the effective loading until the bubble function arrives. Once the bubble arrives, the shock may be ignored, as it may not propagate through the cavitated region. This is illustrated for the first crossing case in Figure 2-19, and a typical input for this is shown in input 2.9. Parameters of the blended function are shown in Table 2-30. Several examples for the more general Nth crossing case are shown in Figure 2-20, with a typical input shown in input 2.10.

Keyword	Parameter	Description
Method	<i>string</i>	“first_crossing”, “second_crossing” or “Nth_crossing”
N	<i>int</i>	(for “Nth_crossing”) number of crossings
Primary Function	<i>string</i>	shock function
Secondary Function	<i>string</i>	bubble function

Table 2-30. – Blended Function Parameters.




Figure 2-19. – Illustration of first crossing blended function.

```

Loads
  sideset 55
    pressure 1
    function 55
  end

function 55          // blended shock/bubble
  type = blended
  method = first_crossing
  primary function 551
  secondary function 552
end

function 551
end

function 552
end

```

Input 2.9. Blended First Crossing Function Example




Figure 2-20. – Illustration of Nth crossing blended functions.

```

Loads
  sideset 55
  pressure 1
  function 55
end

function 55          // blended shock/bubble
  type = blended
  method = Nth_crossing
  N = 2
  primary function 551
  secondary function 552
end

function 551
end

function 552
end

```

Input 2.10. Blended Nth Crossing Function Example

2.8.17. Matrix-function

This section provides for input of a matrix function as is used in a cross correlation matrix for input to a random vibration analysis. In the limit of a single input these reduce to a single function. Note that a matrix-function can have arbitrary symmetry and can be complex. An important feature of the matrix-function is that each entry of the matrix is a function of frequency (or time).

The Matrix-Function is illustrated in the following example.

```
Matrix-Function 1
  name 'cross-spectral density'
  symmetry=Hermitian
  dimension=2x2
  nominalt=20.1
  data 1,1
    real function func11 scale 1.0
  data 1,2
    real function func12
    imag function func121 scale -3.0
  data 2,2
    real function func22 scale 0.5
End
```

Matrix functions have the following parameters.

NAME allows you to optionally enter a string by which the matrix-function will be identified in subsequent messages.

SYMMETRY identifies the matrix symmetry. Options are “none”, “full”/“symmetric”, “asymmetric” and “Hermitian”. If the matrix is not square, only “none” can apply. The default for this optional parameter is “symmetry=none”.

DIMENSION specifies the dimension of the matrix. If not specified, it defaults to 1×1 . The dimension is specified as the number of rows, an “x” and the number of columns. No space should be entered between the terms.

DATA Each data entry specifies one entry in the matrix-function. The **data** entry must be immediately followed by the matrix location specified as a row, column pair. Again, no spaces may be inserted in the location entry. The **data** parameters uses two keywords.

- “real” identifies the real component of the entry. It must be followed by a function reference (see Section 2.8), and optionally by a scale factor.
- “imag” identifies the imaginary component of the entry. It must be followed by a function definition, and an optional scale factor.

nominalt `nominalt` Used only for echoing the matrix values. If *input_summary* is specified as an “ECHO” option (see Section 7.7) general information from the matrix function are written to the log file (the *.rslt* file). If, a *nominalt* entry also exists, then the matrix entries are written for that nominal time (or frequency). Only one such output can be specified. It provides a means of checking the input to assure the matrix values are correct at a single time (or frequency) value.

2.8.18. Alternate Table Interface

An alternate **Table** input is provided. See Section 2.8.19 for details about tables. However, for many inputs, the individual specification of each function on each matrix element is both tedious and inefficient. Table input is provided primarily for efficiency reasons. It cannot be mixed with the individual methods, i.e. if the table keywords are used, the “data” keyword must not be used.

Application of table input to matrix-functions requires three tables for:

1. Real valued data
2. Imaginary valued data
3. A table which associates each nonzero row and column of the matrix-function with appropriate rows of the real valued and imaginary valued data.

Three new keywords are introduced:

Real Table which is a two-dimensional table containing all the real valued entries for each entry in the matrix. Each column contains the frequency data for that entry.

imag Table which is a two-dimensional table containing all the imaginary (complex valued) entries for each entry in the matrix. Each column contains the frequency data for that entry.

Table Index which is a two-dimensional table providing a map from the matrix elements to the data columns in the real and imag tables. This index is a 4 column table. Columns 1 and 2 are the row, column index of the matrix-function. Column 3 is the row index of the real data, while column 4 is the row index of the imaginary data. If the value in column 3 or 4 is zero then the corresponding data is zero See the example in input 2.11.

Note that the table entry has a fixed step, so the same number of values must be provided in each column.

```
Matrix-Function 1
  name 'spectral density'
  dimension=2x2
  symmetry=Hermitian
  real Table RealData
  imag Table imagData
```

```

    Table INDEX idx
END

Table RealData
    size=3 550 // 550 freq samples, 3 matrix locations
    delta=1 0.5
    datafile='RealData.txt'
END

Table imagData
    size=1 550 // 550 freq samples, 1 matrix location
    delta=1 0.5
    datafile='imagData.txt'
END

Table idx
    size=3 4
    rowfirst // transpose matrix for simpler input
    dataline
    // row col real imag
    1 1 1 0 // row=1, col=1 with 1st real data, no imag
    1 2 2 1 // row=1, col=2 2nd real data, 1st imag
    2 2 3 0 // row=2, col=2, 3rd real data row, no imag
END

```

Input 2.11. Example Input for a Matrix-Function using Tables

Each matrix entry in the matrix-function must reference a row of a two-dimensional table. In the table columns contain the frequency response for that entry. The number of rows required for each table depends on the matrix symmetry and on the index in the “Table Index”.

2.8.19. Table

A (1 dimensional) table is implemented by including a file with one value per line. 1-dimensional tables have identical behavior to linear functions (section 2.8.2) while typically being much faster and more memory efficient, especially for a large number of data points.

A small section in the main input deck specifies the initial time and time increment. The data must be sampled at a uniform interval.

Tables are used by being referenced in other sections of the input deck. Tables offer support for multi-dimensional data (up to dimension 4). Note that tables of dimension greater than 1 are complicated.

Each Table includes a number of required and optional parameters, as shown below.

Table 2-31. – Table Section Options.

Parameter	Default	Description
dimension	<i>optional</i>	number of dimensions in the table
size	<i>required</i>	table size in each direction
datafile	<i>required</i>	ASCII file containing the values at each point
dataline	<i>required</i>	flag indicating that all data values will follow.
origin	zero	origin of the table (for scaling)
delta	1	interval between points in each direction
rowfirst		transpose data on input

The **dimension** identifies the table shape. For example, **dimension=2** indicates a table of xy values. If this parameter is not defined, the dimension will be automatically inferred from the number of entries in **size**. If it is defined, it must match the inferred value, and thus only serves as a check on the table size.

The **size** parameters indicate the individual table hyper-cube dimensions. For example, in a table of **dimension=2**, the **size** parameter indicates the number of rows and columns in the table. The total number of entries is the product of all the terms in the size.

The text file containing the table data values is specified using the **datafile** parameter. Data values are separated by white space. The layout of the file is not important, but the order is important. The first dimension cycles the fastest. For a **dimension=2** table, the file list begins with the entries for column 1. The number of entries in the file must match the table size. Comments are not permitted in the **datafile**, but white space is permitted.

The **dataline** parameter indicates that the tabular data is included in this file following the parameter. If **dataline** is specified then **datafile** must *not* be specified. The format is identical to the datafile. For performance reasons, it is best to use **dataline** for smaller data sets, and **datafile** for larger.

The **rowfirst** is provided to transpose the data on input. It applies only to 2D tables. If this keyword is present, then the table values will be interpreted as if the table had been transposed.

Both the **origin** and the **delta** parameters are optional values provided for interpolation. The implicit integer entries of the table are converted to real values for function evaluation by use of these parameters.

Function evaluations within the range of the table can be linearly interpolated. The range in each direction is determined by the following.

$$\text{origin}_i < \text{range}_i < \text{origin}_i + (\text{delta}_i \cdot \text{size}_i) \quad (2.6)$$

Evaluations of the table for regions outside the valid range will use of the value of the nearest data point, just as with linear functions.

In contrast to a **function** (see Section 2.8), tables require memory only as needed. All processors store the full input file in memory. However, tables can store a large amount of data in the **datafile**. This file is opened and data is read from it only as needed. For this reason, tables are preferred over functions when only a few processors may need access to a large amount of data. Tables are the only option when a function of more than one variable is required.

An example of a two-dimensional table definition is shown below.

```
Table example-2D-table
  dimension=2
  size      = 200 300      // note: don't put in an x
  origin     1.0 0.0      // optional. defaults to 0 0
  delta      1.0 0.9      // optional. defaults to 1 1
  datafile   'multidimensionaltable.txt'
END
```

2.9. *Multipoint constraints*

A multipoint constraints (MPC) is a constraint equation applied directly to the stiffness matrix. Some analysis codes treat them as pseudo elements. An MPC is not an element, and is inaccessible through **Exodus**.

An MPC is a displacement constraint

$$\sum_i^n c_i u_i = r,$$

The (nonzero) coefficients c_i are real. The number of (nonzero) coefficients is assumed to be approximately 1. The u_i are displacement of degrees of freedom. By default, r vanishes.

Unlike many Finite Element programs, **Sierra/SD** does not support user specification of constraint and residual degrees of freedom (DOF). In serial solvers the partition of constrained and retained degrees of freedom is performed simultaneously by Gauss elimination with full pivoting so the constrained degrees of freedom are guaranteed to be independent. For GDSW the constraints are specified as Lagrange multipliers which involves no such partitioning. Redundant specification of constraint equations is handled by elimination of the redundant equations and issue of a warning. User selection of constrained DOF in NASTRAN has inconvenienced analysts who must ensure that the constrained DOF are independent and never specified more than once.

Each MPC is specified in the input file with a section descriptor. Note that a separate section is required for each equation (or degree of freedom eliminated). An optional

coordinate system may be specified on the input ²; see section 2.7. The MPC will be stored internally in the basic coordinate system (coordinate frame 0). The input consists of a triplet listing the node or nodeset, a degree of freedom string, and the coefficient of that degree of freedom. The degree of freedom strings are x , y , z , Rx , Ry , Rz . They are case-insensitive. If the global ID of the node in the MPC does not exist in the model, the code will exit with a fatal error.

GDSW is required for inhomogeneous MPCs. The solution method must not be QEVP. **Sierra/SD** will exit with a fatal error if a model including non-homogeneous MPCs violates either of these requirements.

The MPC can apply some general commands that apply to the whole MPC as shown in Table 2-32. Additionally, the MPC can include multiple lines that define the MPC equation entries as shown in Table 2-33.

Table 2-32. – General MPC commands.

Keyword	Type	Description
coordinate	<i>string</i>	Optional coordinate frame for MPC
rhs	<i>real</i>	Optional right-hand side constant

Table 2-33. – MPC Equation lines.

Command Line	Description
<i>integer x/y/z/rx/ry/rz real</i>	Equation entry using single global node id
<i>nodeset nodeset_id x/y/z/rx/ry/rz real</i>	Equation entry using nodeset with EXACTLY one node

In this first example the MPC is defined with respect to a coordinate system **dir1**. The displacement at the x degree of freedom of node 4 is constrained to be equal to the average x degree of freedom of nodes 2 and 3 plus 5, i.e. $u_{(4,x)} - 0.5 * u_{(2,x)} - 0.5 * u_{(3,x)} = 5.0$.

```

MPC
  coordinate dir1
  4 x 1.0
  2 x -0.5
  3 x -0.5
  rhs 5.0
END

```

²At this time, all the nodes in an MPC must be associated with the *same* coordinate system.

In this next example, the x degree of freedom of the node in the nodeset 101 is constrained to be equal to the y degree of freedom of the node in the nodeset with name `aftPlate`. Note that both nodesets must contain exactly one node.

```
MPC
  nodeset 101 x 1.0
  nodeset aftPlate y -1.0
END
```

IMPORTANT

Constraints are handled in various ways by the linear solvers. In the serial solver, the dependent degrees of freedom are eliminated before the matrices are passed to the solver. In parallel, we use Lagrange multipliers to handle the constraints. There is currently no user control of constraint handling methods.

Note also that there are practical differences between rigid elements (described in the following sections) and constraint equations that are nominally identical. For parallel solutions, we are currently using an augmented Lagrange type solution method with the rigid links. This means that terms are added to the stiffness matrix in parallel with the constraints. In most cases, this renders the matrices positive definite, and increases robustness and solution performance with no penalty for accuracy. Thus, rigid links are recommended whenever possible in parallel solutions.

Replacing rigid links with stiff beams may cause ill conditioning and simulation inaccuracies.

3. SOLUTION CASES

Sierra/SD supports a wide variety of different analyses or solution methods. Input consists of an **Exodus** mesh file and a text input file. Solution methods are specified in the text input file in the solution section.

The **Solution** section defines the type of physics to simulate. Analysis types are shown in Tables 3-34, 3-35, 3-36, 3-37. Relevant options are given in the detailed description of each solution case. Also, general options are shown in Table 2-10 and are described in Section 2.4.

Table 3-34. – Eigen Solution Types.

Solution Type	Description
eigen	Eigen solution to extract natural vibration modes of K,M
aeigen	Eigen solution with Anasazi
buckling	Eigen solution solving for buckling modes
CBR	Craig-Bampton reduction for creation of superelements
geometric_rigid_body_modes	Exact analytic definition of the rigid body modes
blk_eigen	Eigen solution on a block by block basis
residual_vectors	Modal truncation augmentation
eigen-inverse	Inverse solution to find material properties to produce given eigen solution
largest_ev	Largest eigenvalue of K,M
qevp	Solution to quadratic eigenvalue problem

Table 3-35. – Modal Solution Types.

Solution Type	Description
ModalFrf	Frequency response using modal displacement or modal acceleration
modalranvib	Random vibration using modal superposition
modalshock	Shock response spectra using
modaltransient	Transient analysis using modal superposition
ModalFrf-inverse	Inverse solution to find load or power spectral density (PSD) to produce given modal frf
modaltransient-inverse	Inverse solution to find load to produce given modal transient

Table 3-36. – Direct Solution Types.

Solution Type	Description
directfrf	Direct computation of frequency response functions
NLstatics	Nonlinear static solution
NLtransient	Nonlinear transient solution
statics	Static solution
transhock	Shock response spectra using direct implicit transient
transient	Linear transient solution
directfrf-inverse	Inverse solution to find load or material properties to produce given frequency response
transient-inverse	Inverse solution to find load or material properties to produce given transient solution

Table 3-37. – Preprocessing and Postprocessing Solution Types.

Solution Type	Description
receive_sierra_data	Import stress, displacement, and material state from a preload Sierra/SM analysis
checkout	Check input for errors, skips matrix formation and solves
dump	Check input for errors, form matrices, skips solve
CJdamp	Approximate modal damping contributions
ddam	Dynamic design analysis method (U.S. Navy)
preddam	Gather data for use by ddam solution case
gap_removal	Do contact search and apply the contact gap removal algorithm
superposition	Expand superelement results to physical degrees of freedom
tangent	Compute tangent stiffness after a nonlinear load step
tsr_preload	Import data for thermal structural response
fatigue	Postprocess modal random vibration results to predict high cycle fatigue life
MPF	Compute and output modal participation factors
waterline	Determine waterline of a floating structure

3.1. *Defining Solution Cases*

It is important to distinguish the two different kinds of Solution section. A Solution section may do one analysis. In this case a **loads** section is used. Or a Solution section may do multiple analyses. Here each case may specify a **load** section. Here's an example of a Solution section that does one analysis. The eight lowest eigen modes of a structure are requested.

```
Solution
  eigen
  nmodes 8
end
```

A more complex analysis could perform a static preload, followed by computation of updated tangent stiffness, and then a linearized eigen analysis with this input:

```
Solution
  case 'NonlinearStatics'
    nlstatics
    load=10
  case 'tangent'
```

```
tangent
case 'eig'
  eigen
  nmodes 8
end
```

3.2. Multicase Solutions

The solution methods of Tables 3-34, 3-35, 3-36, and 3-37 may be a part of a multicase solution. Using one input deck, it is possible to run any valid sequence of simulations. For example, a static preload, followed by computation of updated tangent stiffness, and then a linearized eigen analysis.

In a multicase analysis, each solution step is delineated by the **case** keyword. The word **case** indicates the beginning of a solution case. Cases are run sequentially from top to bottom. Indenting to clarify this hierarchical relationship is recommended, but not required. Each **case** keyword must always be followed by a label. The label is used to set the output file name and is used in a variety of informational, warning, and error messages. Note: the label name for each case must be unique.

The results and outputs from one case will be used to drive subsequent cases. For example the modes computed by an **eigen** solution case will be used in a subsequent **modaltransient** solution case. As another example the stresses and displacements read in from **receive_sierra_data** solution case will be used to define the initial geometry and geometric stiffness for a subsequent eigen solution case.

In a multicase solution, the system matrices (mass, stiffness and damping) will typically be computed once. Matrix updates between solutions may be specified by selecting the **tangent** keyword (see Section 3.30).

3.2.1. Multicase Options

A solution type has specific options. Only transient solution cases have time steps for example. On the other hand, Table 3-38 shows options that may apply to each solution case. They may be specified either above the case control sections, or within the section. The specification above the case control section is the default value. Specifications within the case sub-blocks apply to that sub-block. In the example given in 3.2.2, the **restart** options are thus “none” for most sub cases, but “read” for the eigen analysis and **auto** for the linear transient.

The default load applied to each solution case in a multicase solution is prescribed by the **loads** block if present in the input file. Loading prescribed by a **loads** block can be overridden by specifying a **load** in an individual solution case in the multicase solution. Only one **load** can be applied in each solution case. In the example given in 3.2.2, load

Table 3-38. – Multicase Options.

These parameters may be specified as defaults above the case specifications, or they may be specified for each subcase to which they apply.

Option	Description	Options
restart	Restart options	see Section 2.4.2
solver	selection of solver	see Section 2.4.3

‘10’ will be applied during the nonlinear statics solution case (case ‘NonlinearStatics’). For more information on **loads** see Section [6.3](#) and for **load** see Section [6.3.1](#).

3.2.2. Multicase Example

In the example which follows, a nonlinear statics computation is followed by a tangent stiffness matrix update. The updated matrix is then used in an eigen analysis. Two sets of **Exodus** output files will be written. Output from the statics calculation will be in files of the form *'example-nls.exo'*. Eigen results will be in the form *'example-eig.exo'*. The **tangent** solution normally produces no output in the **Exodus** format.

Transient solution cases are run sequentially in time. For the example below 'trans1' will start at time 0.0 and step through 100 steps of at a step size of $1e-8$ and then through 4000 steps at a step size of $1e-6$. The final time value of case 'trans1' will be $10^{-8}10^2 + 10^{-6} \times 4 \times 10^3 = 0.004001$. Case 'trans2' will start at 0.004001 and run an additional $10^{-4} \times 210^2$. This will end at time 0.024001.

```
Solution
  restart=none
  title='example multicase'
  case 'nls'
    nlstatics
    load=10
  case 'tangent'
    tangent
  case 'eig'
    eigen
    restart=read
  case 'trans1'
    transient
    restart=auto
    time_step 1e-8 1e-6
    nsteps    100  4000
    flush 50
    rho=0.9
    load=20
  case 'trans2'
    transient
    restart=auto
    time_step 1e-4
    nsteps    200
    flush 10
    load=20
end
```

3.3. Checkout Solution Case

Parameter	Type	Default	Description

Table 3-39. – Checkout Solution Case Parameters.

The **checkout** solution method tests out various parts of the code without forming the system matrices or solving the system of equations. This solution method may be used to check for syntax errors, use of nonexistent blocks and sidesets, or completeness of input. Generally checkout would be used to fix all such errors on a local serial platform before allocating expensive resources or waiting in a queue to run a full solution. The checkout solution case performs the same type of input checking as the `--check-syntax` command line argument.

3.4. CJdamp Solution Case

Parameter	Type	Default	Description

Table 3-40. – CJdamp Solution Case Parameters.

The **CJdamp** solution provides a method to compute the equivalent modal damping terms introduced from material damping in lightly damped viscoelastic materials. It is based on a development by Conor Johnson *et al.*³⁷ **CJdamp** is an approximate method which assumes that the mode shapes and frequencies are not modified by the damping. The modal damping is related to the fraction of energy in block.

The **CJdamp** method is effectively a post processing step following an eigen analysis and must be run in a multcase solution. For each of the modes in the eigen analysis, a strain energy is computed on an element basis. These are summed at the block level.

$$SE_j^i = \sum_{\text{elem}}^{\text{in block j}} \phi_i^T K^{\text{elem}} \phi_i \quad (3.1)$$

The total strain energy TSE^i is the sum of the strain energy contributions in mode i from all blocks. We define the block strain energy ratio for mode i as,

$$R_j^i = SE_j^i / TSE^i \quad (3.2)$$

The **CJdamp** contribution for the modal damping of mode i , is given by,

$$\zeta_i = \frac{1}{2} \sum_j R_j^i \eta_j(f_i) \quad (3.3)$$

$\eta_j(f_i)$ is the **CJetaFunction** contribution from block j evaluated at the natural frequency of mode i . More information is available in Section 4.4.9.

*Note that cases following the **CJdamp** solution will include the computed damping as part of their damping calculation.*

Example,

```
Solution
  case eig
    eigen nmodes=30
  case johnson
    CJdamp
  case frf
    ModalFrf
end
```

3.5. **Craig-Bampton reduction Solution Case**

Parameter	Type	Default	Description
nmodes	<i>integer</i>		Number of fixed interface modes
shift	<i>Real</i>	-1.0e6	Shift to apply to matrix system to allow solving singular systems. See Section 3.10.2
correction	<i>node/values/vectors</i>	values	Correction method for rigid body modes
RbmDof	<i>string</i>	123456	Defines which rigid body modes to which correction=vectors applies
ModalFilter	<i>string</i>	none	Modal filter to define modes to retain. See Section 3.10.2.2

Table 3-41. – CBR Solution Case Parameters.

It can be advantageous to reduce a model to its interface degrees of freedom. This is important in satellite work, where the model of the satellite may be much larger than the

model of the remainder of the missile. Reduction of the satellite model to a linearized, Craig-Bampton model makes it possible to share the dynamic properties of the model without requiring details of the interior. There are many types of component mode synthesis techniques (or **CMS**), of which the Craig-Bampton approach is one of the more popular. In this approach the model is reduced to a combination of fixed interface and constraint modes. The fixed interface modes are eigenvectors of the system with all interface degrees of freedom clamped. A constraint mode is the deformation that results if one interface degree of freedom receives a unit displacement, and all other interface degrees of freedom are zero. There is one constraint mode per interface degree of freedom.

The **Craig-Bampton reduction** solution reduces an entire structural model to its reduced system and transfer matrices. Options are listed in the table below, and correspond to the parameters required for an eigen analysis (Section 3.10). In addition, a **CBModel** section must be defined elsewhere (see Section 3.5.1). Any boundary conditions specified are applied before reducing the model.

We note that sensitivity analysis can be performed in **Craig-Bampton reduction** , though the process is somewhat different from other types of sensitivity analysis. Section 3.5.1 contains more information about sensitivity analysis in Craig-Bampton models.

The method will write system matrices and general information. Each of the parameters is described below.

nmodes: The CB model is composed of fixed interface modes and constraint modes. The number of constraint modes is determined by the interface. **nmodes** selects the number of fixed interface modes. The fixed interface modes are eigenvectors of the interior of the structure, and provide a basis for internal deformation. Any number of these modes may be specified. Typically, frequencies up to about twice the system frequency are required for accuracy.

correction: As shown in the theory manual, the null space of the stiffness matrix is determined by the sum of two large terms: $\kappa_{cc} = K_{cc} + K_{cv}\psi$. With parallel iterative solvers, it may be difficult to determine this quantity as accurately as desired. In particular, it is possible for errors in the solver to render the reduced matrices negative definite, which can cause instability in subsequent transient analysis. It is strongly recommended that low solution tolerances be used in developing CB models. In addition, the matrix may be post processed to correct these errors. The post processing options are as follows:

none no correction will be applied.

values (default) no corrections will be made to the eigenvector space, but the negative eigenvalues will be adjusted to zero.

vectors This option is available in models that have no boundary conditions, so that the CBR model is floating and has six rigid body modes. If boundary conditions are present, there is a fatal error. Additionally, if the model does not have six rigid body modes for another reason (for example a gap contact constraint that impedes rotation) then the correction=vectors option is not well posed, should

not be used, and could cause a serious degradation of model behavior. The `correction=vectors` will compute Zero energy eigenvectors geometrically (exactly), and these are used to correct both the eigenvalues and the eigenvectors. This is more involved than correcting the eigenvalues alone, but it is not a significant computational cost, and can improve the usefulness of the resulting model.

If `correction=vectors` is selected, one may also optionally determine which zero energy modes are required. This word `RbmDof` activates this. It is followed by a string indicating which dofs are active on the interface. The string contains the numbers 1 through 6, where 1 represents translation in the x coordinate direction. These specifications apply in the basic coordinate frame.

```
solution
  CBR
    nmodes=20 shift=-4e6
    correction=vectors
    RbmDof='123'
end
```

Inertia Tensor for Craig-Bampton Reduction. A reduced inertia matrix, $\mathbf{I}_{\mathbf{v}\mathbf{v}}$ may be output from a CBR (Craig-Bampton Reduction) analysis. The Inertia Tensor may be used to apply initial conditions to the superelement in some applications. The input file syntax is described in the **CBModel** section, 3.5.1. Φ is the matrix of mode shapes used for the CBR analysis. It consists of both fixed-interface modes and constraint modes. $\mathbf{I}_{\mathbf{v}\mathbf{v}}$ is defined by

$$\mathbf{I}_{\mathbf{v}\mathbf{v}} = \Phi^T \mathbf{R},$$

The number of rows in $\mathbf{I}_{\mathbf{v}\mathbf{v}}$ is the number of CBR modes, and the number of columns is the number of rigid body vectors. For example, for the three translational rigid-body modes and assuming three degrees-of-freedom per finite element node,

$$\mathbf{R} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ \vdots & \vdots & \vdots \end{bmatrix}.$$

Mass Inertia Matrix for Craig-Bampton Reduction. The Mass Inertia Matrix is used in some applications to apply a load to the interior degrees of freedom of a model. ³ The mass inertia matrix is defined as,

$$I_m = \Phi^T M \mathbf{R}.$$

³Neither the inertia tensor, nor the mass inertia matrix may be applied in **Sierra/SD**. They are output quantities.

where M is the mass matrix of the unreduced model.

Note the following limitations for the CBR method.

- *MPCs may no share nodes with interface nodes.*
- *The entire reduced order model and associated transfer matrix must fit into memory. On a parallel machine, this memory is required on every processor. The dimension of the model is the sum of the numbers of constraint and fixed interface modes.*
- *Static solutions with all interface degrees of freedom clamped are part of the reduction process. If the interface dofs do not fully constrain the system, then the linear system may be singular. In such cases the solution is not reliable. Due diligence includes verifying the reduced order eigen solution against the full system. One may also compare the retained mass.*

3.5.1. CBModel

The **CBModel** section provides a method of specifying information related to a Craig-Bampton model reduction of the entire structure. It is required by the CBR method (section 3.5).

The “interface” is that portion of the model which will interface to the external structure. The interface is defined by collections of nodes specified as nodesets or sidesets. After eliminating boundary conditions, the active degrees of freedom on the nodes become the interface.

The input deck keywords shown in Table 3-42 are described below.

nodeset: The **nodeset** keyword specifies the nodes to be placed in the interface.

Nodesets are defined in the **Exodus** file. A nodeset ID or name must follow the nodeset keyword. Alternatively, a list of nodesets (in MATLAB type format) can be specified. This is identical to the **history** file definition of Section 7.3, and follows the rules for integer lists detailed in Section 2.1.

sideset: A **sideset** may also be used to specify the interface nodes. Any number of nodeset and sideset combinations are allowed. The interface is the union of all such entries.

format: The preferred format is the **netcdf** , format. This is a superset of the **Exodus** format. It is the format that must be used if the reduced model is to be inserted into another **Sierra/SD** model as a superelement. The **DMIG** format is for use with NASTRAN. It contains only the reduced system matrices (no maps, coordinates, etc). The MATLAB format is a convenience.

Table 3-42. – CBModel Parameters.

Keyword	type	Description
nodeset	<i>int/string/list</i>	Exodus nodeset name(s) and/or id(s).
sideset	<i>int/string/list</i>	Exodus sideset name(s) and/or id(s).
format	<i>string</i>	specifies the output format. MATLAB - MATLAB.m format DMIG - NASTRAN DMIG format DMIG* - NASTRAN long DMIG format netcdf - netcdf format [†]
file	<i>string</i>	specifies the file name for output.
GlobalSolution	<i>bool</i>	‘yes’ to compute the eigen solution of the reduced system.
inertia_matrix	<i>bool</i>	‘yes’ to compute the inertia tensor
sensitivity_method	<i>string</i>	specifies the method to compute CBR sensitivities. constant_vector - constant vector method finite_difference - finite difference method
spoint_offset	<i>integer</i>	offset for spoints in DMIG format

[†]The netcdf format is the database upon which exodusII is built. Seacas tools, Matlab and Python can all read netcdf files.

file: The **file** keyword is required to specify the output file name.

GlobalSolution: As a convenience, we will optionally compute the eigenvalues of the reduced system. It is strongly recommended that these values be compared with the eigenvalues of the full system to ensure that the model has converged over the frequency of interest. Moreover, users are strongly advised to review section Craig-Bampton Model Reduction of **Sierra/SD How To**⁵³ before computing superelements.

inertia_matrix: The inertia matrix defined in Section 3.5 is optionally computed and written to the super-element files with i the reduced mass and stiffness matrices.

sensitivity_method Currently, the constant vector and finite difference methods are available for computing sensitivities for Craig-Bampton reduction. The default is the constant vector method.

spoint_offset NASTRAN DMIG output defines “spoint” variables that store generalized degrees of freedom. The identifiers the generalized output variables must be unique. The range of identifier output may be specified with this option. By default, “spoint_offset” is 10000 which means spoints are numbered as 10001, 10002, 10003, etc.

Specify “displacement” in the **outputs** section (7) to output the constraint modes and fixed interface modes that were used as a basis to generate the reduced order system to the **Exodus** file. First the fixed interface modes are output, followed by the constraint modes. The eigenvalues of the fixed interface modes correspond to the mode shapes. For the

constraint modes an integer index replaces the eigenvalues. These modes may be visualized and evaluated using any of the standard tools.

Data in Table 3-43 will be written to a file. The Output Transfer Matrix (or OTM) depends on data in the **history** section (see Section 7.3). Specifically, the output nodes and elements, and the output variables are specified in the history file *as if they were to be output to a history file*. For simplicity, and because the OTM describes a linear transfer matrix, only a limited subset of results are provided. In particular, displacements and the natural strains and stresses may be written. The transfer matrix provides the following computation.

$$\begin{aligned} \begin{bmatrix} u \\ \epsilon \\ \sigma \end{bmatrix}_{out} &= \begin{bmatrix} \text{OTM} \end{bmatrix} \begin{bmatrix} q \\ u_{int} \end{bmatrix} \\ &= \begin{bmatrix} \Phi_u & \Psi_u \\ \Phi_\epsilon & \Psi_\epsilon \\ \Phi_\sigma & \Psi_\sigma \end{bmatrix} \begin{bmatrix} q \\ u_{int} \end{bmatrix} \end{aligned}$$

Here q is the amplitude of the internal constraint modes (typically computed in the next level analysis), and u_{int} is a vector of interface displacements. The fixed interface modes (eigen modes of a clamped boundary) are represented by Φ , and the constraint modes by Ψ .

The left-hand side vectors represents internal results (displacement, strain and stress) which are computed from the interface results. Any of the output results may be omitted, and the OTM will retain only nonzero components. For example, if only displacements are required, the matrix reduces to $[\Phi_u \Psi_u]$. The OTM matrix is a rectangular matrix, and it is typically full. An example **CBModel** section follows.

```
CBMODEL
  nodeset=1:2          // nodes from nodeset 1 and 2
  format=netcdf         // use a netcdf format file
  file='junk.ncf'
END
```

The reduced inertia matrix **I** for Craig-Bampton Reduction defined in Section 3.5 may be computed and written to the super-element files with the reduced mass and stiffness matrices. **I** can be written to the results file in either netcdf, MATLAB or DMIG format. In the **CBModel** section

```
CBMODEL
  nodeset 1
  format = netcdf
  file = cbmodel.ncf
  GlobalSolution = yes
  inertia_matrix = yes
END
```

Table 3-43. – Data output for Craig-Bampton Reduction.

Variable	Description																
NumC	number of constraint modes																
NumEig	number of fixed interface modes																
Kr	Reduced stiffness matrix.																
Mr	Reduced mass matrix.																
Cr	Reduced damping matrix. Only available for dashpots and block proportional damping.																
cbmap	A two column list providing a map from each interface degrees of freedom to the node and coordinate direction of the global model. The first column of this list is the node number (1:N) in the structure. The second column indicates the coordinate direction as follows. <table> <tr> <th>Number</th><th>Description</th></tr> <tr> <td>1</td><td>x</td></tr> <tr> <td>2</td><td>y</td></tr> <tr> <td>3</td><td>z</td></tr> <tr> <td>4</td><td>Rotation x</td></tr> <tr> <td>5</td><td>Rotation y</td></tr> <tr> <td>6</td><td>Rotation z</td></tr> <tr> <td>7</td><td>acoustic pressure</td></tr> </table>	Number	Description	1	x	2	y	3	z	4	Rotation x	5	Rotation y	6	Rotation z	7	acoustic pressure
Number	Description																
1	x																
2	y																
3	z																
4	Rotation x																
5	Rotation y																
6	Rotation z																
7	acoustic pressure																
OutMap	The “cbmap” has the same number of rows as Kr or Mr. A map of the nodes in the output transfer matrix. OutMap(i) is the global node number for each node in the output. There are always 6 rows of output for each node. Thus, OutMap(1) corresponds to rows 1 through 6 in the OTM.																
OTM	Output Transfer Matrix to provide a transfer function from the interface dofs to internal degrees of freedom or other results.																
OutElemMap	A map of the <i>elements</i> in the output transfer matrix, OTME. OutElemMap(i) is the global element number for each element in the output. There are always 6 rows of output for each element.																
OTME	Output Transfer Matrix to provide a transfer function from the interface dofs to internal elements.																




Figure 3-21. – An initial analysis using CBR can be applied to reduce a complex component to much smaller matrices. In subsequent analyses the superelement replaces the complex component in the system analysis. There is little loss of accuracy, but significant computational benefit.

the `inertia_matrix = yes-no` line requests the output of the inertia tensor. The default value of `inertia_tensor` is no.

NOTE:

*The OTM output capability permits an analyst to output the reduced order model of the entire structure for use in another code that supports superelements (such as MSC/NASTRAN). In a subsequent release, we will add the capability to input these matrices as a superelement in **Sierra/SD**. At that point one could perform a Craig-Bampton reduction to generate a reduced order model of that portion of the structure. A follow-up analysis could use this as a superelement. See details in [Figure 3-21](#).*

3.5.2. Sensitivity Analysis for Craig-Bampton models

Sensitivity output for Craig-Bampton reduction requires both the **sensitivity** block described in Section 2.6, and the **sensitivity_method**, keyword in the **CBModel** block. The **sensitivity_method** defaults to the constant vector method unless otherwise specified.

The sensitivity output from a Craig-Bampton reduction is different from that typically seen in eigenvalue or transient solutions. In the case of a Craig-Bampton reduction, the sensitivities that are output consist of partial derivatives of the reduced mass and stiffness matrices with respect to the parameters. We give a brief description here, and refer to the CBR Sensitivity Analysis discussion in the section Solution Procedures of the Theory Manual for further details.

The reduced stiffness matrices in Craig Bampton reduction is computed via the transformation

$$\kappa = T^T K T \quad (3.4)$$

where T is the Craig-Bampton transformation matrix, K is the stiffness matrix, and κ is the reduced stiffness matrix. A similar expression can be written for the reduced mass matrix. Sensitivities of κ with respect to a parameter p can be computed with the constant vector and finite difference approaches.

The **sensitivity_method**, approach makes the approximation that the transformation matrix $T = T_o$ does not change with the parameter p . Consequently, the sensitivity can be computed as

$$\frac{d\kappa}{dp} \approx \frac{T_o^T (K(p + \Delta p) - K(p)) T_o}{\Delta p} \quad (3.5)$$

Since T_o represents a truncated modal space, this approximation is not always accurate.

The **finite_difference** method uses forward differences. Given updated system stiffness $K_1 = K(p + \Delta p)$ and transformation matrix $T_1 = T(p + \Delta p)$, direct forward differences are used to evaluate the sensitivity

$$\frac{d\kappa}{dp} \approx \frac{T_1^T K(p + \Delta p) T_1 - T_o^T K(p) T_o}{\Delta p} \quad (3.6)$$

Unlike the constant vector approach, the finite difference method will converge to the exact sensitivities as Δp goes to zero. This is true provided that there are no repeated modes in the system.

One additional complexity that arises in sensitivity analysis of Craig-Bampton reduced models is when there are repeated modes in the transformation matrix T . Since the ordering of repeated modes in any eigenvalue problem is arbitrary, the perturbed transformation matrix T_1 could have a different ordering of the repeated modes than T_o . This would corrupt the difference operation in equation 3.6. However, in the case of the constant vector method this is not an issue since there is only one transformation matrix T_o . For these reasons, in the case where the user suspects repeated modes may be present

in the structure, we currently recommend using the constant vector rather than the finite difference approach.

The output from a sensitivity analysis of a Craig-Bampton model is different from other types of sensitivity analysis. The output is written to either a MATLAB or **netcdf** format, depending on the **format** parameter (see Table 3-42). The default is **netcdf**.

The outputted quantities are the derivatives of the stiffness and mass matrices with respect to the various parameters. Thus, if there were two sensitivity parameters p_1 and p_2 , the output quantities would be

$$\frac{\partial \kappa}{\partial p_1}, \frac{\partial \kappa}{\partial p_2} \quad (3.7)$$

where κ would be the reduced mass and/or stiffness matrix. The dimensions of these sensitivity matrices would be the same as the dimensions of the corresponding reduced mass and stiffness matrices.

The output matrix derivatives given in equation 3.7 are useful for studying how the reduced matrices change with the parameters. These matrix derivatives can also be used in subsequent analysis with the corresponding superelements. For more details, we refer to Section 5.31.

3.6. DDAM Solution Case

Parameter	Type	Default	Description
nmodes	<integer>	10	case 1 (eigen, 3.6.0.1)
shift	<real>	-1.0e6	case 1 (eigen, 3.6.0.1)
MODALFILTER	<integer>	10	case 2 (preddam, 3.6.0.1)
load	<integer>		case 2 (preddam, 3.6.0.1)
remove	<integer list>		case 2 (preddam, 3.6.0.1)
cumulative mef	<real>(6)		case 2 (preddam, 3.6.0.1)
add	<integer list>		case 2 (preddam, 3.6.0.1)
analysis__direction	vertical/ athwartship/ fore_aft		case 3 (DDAM, 3.6.0.1)
ship_type	surface_ship/ submarine		case 3 (DDAM, 3.6.0.1)
mount_type	hull/deck/ shell_plating		case 3 (DDAM, 3.6.0.1)
velocity__coeffs	<real>(4)		case 3 (DDAM, 3.6.0.1)
acceleration__coeffs	<real>(4) (<real>)		case 3 (DDAM, 3.6.0.1)

Table 3-44. – DDAM Solution Case Parameters.

The U.S. Navy Dynamic Design Analysis Method (DDAM) is an established procedure employed in the design of ship equipment and foundations for shock loading requirements. The details of the formulation, specific procedures for application, acceptance criteria, etc., are documented in NAVSEA Report 250-423-30 and NAVSEA 0908-LP-000-3010. Support for performing DDAM calculations, as implemented in the **Sierra/SD** Finite Element Code, is documented both in the **Sierra/SD** User's manual and in the DDAM Primer. The user is expected to be fully familiar with both NAVSEA publications.

DDAM is focused on five main phases: the problem formulation phase, mathematical modeling phase, coefficient computation phase, dynamic computation phase, and evaluation phase. DDAM as implemented in **Sierra/SD** will focus primarily on the *evaluation* phase. This phase assesses modal analysis, modal filtering, displacement, velocity, acceleration, force, von Mises stress calculations and all NRL sums.

DDAM does have a few limitations. The equipment to be analyzed must be represented as a linear elastic system with discrete modes. Damping is neglected. For very low frequency (VLF) systems, DDAM may not be appropriate, and, where closely-spaced modes exist, DDAM may produce excessive responses.

Current DDAM capabilities in **Sierra/SD** are as follows:

1. Modal analysis.

Sierra/SD doesn't have the capability to select modes by frequency limits. The user must choose the number of modes, `nmodes`, in case 1 (`eigen`, 3.6.0.1), and continually increase `nmodes` and restart or rerun eigen analysis until the desired frequency level is reached. This step is recommended to be completed before case 2 (`preddam`, 3.6.0.1) and case 3 (DDAM, 3.6.0.1) are included and analyzed.

2. Include modal masses that include at least 1% individually of the total modal mass.

The user must manually add any extra needed modes using `DDAM!add` in the `modalfilter` block.

3. The user may verify `preddam` (3.6.0.1) and DDAM (3.6.0.1) by examining filtered modes, participation factors, modal weights, shock design coefficients and values found in the following text files.

a) `preddam` → `PREDDAM_RESULTS.txt`

b) DDAM → `DDAM_RESULTS.txt`

4. Incorporate a 6g minimum load requirement. This is not a user option.
5. DDAM with complex models that consist of multiple element types.
6. DDAM analysis using Superelement capability.
7. Nodal displacements, velocities, accelerations, and forces written to the output `exodus` file.
8. Element-wise von Mises stresses written to the output `exodus` file.
9. Ship directions must match coordinate direction (`vertical` → Z-direction, `athwartship` → Y-direction, `fore_and_aft` → X-direction).

Directions must match between `preddam` inputs `modalfilter` block, `load` block, and DDAM input `analysis_direction`.

10. DDAM Analysis using symmetric boundary conditions.
11. DDAM Analysis of all models in parallel.
12. Postprocessing available for all DDAM modal variables and all DDAM NRL summed variables (Paraview — graphical, `explore`, `blot`, `exo2mat` — data extraction)

3.6.0.1. DDAM Solution Cases

The DDAM solution is divided into three cases: case 1 (**eigen**), case 2 (**preddam**), and case 3 (**DDAM**).

Case 1 (eigen) is modal analysis already implemented into **Sierra/SD**. The parameters, as they pertain to DDAM, are shown in table 3-44. More information on the keyword **eigen** and its inputs may be obtained in Section 3.10.

Note: The standard **Sierra/SD** specification for eigen analysis is **nmodes=<number>**, where **<number>** is an integer value for the requested number of modes. This capability can compute approximately half the modes of the system at most. For verification purposes a limited capability exists to compute all the eigen modes. This capability runs only in serial, and only on small models. It is selected by **nmodes=all**.

Case 2 (preddam) utilizes the eigenvectors and system mass matrix produced in case 1 (**eigen**) to calculate and filter the modal participation factors, modal weights, individual modal weight percentage, cumulative modal weight, and cumulative modal weight percentage; see Table 3-44

modalfilter is implemented as a part of **Sierra/SD** and may be used as a part of other solution methods. **modalfilter** provides a means of filtering data taken from the modal analysis and the participation factors. More information may be found in Section 3.10.2.2.

Finally, a **load** block should be defined. This **load** block (in **Sierra/SD** DDAM analysis) specifically applies to the value of gravitational loading (-386.4) and its direction must match DDAM analysis direction defined in the **modalfilter** block and in case 3 (**DDAM**).

preddam must be run following **eigen** in order to obtain necessary modal data for **preddam** calculations. **Preddam** may be run without case 3 (**DDAM**).

Case 3 (DDAM) uses filtered eigenvalues and mode shapes from case 1 **eigen** and filtered modal participation factors and modal weights from case 2 **preddam** to calculate shock design coefficients and values, filtered modal displacements, velocities, accelerations, forces, von Mises stresses, and the NRL sum. The inputs are as follows:

Case DDAM may only be run following **preddam** and **eigen**.

An example follows.

```
SOLUTION
  case 1
    eigen
      nmodes = 4
  case 2
    preddam
      MODALFILTER vertical
      load 1
```

```

case 3
  DDAM
    analysis_direction VERTICAL
    ship_type SURFACE_SHIP
    mount_type HULL
    response_type
    velocity_coeffs 1.4 5.2 220.1 12.2
    acceleration_coeffs 1.0 2.0 3.0 4.0 5.0
  END
MODALFILTER vertical
  remove 1:500
  //
  cumulative_mef      x      y      z      Rx      Ry      Rz
  cumulative_mef      0.0    0.0    1.0    0.0    0.0    0.0 //VERTICAL
END
LOAD 1
  body
    gravity
      0.0      0.0    1.0
    scale -386.4
  END

```

3.7. DirectFRF Solution Case

Parameter	Type	Default	Description
interpolate points	<i>integer</i>	0	Number of additional points to interpolate for Padé expansion. If zero, no interpolation is performed.
interpolate order	<i>integer</i>	20	Order of the rational function for Padé expansion.
flush	<i>integer</i>	50	Defines how often results are written to the exodus results file. See Section 2.4.1

Table 3-45. – DirectFRF Solution Case Parameters.

Option **directfrf** is used to perform a direct frequency response analysis. In other words, we compute a solution to the Fourier transform of the equations of motion, i.e.,

$$\left(\underbrace{K + i\omega C - \omega^2 M}_{\equiv A(\omega)} \right) \bar{u} = \bar{f}(\omega) \quad (3.8)$$

where $\bar{u}$ is the Fourier transform of the response u , and $\bar{f}$ is the Fourier transform of the applied force. The matrix equation is then solved for each frequency. When a direct solver is used, this means that a complex factorization must be performed once per output. This is time-consuming, and the **ModalFrf** may be a better option for many situations (see Section 3.18).

The force function must be explicitly specified in the load section, and *must* have a **function** definition. Note that the force input provides the real part of the force at a given frequency, i.e., it is a function of frequency, not of time.

The frequency response function is evaluated at the frequencies specified as described in Section 7.4. In a **frequency** section set **freq_step**, **freq_min**, and **freq_max**. Also set an application region. Examples are presented in inputs 3.9, 3.10 and 3.16.

In addition to the output that is sent to the **frq** file, output may also be written to the **Exodus** file, provided that the keywords (such as acceleration) are specified in the **outputs** section. If nothing is specified in the **outputs** section, then nothing is written to the **Exodus** output files.

The expression “frf” is often interpreted as the *ratio* of output/input. There are reasons for using that ratio, including the confusion that can come from scaling the Fourier transform. The **Sierra/SD** code computes the output and does not compute a ratio. If the ratio is required, use a function with unit load as the input.

3.7.1. Padé Expansion

Computation of each frequency response is expensive because the system matrices must be computed, factored and solved once at each frequency. A cost-effective approach is to use a much coarser computational grid for full computation, and use a rational function (or Padé) expansion for intermediate points.⁴ The two additional parameters required for the expansion are listed in table 3-45. They are described below, and an example is shown in input 3.1. The theory is described in reference.⁶

```
solution
  case out
    directfrf
      Interpolate Points = 50
      Interpolate Order = 18
  end
```

⁴A rational function expansion is similar to a Taylor series expansion, but is capable of approximating resonant behavior.

Input 3.1. Padé Expansion Input Example. In this example, each exactly computed direct frequency response point will be separated by 50 interpolated values. These values will be determined using a Padé expansion of order 18.

3.8. *DirectFRF-Inverse Solution Case*

Parameter	Type	Default	Description

Table 3-46. – DirectFRF-Inverse Solution Case Parameters.

The `directfrf-inverse` solution method is used to solve an inverse problem (see Sec. 3.39) for a direct frequency response analysis. As in a forward solution, most of the parameters of an inverse frequency response method are found in other sections⁵. The user provides complex displacements and/or pressures at a set of nodes in the model, and the solution to the inverse problem is a set of loads, materials, etc. that best correspond with the user’s input.

The forward problem is defined in Eq. (3.8). The inverse equation is identical, but must be solved with optimization subject to regularization because measurements are available only at a subset of the analysis degrees of freedom.

The basic requirements for a `directfrf-inverse` simulation are as follows:

Optimization: Control over the optimization problem is specified in the `optimization` block. See Sec. 3.39.1 for further details.

Inverse-Problem: The `inverse-problem` block provides the connection to the measurement data. It is also where `design_variable` is specified (e.g., load, material, etc.). See Sec. 3.39 for details.

Truth Table: The truth table `data_truth_table` from the `inverse-problem` block is a list of the indices of the global node numbers (a.k.a. target nodes) where displacements or acoustic pressures are measured. See Sec. 3.39.3 for file format details.

Data File: Experimentally determined “target” displacements are read from `real_data_file` and `imaginary_data_file` specified in the `inverse-problem` block. See Sec. 3.39.3 for file format details.

Frequency: The frequencies at which the problem is solved are specified in the `frequency` block (Sec. 7.4).

⁵The forward solution supports a Padé expansion. This is not supported for inverse methods.

Sierra/SD uses the Rapid Optimization Library (ROL) for solving optimization problems. During the optimization solution ROL writes an output file, *ROL_Messages.txt* that contains convergence information. Section 3.39.8 contains a discussion of the output file that is written by ROL.

3.8.1. Load Identification

```
solution
  directfrf-inverse
end
inverse-problem
  design_variable = load
  data_truth_table = ttable.txt
  real_data_file = dataReal.txt
  imaginary_data_file = dataImag.txt
end
optimization
% optimization_package = ROL_lib
  ROLmethod = linesearch
  LSstep = Newton-Krylov
  LS_curvature_condition = null
  Max_iter_Krylov = 20
  Use_FD_hessvec = false
  Use_inexact_hessvec = false
  opt_tolerance = 1e-12
  opt_iterations = 2
end
loads
  sideset 301
    inverse_load_type = spatially_constant
    pressure=10
    function = 1
end
```

Input 3.2. Direct frequency response load identification example input

Specifying `design_variable = load` applies inverse methods to determine sideset loads which best correspond with the measured displacements and/or acoustic pressures provided by the user. The material and model parameters do not change during the solution. For structures, the loads are pressures or tractions⁶, and for acoustics, the loads are acoustic

⁶Moments and point forces are not currently supported.

accelerations. Note that for structures, inversion is based on the signed magnitudes of the tractions; the direction of each traction is fixed.

An example input deck is given in input 3.2. In addition to the input blocks discussed in the beginning of this section, there are several others specific to `design_variable = load`:

Loads: See Sec. 3.39.6 for a description of the inverse parameters in the **loads** block for load identification problems.

Frequency: For unknown loads, the frequencies at which the problem is solved are independent. That is, a separate load identification is performed at each frequency.

Pressure or traction load identification through *modal* frequency response is also supported for structures (see Sec. 3.19.1). Because modal solutions are significantly cheaper than direct solutions, one approach might be to begin the inverse optimization with `modalfrf-inverse`, then use the output as the initial guess for a follow-up `directfrf-inverse` case. Section 3.39.7 contains a discussion of the current limitations with inverse load methods.

3.8.2. Material Identification

Specifying `design_variable = material` with the `directfrf-inverse` method applies inverse methods to determine material parameters when provided with both loads and structural displacements and/or acoustic pressures in a given finite element model⁷. The load parameters do not change during the solution. As in the previous section, the forward problem is defined in Eq. (3.8), and the inverse equation is identical but must be solved with optimization subject to regularization because measurements are available only at a subset of the analysis degrees of freedom. The solution provides the material parameters for elements in the model that are specified to have unknown materials.

In addition to the input blocks discussed at the beginning of this section, there are several others specific to `design_variable = material`:

Block: See Sec. 3.39.4 for a description of the **block** specifications for material inverse problems.

Material: See Sec. 3.39.5 for a description of the provides **material** specifications for material inverse problems.

Frequency: For unknown materials, the same set of material properties apply for every frequency in the simulation, except in the case of frequency-dependent material properties.

⁷As the system matrices (and consequently the modes) change at every inverse iteration, `design_variable` cannot be set to `material` for `modalfrf-inverse` problems.

Viscoelastic material identification is also supported using measured homogenized complex bulk and shear moduli. This capability is limited to structural-only problems where all material blocks are isotropic viscoelastic.

3.8.3. Multi-Experiment Material Identification

In the same manner as `design_variable = material` described in the previous section, `design_variable = multi_material` may be used to apply inverse methods in the frequency domain to determine material parameters. Here, multiple inverse problems are combined. For instance, if two different load and displacement conditions result in two separate responses for the same set of material properties, this method will use both responses to determine a single set of material properties.

Most parameters for a multi-experiment inverse frequency method are similar to those for a single-experiment inverse frequency method. The differences occur in the following sections:

Loads: The **loads** block must be empty for this solution case; anything the user specifies here will be overwritten by what they specify in the **load** block.

Load: Each **load** block provides a separate set of loads for each experiment individually.

Inverse-Problem: In addition to the parameters discussed earlier, the **inverse-problem** section must include values for **nresponses** for the number of experiments and **loadID** to specify a list of loads, one for each experiment.

3.9. Dump Solution Case

Parameter	Type	Default	Description

Table 3-47. – Dump Solution Case Parameters.

The keyword **dump** will cause **Sierra/SD** to form matrices. No solve will be done. The main reason to use this solution case is to output the matrices to mfiles [2.3](#) for debugging or custom post processing.

3.10. Eigen Solution Case

Parameter	Type	Default	Description
nmodes	<i>Integer</i>	10	Number of modes to extract. See Section 3.10.1
shift	<i>Real</i>	-1.0e6	Shift to apply to matrix system to allow solving singular systems. See Section 3.10.2
untilfreq	<i>Real</i>	Inf	Target frequency to reach. See Section 3.10.2.1
ModalFilter	<i>string</i>	none	Modal filter to define modes to retain. See Section 3.10.2.2
fluidloading	<i>yes/no</i>	no	Turns on added mass for approximate Wet Modes calculation.

Table 3-48. – Eigen Solution Case Parameters.

The **eigen** solution case computes eigenvalues and mode shapes of a system representing the natural vibration modes. The parameters **NegEigen** , **eig_tol** , and **eigen_norm** are described in sections [2.3](#) and [2.3](#)

Eigenvectors, Φ , are stored in **Exodus** files as time dependent displacements. The time of an eigenvector is the frequency. Seacas library tools expect times to be unique, occasionally causing benign warning messages when there are multiple modes at the same frequency. The mode shapes multiplied by the consistent mass matrix [7.1.10](#) are also available.

3.10.1. Option nmodes

Nmodes is the number of modes to compute. The eigenvalues are computed beginning with the lowest frequency mode and working up. The calculation continues until **nmodes** have converged. Modal analysis uses iterative Lanczos procedures. These methods build a Krylov subspace from which the solution is determined. The Krylov subspace method is designed to find the lowest frequencies, which is typically what is needed for structural analysis. More linear solves and longer run times are required to compute more modes and reach higher frequencies.

If **nmodes=all**, then all the finite modes of the structure are computed using dense linear algebra. No shift is required. The mass matrix may be either full rank or singular. Constraints may be present. This option is only available in serial. The number of dofs must be < 1000 .

3.10.2. Solving Singular Systems with Shifts

There is a trade off between reducing the number of linear solves required to solve an eigenvalue problem or reducing the cost per linear solve. This trade off is controlled using the shift. The number of solves per eigenvalue decreases as the shift decreases in magnitude, and at the same time the cost per linear solve increases. Direct linear solvers require a shift for singular (floating) systems. Iterative methods are more reliable when applied to positive definite linear systems than they are when applied to positive semi-definite linear systems. The eigenvalue problem is defined as,

$$(K - \omega^2 M)\phi = 0. \quad (3.9)$$

K and M are positive semi-definite matrices. If the shift $\sigma < 0$, then $K - \sigma M$ is positive definite. Solution cases generally involve positive definite matrices except direct frequency response.

Here K and M are the stiffness and mass matrices respectively, and ω and ϕ are the eigenvalues and vectors to be determined. The problem may be solved using a variety of methods — the Lanczos algorithm is used in **Sierra/SD**. In this method, a Krylov subspace is built by repeated solving equations of the form $Ku = b$. For floating structures, or structures with zero-energy mechanisms, K is singular and special approaches are required to solve the system. The two approaches used in **Sierra/SD** are described below.

Deflation. If it is possible to identify the singularity in K , then the null vectors of K are eigenvectors (with $\omega = 0$), and the system can be solved by ensuring that no component of the null vectors ever occurs in b . This approach is equivalent to computing the pseudo inverse of K .

The strength of deflation is that if the eigenvectors can be determined exactly, the Lanczos algorithm is unaltered and the remaining vectors can be determined somewhat optimally. The difficulty is ensuring that we have correctly determined the eigenvectors, especially when mechanisms or multipoint constraints exist in the model.

As an example, the geometric rigid body modes capability 3.36 can be used to analytically compute the standard rigid body modes and deflate them out of a singular system.

Shifting. The second method involves solution of a modified (or shifted) eigenvalue problem.

$$((K - \sigma M) - \mu M)\phi = 0. \quad (3.10)$$

This system has the properties that the eigenvectors, ϕ , are unchanged from the original equation, and the eigenvalues, μ , are related to the original values. Namely, $\mu = \omega^2 - \sigma$.

On serial platforms, a small negative shift is normally sufficient to solve the problem due to the high accuracy of serial direct solvers. For parallel solution, a reasonable shift value is usually given by $\sigma = -\omega_{elas}^2$, where ω_{elas} is the expected first nonzero (or elastic) eigenvalue. Because **Sierra/SD** cannot compute an optimal shift value a priori, a default of -1 is used and a warning is written if the shift used is well outside of the expected range. In practice, a shift of $-1e6$ is often used in **Sierra/SD** input, as this gives reasonable results for extracting frequencies anywhere near 150Hz, assuming the time unit is seconds. This is often the general frequency range of engineering interest. However, an adequate shift may vary some depending on the units and the properties of an analysis.

The shifted problem benefits from the fact that $K - \sigma M$ can be made non-singular (except in rare situations). This is done by choosing σ to be a large negative value. Unfortunately, the Lanczos routine convergence is affected if σ is chosen to be too far from the recommended value of $\sigma = -\omega_{elas}^2$.

If σ is too large, many solves will be required to determine the eigenvalues, which consequently slows convergence. If σ is too small each linear solve may be near singular, requiring many solver iterations or being unable to reach the target residual at all resulting in a 'SOLVER OUT OF BOUNDS' error.

Another consequence of an excessively large or small negative shift is that potentially not all redundant zero eigenvalues may be found. These modes may be found by correcting the shift, tightening tolerances, or by restarting.

The shifted eigenvalue problem is more reliable. Set the `grbm_tol` to a small value (e.g. $1e-20$)(or use the default), and manually enter a negative shift. The output should still be examined to ensure that no global rigid body modes are detected.

If the model is not floating and has no mechanisms and no zero energy modes, the system is not singular and a shift is unnecessary.

Example

A representative **Solution** section for an eigen analysis with a **shift** of -10^6 followings. A dozen modes are requested. This shift would be appropriate for a system where the first elastic mode is approximately 150Hz.

```
Solution
  eigen
  nmodes 12
  shift -1.0e6
end
```

3.10.2.1. UntilFreq Option and Eigen Restart

The **untilfreq** keyword provides an additional method of controlling the eigen spectrum to be computed. If this value is provided, then the analysis will be automatically (and internally) restarted until the frequency of the highest mode is at least the value of the **untilfreq**. This restart capability is somewhat crude. There are always **nmodes** new modes computed on each calculation. Also, because there can be inaccuracies associated with restarting the eigen solver, **Sierra/SD** restarts a maximum of **5** times.

Sierra/SD uses the **ARPACK** Lanczos solver for the eigen problem. This solver maintains the orthogonality of the eigenvectors for a single batch of modes. However, when restarted, the solver must deflate out the previously computed modes. There can thus be a slight loss of orthogonality due to round off. With repeated restarts, the effect can significantly reduce accuracy.

Additionally, eigen restart can be manually controlled via the multicase solution. In the example below, the first 1000 modes are computed in 'eig1'. Next, **eig2** restarts and reads and deflates those 1000 modes from the system. In this solution case, an additional 500 modes are computed, for a total of 1500 modes. Restarting eigen analysis in this way can compute modes with less memory than computing all modes at once. An excessive number restarts of this nature (more than about five) will lead to accuracy loss from round off errors during deflation.

```
Solution
  case 'eig1'
    eigen
    nmodes 1000
  case 'eig2'
    eigen
    restart = read
    nmodes 1500
end
```

3.10.2.2. ModalFilter Option

The optional **modalfilter** keyword provides a means of reducing the modes retained for output and for subsequent analysis. For more details, see Section 3.16.

Use of this parameter within the **eigen** solution case is *deprecated* in favor of the new **modalfiltercase** solution case (Section 3.16).

You can also put the modal filter into a separate case, called **preddam** , 3.6.

3.10.2.3. Fluid loading Option

Although a coupled structural/acoustic eigenvalue problem is quadratic, under certain approximations, the fluid may be represented as an added mass on the structure, and a real eigen problem results as described in subsection Wet Modes or Added Mass section

Solution Procedures of the Theory Manual. The `fluidloading` enables that added mass calculation. For wet modes all material blocks that use an acoustic material are treated alike as a fluid domain for mass loading.

In an early implementation the lowest modes of the acoustic stiffness matrix were used to approximate the added mass. As a sanity check, one mode of the acoustic stiffness matrix is approximated.

Fluid loading is not applicable to models that use infinite elements. There are two reasons for this. First, infinite elements lead to a nonsymmetric eigenvalue problem, which is not supported. Second, with infinite elements, the fluid region stiffness matrix is singular, implying an infinite added mass. A fatal error is returned if infinite elements are detected.

3.10.2.4. Rigid Body or Zero Energy Modes

Rigid body modes represent rigid body translation or rotation of a structure. A normal free-free structure would be expected to have 6 rigid body modes. A structure with constraints may have fewer than 6 rigid body modes. On the other hand a structure with multiple disconnected pieces may have more than 6 rigid body modes. To compute *any* rigid body modes, one must request all of them. For structures with 6 rigid body modes, `nmodes` $\geq$ 6. Otherwise, the solution case is likely to fail. Prior to computing the wet modes, rigid body modes are generated by adding a Geometric Rigid Body Modes solution case as described in Section 3.36. As described there the parameter `num_rigid_modes` must be set to 6.

3.11. Eigen-Inverse Solution Case

Parameter	Type	Default	Description
nmodes	<i>Integer</i>	10	Number of modes to extract. See Section 3.10.1
shift	<i>Real</i>	-1.0e6	Shift to apply to matrix system to allow solving singular systems. See Section 3.10.2
untilfreq	<i>Real</i>	Inf	Target frequency to reach. See Section 3.10.2.1
ModalFilter	<i>string</i>	none	Modal filter to define modes to retain. See Section 3.10.2.2
modalAdjoint-Solver	<i>gdsw/ camp/ both</i>	camp	Select solver for the inverse problem (eigenvector material identification only)

Table 3-49. – Eigen-Inverse Solution Case Parameters.

The **eigen-inverse** solution method is used to solve an inverse problem (see Sec. 3.39) for an eigen analysis. In this solution method, only material identification is currently supported. Specifying **design_variable = material** applies inverse methods in the modal domain to determine material properties on a block or element when provided with modal frequencies and mode shapes. The user specifies some of the lowest modes of the structure, and optionally the mode shapes of the structure at locations in the model.

The standard parameters for modal analysis also apply here (see Sec. 3.10. Modal filters (3.16) may not be applied in **eigen-inverse** analysis. The analysis requires input both for measurement data and for control of various optimization parameters. See the following sections for details:

Optimization: Control over the optimization problem is specified in the **optimization** block. See Sec. 3.39.1 for further details.

Inverse-Problem: The **inverse-problem** block provides the connection to the measurement data. Of particular importance are the parameters **modal_data_file** and **modal_weight_table**, which are described further in Sec. 3.39.3. Also necessary is the parameter **design_variable = material**.

Sierra/SD uses the Rapid Optimization Library (ROL) for solving optimization problems. During the optimization solution ROL writes an output file, *ROL_Messages.txt* that contains convergence information. Section 3.39.8 contains a discussion of the output file that is written by ROL.

3.11.1. Eigenvalue Material Identification

```
solution
  eigen-materialid
    nmodes=12
    shift=-1e5
end
inverse-problem
  design_variable = material
  modal_data_file = modal_data.txt
  modal_truth_table = modal_truth.txt
end
optimization
  optimization_package = ROL_lib
  ROLmethod = linesearch
  LSstep = secant
  scaleDesignVars = yes
  LS_curvature_condition = null
  Max_iter_Krylov = 50
  Use_FD_hessvec = false
```

```

Use_inexact_hessvec = true
opt_tolerance = 1e-10
opt_iterations = 30
end

```

Input 3.3. Eigen material identification example input

In the case of eigenvalue optimization, only the modal frequencies are included in the objective function. An example input is shown in input 3.3. The theory for this problem is available in [16]. The objective function for the eigen value problem is given as:

$$J(\boldsymbol{\lambda}, \mathbf{u}, \mathbf{p}) = \sum_{i=1}^{i=m} \left[\frac{\beta_i}{2} \left(\frac{\lambda_i - \lambda_{mi}}{\lambda_{mi}} \right)^2 \right] + \mathcal{R}(\mathbf{p}). \quad (3.11)$$

Where m is the number of modes, β_i is zero or one, λ_i is the computed eigenvalue, λ_{mi} is the measured eigenvalue, and $\mathcal{R}(\mathbf{p})$ is the regularization term.

3.11.2. Eigenvector Material Identification

In this case, both the eigenvalues (modal frequencies) and eigenvectors (mode shapes) are matched in the inverse solution. A detailed description of the theory and implementation details of this solution case is given in [16]. To use this capability, it is necessary to specify the keywords `data_file` and `data_truth_table` for the eigenvector data and eigenvector truth table, respectively, which specify the modal shape amplitude data and truth table information. The latter allows one to differentiate between a tri-axial and uni-axial accelerometer.

The objective function extends Eq. (3.11) by adding the eigenvector term.

$$J(\boldsymbol{\lambda}, \mathbf{u}, \mathbf{p}) = \sum_{i=1}^N \left[\frac{\beta_i}{2} \left(\frac{\lambda_i - \lambda_{mi}}{\lambda_{mi}} \right)^2 + \frac{\kappa_i}{2} \frac{\|\mathbf{u}_i - \mathbf{u}_{mi}\|_{\mathbf{Q}}^2}{\|\mathbf{u}_{mi}\|_{\mathbf{Q}}^2} \right] + \mathcal{R}(\mathbf{p}), \quad (3.12)$$

where $\mathbf{u}_i$ is the computed eigenvector, $\mathbf{u}_{mi}$ is the measured eigenvector, and $\mathbf{Q}$ is the observation matrix obtained from the truth table.

The issues an implementation must handle include repeated modes, crossing modes, the singular adjoint linear system, and eigenvector scaling.

3.11.2.1. Repeated Modes When the computed or measured data contain two repeated eigenvalues, the associated eigenvectors are not deterministic and steps are taken to orthogonalize these modes with respect to the measured data. The measured data must be sufficient enough for this orthogonalization for the inverse problem to converge. Note that repeated modes can occur at iterations in the inverse process even when no repeated modes are present in the measured data.

3.11.2.2. Mode Swapping/Crossing If the computed eigenvalues of a structure are ordered from smallest to largest, the ordering of mode shapes will typically change as the material parameters are varied. This also causes non-differentiability in the objective function, which causes difficulties in gradient-based optimization. Mode tracking refers to maintaining a correspondence of eigenpairs (eigenvalue and eigenvectors) between an original and an updated system throughout changes in the eigenproblem. Measured data is often incomplete, having only a few measured data points (physical accelerometer locations) on a model with millions of degrees of freedom. An incorrect mode swap results in a discontinuity in the slope of the objective function.

A mode tracking algorithm is used to minimize eigenvector misfit at each optimization step.

3.11.2.3. Singular Solve The Adjoint Solution is singular due to the fact that the eigenvector u_i is in the kernel of the coefficient matrix. In order for a solution to exist, the right hand side must be orthogonal to u_i . Additionally, if rigid body modes ($\lambda = 0$) or repeated mode are present, components of the corresponding eigenvectors must also be removed from the right hand side before the solve. Even when this is done, however, the resulting system of equations is singular and a Helmholtz (indefinite) problem, which presents significant computational cost and robustness challenges for iterative linear solvers.

The `modalAdjointSolver = camp` (default for eigenvector inversion) option enables a new solver that uses a modal superposition of the previously computed eigenvectors to solve this system of equations. When using the `camp` solver, it is recommended to request more modes than contained in the measured data, and use the truth table file to remove these modes from the optimization part of the solution.

The *CAMP* block is used to define the parameters for the `camp` solver.

```
CAMP
    SC_OPTION 0
END
```

Input 3.4. CAMP solver block

Note that it is recommended to set the `sc_option` parameter to 0 for the `camp` solver.

3.11.2.4. Computed Eigenvector Scaling An important consideration in eigenvector optimization is that the mode shapes computed in Sierra/SD are by default *mass normalized*. Measured modal shape amplitudes, on the other hand, could present with very different scalings, since any eigenvector can be scaled by an arbitrary scale factor and will still be a valid eigenvector. Thus, the eigen-inverse solution method includes an automatic re-scaling of the computed mode shapes in the optimization so that they have the same

norm as the measured mode shapes. This re-normalization allows them to be properly differenced in the objective function. We note that this internal re-scaling requires no user intervention.

If the norms of the measured eigenvectors differ substantially from the norms of the eigenvectors computed in Sierra/SD, then the re-scaling described in the previous paragraph is necessary to correctly determine the next iteration of the design variables. The scaled computed eigenvector $\tilde{\mathbf{u}}_i$ can be written as

$$\tilde{\mathbf{u}}_i = \alpha_i \mathbf{u}_i, \quad (3.13)$$

where $\alpha_i = \|\mathbf{u}_{mi}\|/\|\mathbf{u}_i\|$ such that the norm of $\tilde{\mathbf{u}}_i$ is identical to the norm of eigenvector $\mathbf{u}_{mi}$. With this change, the eigenvector term in Eq. (3.12) becomes

$$J(\mathbf{u}) = \sum_{i=1}^N \frac{\kappa_i}{2} \frac{\|\alpha \mathbf{u}_i - \mathbf{u}_{mi}\|_{\mathbf{Q}}^2}{\|\mathbf{u}_{mi}\|_{\mathbf{Q}}^2}, \quad (3.14)$$

A corresponding change to the gradient $J_{\mathbf{u}}$ is also required, but this change is not discussed here.

3.12. *AEigen Solution Case*

Parameter	Type	Default	Description
nmodes	<i>Integer</i>	10	Number of modes to extract. See Section 3.10.1
shift	<i>Real</i>	-1.0e6	Shift to apply to matrix system to allow solving singular systems. See Section 3.10.2
untilfreq	<i>Real</i>	Inf	Target frequency to reach. See Section 3.10.2.1
ModalFilter	<i>string</i>	none	Modal filter to define modes to retain. See Section 3.10.2.2
implicit_restarts	<i>integer</i>	5	Number of allowed restart steps, default same as eigen
anblocksize	<i>integer</i>	1	Number of subspace vectors added each step, default same as eigen
Convergence tolerance for solver iterations			
anrho_shift	<i>Real</i>	0.0	Shift of infinite eigenvalues of constrained problem
anverbosity	<i>integer</i>	17	Level of verbosity of log files, see below for details
subspace_size	<i>integer</i>		Maximum number of subspace BKS basis vectors. By default, picks based on number of requested eigen modes.

Table 3-50. – AEigen Solution Case Parameters.

The symmetric generalized eigenvalue problem of equation 3.9 is solved for the lowest frequency modes. Eigenvalue problems are solved using a linear solver. The iterative solution of the eigenvalue problem is a nonlinear iteration that requires multiple linear solves.

The standard eigenvalue methods of Section 3.10 use the ARPACK⁴¹ package. It is a powerful, public domain solver and has been successful. Other approaches are being developed. The **AEigen** keyword selects the Krylov-Schur method in the Anasazi⁷ package.

Modal solution methods all control accuracy using the parameter `eig_tol`.

A Block Krylov-Schur (BKS) method is used that is similar to the ARPACK solver, discussed in Section 3.10. A subspace is built by repeatedly solving systems of the form $(K - \sigma M)$, where σ is the shift specified by keyword `shift`; see 3.10.2 for more discussion on shifts.

Unlike ARPACK, Anasazi optionally employs block solvers. That is, the subspace generated by the BKS iteration adds multiple vectors at each step, where ARPACK adds a single vector. The number of vectors is specified by the keyword `anblocksize`.

The amount of diagnostic information or level of verbosity is changed used `anverbosity`. The default value is 17, implying that the Anasazi solvers will output errors, warnings, and timing details. Each verbosity type is controlled by a single bit in the integer `anverbosity`. These types are listed in Table 3-51. Each combination is valid, the combinations being formed by adding different verbosity values. For example, setting `anverbosity` to $25 = 0 + 1 + 8 + 16$ requests output for Errors, Warnings, Final Summary and Timing Details.

Verbosity type	Value
Errors	0
Warnings	1
Iteration Details	2
Orthogonalization Details	4
Final Summary	8
Timing Details	16
Status Test Details	32
Debug	64

Table 3-51. – AEigen Verbosity Table.

Input 3.5 illustrates the “Solution” section of an input file using the Krylov-Schur method, “`ansolver BKS`”. A dozen of the lowest frequency modes are requested. The `shift` is -10^6 . The shift is appropriate for a floating system where the first elastic mode is approximately 150 Hz. This produces an eigen analysis equivalent to the example given in 3.10 for `eigen`. The verbosity level specifies that after finishing the eigen analysis, the solver will print status information (number of iterations, current eigenvalues) and timing statistics.

```
Solution
  aeigen
    nmodes 12
    shift -1.0e6
    ansolver BKS
    anblocksize 1
    anverbosity 25
end
```

Input 3.5. Aeigen example

3.13. *LargestEv Solution Case*

Parameter	Type	Default	Description

Table 3-52. – LargestEv Solution Case Parameters.

The `Largest_Ev` analysis determines the largest eigenvalue of the system, i.e.

$$(K - \lambda M)\phi = 0$$

There are no arguments to the solution method. The eigenvalue not physically meaningful. This solution is used primarily for debugging.

3.14. *Fatigue Solution Case*

Parameter	Type	Default	Description
method	<i>narrow-band/wirsching</i>	narrow-band	Fatigue calculation algorithm to use.
duration	<i><real></i>	1	Time duration, τ over which to integrate fatigue damage.

Table 3-53. – Fatigue Solution Case Parameters.

High Cycle Fatigue occurs after long periods of alternating stresses in the elastic range. The input parameters for fatigue-failure are shown in Tables 3-53 and 4-105. An example is shown in input 3.6. Fatigue analysis requires inputs in several sections.

1. The fatigue section in the **Solution** block must be immediately preceded by a `modalranvib` solution, with `noSVD` option (3.20). This defines the random stress moments needed for computation of the stress crossing rates.
2. The **Fatigue** section in the Solution block defines general fatigue parameters. Relevant parameters are outlined in Table 3-53.
3. The **material** section must include parameters necessary for computation of fatigue and damage. Sections without this input will not have a fatigue parameter output. Typical materials parameters required for fatigue analysis are found in Table 4-105, found in the materials section, 4.4.2.

```

Solution
  title 'High cycle fatigue example'
  case modes
    eigen
    nmodes 10
  case vrms
    modalranvib
    noSVD
    lfcutoff 10
  case failure
    fatigue
    duration = 1.0
    method = Wirsching
end
Block 100
  material aisi4140
end
Block 200
  material steel
end
Material aisi4140
  E = 10e6
  density = 0.00075
  Fatigue_A1 = 31.5805
  Fatigue_A2 = -14.0845
  Fatigue_Stress_Scale = 0.001 // Psi to Ksi
  MaterialType = PeakStress // or StressRange
end
Material steel
  E = 30E6
  nu = 0.28
  density = 0.007
end

```

Input 3.6. High Cycle Fatigue Input Example. Fatigue is specified by block. In this example, there are two blocks. The material properties for block `aisi4140` include fatigue parameters, but the material properties for block `steel` do not include fatigue parameters. The fatigue calculation is performed for block 100.

3.14.0.1. User Output If the `fatigue` method is specified in the Solution section then output for fatigue will be provided. There is no need for a specification in the **outputs** section. The following fields will be output.

NarrowBandDamageRate: There are two possible equations for D_{NB} depending on if the material parameters specified for the S-N curve are based on a stress range or peak stress. The Narrow Band damage per unit time, defined as:

$$\dot{D}_{NB} = \frac{\nu_o^+}{A} (\sqrt{2}\sigma_s F_{SS})^m \Gamma\left(\frac{m}{2} + 1\right) \quad (3.15)$$

for materials defined based on Peak Stress, and

$$\dot{D}_{NB} = \frac{\nu_o^+}{A} (2\sqrt{2}\sigma_s F_{SS})^m \Gamma\left(\frac{m}{2} + 1\right) \quad (3.16)$$

for materials defined base on the Stress Range. Equation (3.15) is used by default. A stress range type material can be specified in the input (MaterialType = StressRange). See the theory manual for details of the parameters.

WirschingDamageRate: The Wirsching damage per unit time. This may be thought of as a scaled value of the NB damage intensity.

$$\dot{D}_W = \lambda \dot{D}_{NB}.$$

ZeroCrossingIntensity: the expected positive zero-crossings intensity.

$$\nu_o^+ = \sqrt{\frac{M_2}{M_0}}$$

Damage: The selected damage rate multiplied by time.

$$D = \tau \dot{D} \quad (3.17)$$

Here $\dot{D}$ can be either the Narrow Band or Wirsching Damage Rate.

PeakFrequency: the expected peak occurrence frequency.

$$\nu_p = \sqrt{\frac{M_4}{M_2}}$$

The even moments, M_x (with $x=0,2,4$), are output as part of the random vibration computation, see (3.20). For example, $M_2 = (\text{VRMS2}/2\pi)^2$.

3.15. Buckling Solution Case

Parameter	Type	Default	Description
nmodes	<i>Integer</i>	10	Number of modes to extract. See Section 3.10.1
shift	<i>Real</i>	-1.0e6	Shift to apply to matrix system to allow solving singular systems. See Section 3.10.2
untilfreq	<i>Real</i>	Inf	Target frequency to reach. See Section 3.10.2.1
ModalFilter	<i>string</i>	none	Modal filter to define modes to retain. See Section 3.10.2.2
bucklingSolver	<i>ARPACK/Anasazi</i>	ARPACK	Eigen solver method to use.

Table 3-54. – Buckling Solution Case Parameters.

The **buckling** keyword is used to obtain the buckling modes and eigenvalues of a system. The parameters which can be specified for a buckling solution are tabulated below. Users are encouraged to review the discussion of buckling analyses in.⁵³

The **shift** parameter indicates the shift desired in a buckling analysis. The shift value represents a shift in the eigenvalue space (i.e. ω^2 space). It is set using the **shift** parameter. Determining an effective shift is problem dependent.

The **nmodes** parameter specifies the number of requested buckling modes. Its default value is 10.

Unlike eigen analysis, buckling solution cases require a **loads** block. This is because buckling is always specified with respect to a particular loading configuration. For example, for a pressure load applied on a sideset, the buckling analysis would indicate the critical amplitude of the applied pressure needed to cause buckling. The critical buckling load is computed as the product of the first (lowest) eigenvalue times the amplitude of the applied load. Thus, for the case

```
Loads
  sideset 1
  pressure = 10.0
end
```

The lowest obtained eigenvalue is 100.0. The critical buckling pressure would be $P_{cr} = 100.0 \times 10.0 = 1000.0$. This would indicate that buckling would occur if the loading were applied as,


```

Loads
  sideset 1
  pressure = 1000.0
end

```

Similar conclusions can be drawn about force loads on nodesets.

Buckling of floating structures is not supported at this time. If global rigid body modes are present, the solution may not be correct. The use of beams and shells in the buckling solution case is not supported. For models with multipoint constraints, the default serial linear solver is not available.

Example A Solution section for buckling analysis with a **shift** of -10^6 looks like the following, if 1 mode is needed (i.e. if the user is confident that the modes are well separated).

```

Solution
  buckling
  nmodes 1
  shift -1.0e6
end

```

3.16. ModalFilter Solution Case

Parameter	Type	Default	Description
ModalFilter	<i>string</i>	none	Modal filter to define modes to retain. See Section 3.10.2.2
write_files	<i>all/ none/ output/ history</i>	all	Controls which result files are written during this solution.

Table 3-55. – ModalFilter Solution Case Parameters.

A ModalFilter filters modes computed by an **eigen** solution case (Section 3.10). The filtered modes are then used by subsequent solution cases.

Since the **eigen** solution case already outputs the full set of modes to disk, the filtered modes are not written to disk by default. The **write_files** parameter selects output of the filtered modes.

Controlling the modes retained for subsequent analyses can significantly reduce run time with little effect on the desired response. For example, a shell structure may have many

hundreds of modes contributing to the normal mode response, and only a few that interact with the loads. ¹

During eigen analysis or CBR reduction, the usual number of modes (**nmodes**) are computed. These modes are filtered, and only a subset are written to the **Exodus** file or used in subsequent analysis. An example input is shown in input 3.7.

```
SOLUTION
  case eig
    eigen
    nmodes=500
  case filt
    modalfiltercase
    ModalFilter=MPF1
END

ModalFilter MPF1
  remove 1:999
  cumulative mef 0.8 0.8 0.8 0.2 0.2 0
  add 99:101,103
END
```

Input 3.7. Example ModalFilter Input

For this example, the following actions are performed in the filter.

1. The first 999 modes are removed. In this case 500 modes are computed, and all modes are removed.
2. The modes contributing the most to a cumulative modal effective mass are added. Modes sufficient for 80% contributions to the x , y , and z directions are added. Modes needed to achieve 20% of the rotational terms for x and y are added. Since the contribution for rotation about z is zero, no modes are added there.
3. Modes 99, 100, 101 and 103 are added if they are not already included.

Each entry in the modal filter section consists of two parts: an action (like *remove* or *add*) and an application space. The *application space* for the “add” and “remove” space is an integer list with a format much like MATLAB. See section 2.1 for more details. Valid action keywords are listed in Table 3-56.

remove Removes modes in the application space from output.

add Adds modes in the action space.

¹The modes of large ship are an example. Only a few of the modes contribute to global bending or torsional modes. The remaining modes are local, and may not be of interest to the analysis.

Keyword	Application Space
remove	integer list or “all”
add	integer list or “all”
cumulative mef	6 fractions
cumulative nmef	6 fractions

Table 3-56. – Modal Filter Keywords.

cumulative mef Adds modes which contribute to the modal effective mass. Following this keyword sequence, 6 fractions are entered, one for each of the 6 rigid body mode contributions. The modes are sorted and modes are kept that contribute most to the modal effective mass. When the fractional contribution exceeds the threshold, no more modes are added for that direction. Contributions from each direction are combined (unioned) and added to the list of modes kept.

The 6 fractions following the keyword indicate the threshold for each coordinate direction. Each fraction must be between 0 and 1, inclusive. A value of zero means no modes are retained. A value of unity retains all modes.

cumulative nmef Adds modes which contribute to the *normalized* modal effective mass. This option is identical to the “cumulative mef” option except that the terms are normalized such that the total contribution from all computed modes sums to one.

The Modal Effective Mass equals the Modal Participation Factor. The Modal Participation Factor is defined in the next section.

The Normalized Modal Effective Mass is not defined.

3.17. *Modal Participation Factor Solution Case*

Parameter	Type	Default	Description
write_table	<i>Yes/No</i>	Yes	If yes write Γ table, if no write summary table
blockwise	<i>Yes/No</i>	Yes	If yes write blockwise summary to result table
RCID	<i>string</i>	123	A string representing rigid body modes to include in the calculation. 123 resents the translational degrees of freedom. 123456 includes all six degrees of freedom.

Table 3-57. – MPF Solution Case Parameters.

A *Modal Participation Factor* (MPF) is a quality of a mode shape. Another name for the *Modal Participation Factor* is the *Modal Effective Mass*. A MPF is a direction cosine of an eigenvector along one of the 6 rigid body modes. It measures the interaction of the modes with a gravity load or a base excitation.

The rigid body modes $\{R_i\}_{i=1}^6$ ignore any boundary conditions.

The modal participation factor of an eigenvector v of the constrained system is determined from the representation of v in the unconstrained space, using a lumped mass matrix 2.4.4,

$$\Gamma_{ij} = \frac{R_i^T M v_j}{\sqrt{(R_i^T M R_i)(v_j^T M v_j)}}. \quad (3.18)$$

Γ_{ij} is a mass normalized measure of the contribution of a given rigid body term, γ_i , to the vector, v_j . A summary term which represents the total fraction of a vector that is spanned by all rigid body modes is also useful.

$$\text{MPF}_j = \sum_i^6 \Gamma_{ij}^2 \quad (3.19)$$

The MPF method computes these participation factors for the eigenvectors of a system. This method *must* be used as part of a multcase solution, and the previous case *must* be an eigenvalue problem (see Section 3.10). Further, this method (by default) computes the modal participation factor on a block by block basis. Thus, those portions of the model that most contribute to the rigid body motion may be determined. ² Then,

$$\Gamma_{ij}^k = \frac{R_i^T M^k v_j}{\sqrt{(R_i^T M^k R_i)(v_j^T M^k v_j)}} \quad (3.20)$$

Options for the MPF method are listed in Table 3-57.

Summary data from the calculation is written to the results file as described in Table 3-58. In addition, unless *write_table=no*, data will be written to an external text file. The format for the file is specified in the results file. It contains the block wise modal participation factors, Γ_{ij}^k of equation 3.20. An example is provided in input 3.8.

The external text file is intended to be easily read by external programs such as the MATLAB “load” command. It therefore has no header information. The data ordering is exactly the same as the table written to the echo file (which contains that header information). Each column is grouped first by block (in the order of the blocks in the genesis file), and then by degree of freedom. Usually there are either 3 or 6 dofs per block entry. Each row corresponds to a single mode.

The optional external text file with filename extension mpf contains block-wise modal participation factors. The data is presented in tabular format, separated by white space.

²The overall modal contribution is not the sum of the block wise contributions, and contributions from individual blocks may cancel other blocks. See Table 3-58.

Data	Value	Description
MPF	$\sum_i (\Gamma_{ij})^2$	Overall mode _{<i>j</i>} MPF
MPF-B _{<i>k</i>}	$\sqrt{\sum_i (\Gamma_{ij}^k)^2}$	MPF for block <i>k</i> , mode <i>j</i>
MPF by RBM _{<i>i</i>}	$\sum_j (\Gamma_{ij})^2$	MPF for direction <i>i</i>

Table 3-58. – MPF Summary data. Each mode, v_j , has contributions from each of these summary values.

Most data analysis software tools can easily import this type of data for analysis and plotting (e.g., Microsoft Excel, OpenOffice Calc, Python, MATLAB, Octave, etc.). The MPF file contains no header information, so it is important to understand what each column and row represents. Each row of data corresponds to a mode. Columns represent modal participation factors calculated for each block and requested coordinate (controlled with “rcid”). The columns are grouped first by block, and then by degree of freedom. Blocks are written out in the order they are found in the genesis file (note: they are not sorted by Block ID or by the order they appear in the input file). Hence, if the exodus file contains two blocks and rcid=123 (default), the *.mpf file will contain six columns in the following order: Block1_x, Block1_y, Block1_z, Block2_x, Block2_y, Block2_z. If rcid=123456, then six columns per block (x, y, z, Rx, Ry, Rz) will be written out and there will be 12 columns.

```
Solution
case eig
  eigen nmodes=10
  shift=-1e5
case out
  mpf
  blockwise=yes
  RCID=123
  write_table=yes
end
```

Input 3.8. Modal Participation Factor (MPF) Example

Modal Effective Mass There are several definitions of the modal participation factor. The value from equation 3.18 is unit normalized such that the sum of the squares of Γ_{ij} over all modes equals unity. A related term is the modal effective mass.

$$MEFF_{ij} = \sqrt{M_{oi}} \cdot \Gamma_{ij} \quad (3.21)$$

Here M_{oi} is total mass associated with each rigid body mode. The sum of the squares of $MEFF_{ij}$ over all modes, j , for a given rigid body mode, i is the total mass associated with

that RBM. Modal truncation decreases the sum. The `modal effective mass` table is output to the result file immediately after the modal participation factors. At the bottom of the table, the sum of the squares of the modal effective mass is given for each rigid body mode. The total mass and moments of inertia of the system are also given.

Lumped or Consistent Mass We always use the lumped mass for computation of the geometric rigid body vectors used in the modal participation factor calculation. These vectors are mass orthogonalized, and use of the consistent mass matrices for these efforts, especially when there are MPCs can be complicated in parallel. There is a small error introduced when the modes are computed using a consistent mass, and the rigid body vectors use a lumped mass. Refining the mesh reduces the problem, but most accurate results are obtained when the lumped mass is used (see [Section 2.4.4](#)).

3.18. ModalFrf Solution Case

Parameter	Type	Default	Description
nmodes	<i>Integer</i>	10	Number of modes to extract. See Section 3.10.1
shift	<i>Real</i>	-1.0e6	Shift to apply to matrix system to allow solving singular systems. See Section 3.10.2
untilfreq	<i>Real</i>	Inf	Target frequency to reach. See Section 3.10.2.1
ModalFilter	<i>string</i>	none	Modal filter to define modes to retain. See Section 3.10.2.2
complex	<i>True/False</i>		If true require use of a complex eigen solution. If false require use of a real valued eigen solution. If unset use the existing eigen solution.
lfcutoff	<i>Real</i>	-Inf	Exclude any modes below this frequency from the modal computation. Often used to exclude rigid body modes.
usemodalaccel			If set, use modal acceleration method
nrbms	<i>Integer</i>	0	Number of rigid body modes, needed for usemodalaccel
write_files	<i>all/ none/ output/ frequency</i>	all	Controls which result files are written during this solution.

Table 3-59. – ModalFrf Solution Case Parameters.

Option **ModalFrf** is used to perform a modal superposition-based frequency response analysis. In other words, **ModalFrf** provides an approximate solution to the Fourier transform of the equations of motion. If $\bar{u}$ is the Fourier transform of the displacement, u , and $\bar{f}$ is the Fourier transform of the applied force, then

$$(K + i\omega C - \omega^2 M) \bar{u} = \bar{f}(\omega).$$

If the damping matrix is zero, or if it can be diagonalized by the undamped modes, then **ModalFrf** uses the undamped modes for the superposition. Otherwise, for general damping matrices C , complex modes are used for the superposition. In either case,

ModalFrf is performed in a multicasel approach, where the modes (real or complex) are computed in a first case, and then **ModalFrf** is computed in a subsequent case.

Modal damping can be applied with either real eigenvalues computed by **eigen** or complex eigenvalues shapes computed by **qevp**. However, proportional damping is currently available only with real modes. For more details on damping, see Section 4.8.

3.18.0.1. ModalFrf with Real-Valued Modes In the case where the undamped real modes are used for the superposition, two options are available for the **ModalFrf** solution: the modal displacement method, and the modal acceleration method. In the case when complex modes are used, the modal displacement method is available. In both the modal displacement and modal acceleration methods, the approximate solution is found by linear modal superposition. Once the modes have been computed, there is little cost in computation of the frequency response. The solution does suffer from modal truncation, but in the case of the modal acceleration method, a static correction term partially accounts for the truncated high frequency terms. Thus, that method is generally more accurate than the modal displacement method. The most accurate approach, though also the most computationally expensive, is the **directfrf** method described in Section 3.7.

For real modes using the modal displacement method, the relation used for modal frequency response is given below.

$$\bar{u}_k(\omega) = \sum_j \frac{\phi_{jk} \phi_{jm} \bar{f}_m(\omega)}{\omega_j^2 - \omega^2 + 2i\gamma_j \omega_j \omega}$$

Here $\bar{u}_k$ is the Fourier component of displacement at degree of freedom k , ϕ_{jk} is the eigenvector of mode i at dof k , and ω_j and γ_j represent the mode frequency and associated fractional modal damping respectively. In the case of complex modes, the equations need to be linearized as described in subsection Quadratic Modal Superposition section Solution Procedures of the Theory Manual.

For the modal acceleration method, the procedure for computing the modal frequency response is more complicated. The response is split into the rigid body contributions, and the flexible contributions. The number of global rigid body modes must be specified in the input file. Also, see subsection Modal Frequency Response Methods section Solution Procedures of the Theory Manual.

The modal acceleration method more accurately computes the poles (or peaks) of the response and much more accurately compares the zeros of a function. The cost is an additional factor and solve. It can be used on floating structures, and the additional factor involves the stiffness terms (which are singular) and has no mass terms to stabilize the solution. Thus, it may be much more difficult to perform than solve than the other solves involved in the eigen analysis. In eigen analysis we recommend a negative shift for floating structures to remove the singularity associated with rigid body modes. No such approach is possible if you are using the modal acceleration method. Thus, significant “tweaking” of the linear solver parameters may be required to accurately determine the global rigid body modes required for success of this method.

The force function must be explicitly specified in the load section, and MUST have a “function” definition. Note that the force input provides the real part of the force at a given frequency, i.e. it is a function of frequency, not of time.

The following table gives the parameters needed for **ModalFrf** section.

The **nmodes** parameter controls the eigen analysis (see Section 3.10). The optional keyword, **usemodalaccel**, is used to determine whether to use the modal displacement or the modal acceleration method. If this keyword is specified, modal acceleration is used, otherwise the modal displacement method is invoked. If **usemodalaccel** is used, then the number of global rigid body modes must be specified using **nrbms**.

The parameters **freq_step**, **freq_min**, and **freq_max** are used to define the frequencies for computing the shock response spectra. They are identified in the **frequency** section with the application region (see Section 7.4). The range of the computed frequency spectra is controlled by **freq_min** and **freq_max**, while **freq_step** controls the resolution. The accuracy of the computed spectra does not depend on the magnitude of **freq_step**. This parameter controls the quantity of output. Examples are shown in inputs 3.9, 3.10 and 3.16.

We note that, in addition to the output that is sent to the **frq** file, output is also written to the **Exodus** file during a ModalFrf, provided that the keywords are specified in the **outputs** section. If nothing is specified in the **outputs** section, then nothing is written to the **Exodus** output files.

In the case of undamped modes, the following is a multcase example of how the ModalFrf could be specified.

```

Solution
  case eig
    eigen nmodes=7
    shift=-1e5
  case out
    ModalFrf
end
Frequency
  freq_step=300
  freq_min=100
  freq_max=2500
  nodeset=12
  acceleration
end

```

Input 3.9. ModalFrf Example Input

3.18.0.2. ModalFrf with Complex Modes In the case when complex modes are used, the modal displacement method is available. In this case the **qevp** solution case is used to compute the modes. There are currently three methods that can be used with the **qevp** solution case, and they are the **sa_eigen** method, the **Anasazi** method, and the **ceigen** method. For more details, we refer to Section 3.24.0.1. We note that in the case of complex modes, modal superposition is currently implemented for the **sa_eigen** method and the **Anasazi** method. The **ceigen** method does not support a subsequent modal superposition.

Also, when computing the complex modes in preparation for a modal superposition, we recommend using the **reorthogonalization** flag. When turned on, this flag searches for repeated modes and reorthogonalizes the eigenvectors of those modes. In many cases, repeated modes coming out of the eigen solvers are linearly independent, but not orthogonal. For more details, we refer to Section 3.24.0.1.

In the case of complex modes, the following is an example.

```

Solution
  case qevp
    qevp
    method = sa_eigen
    reorthogonalize = Y
    nmodes=20
    nmodes_acoustic = 5
    nmodes_structural = 5
  case out
    ModalFrf
end

```

```

        complex = y
    end
    Frequency
        freq_step=300
        freq_min=100
        freq_max=2500
        nodeset=12
        acceleration
    end

```

Input 3.10. Complex ModalFrf Example Input

3.19. *ModalFRF-Inverse Solution Case*

Parameter	Type	Default	Description
nmodes	<i>Integer</i>	10	Number of modes to extract. See Section 3.10.1
shift	<i>Real</i>	-1.0e6	Shift to apply to matrix system to allow solving singular systems. See Section 3.10.2
untilfreq	<i>Real</i>	Inf	Target frequency to reach. See Section 3.10.2.1
ModalFilter	<i>string</i>	none	Modal filter to define modes to retain. See Section 3.10.2.2
lfcutoff	<i>Real</i>	-Inf	Exclude any modes below this frequency from the modal computation. Often used to exclude rigid body modes.

Table 3-60. – ModalFRF-Inverse Solution Case Parameters.

The `modalf rf-inverse` solution method is used to solve an inverse problem (see Sec. [3.39](#)) for a modal frequency response analysis. The modal FRF method is similar to the direct FRF method, except the user must specify the number of modes `nmodes`. As in a forward solution, most of the parameters in an inverse modal frequency response analysis are found in other sections, and as in a `directfrf-inverse` problem, the user provides complex displacements and/or acoustic pressures at a set of nodes in the model.

The forward problem is defined in equation (3.8). The inverse equation is identical, but must be solved with optimization subject to regularization because measurements are available only at a subset of the analysis degrees of freedom.

The basic requirements for a `modalfrf-inverse` simulation are as follows:

Optimization: Control over the optimization problem is specified in the **optimization** block. See Sec. 3.39.1 for further details.

Inverse-Problem: The **inverse-problem** block provides the connection to the measurement data. It is also where **design_variable** is specified (e.g., load, material, etc.). See Sec. 3.39 for details.

Truth Table: The truth table **data_truth_table** from the **inverse-problem** block is a list of the indices of the global node numbers (a.k.a. target nodes) where displacements or acoustic pressures are measured. See Sec. 3.39.3 for file format details.

Data File: Experimentally determined “target” displacements are read from **real_data_file** and **imaginary_data_file** specified in the **inverse-problem** block. See Sec. 3.39.3 for file format details.

Frequency: The frequencies at which the problem is solved are specified in the **frequency** block (Sec. 7.4).

Sierra/SD uses the Rapid Optimization Library (ROL) for solving optimization problems. During the optimization solution ROL writes an output file, *ROL_Messages.txt* that contains convergence information. Section 3.39.8 contains a discussion of the output file that is written by ROL.

3.19.1. Load Identification

```
solution
  modalfrf-inverse
  nmodes 100
end
inverse-problem
  design_variable = load
  data_truth_table = ttable.txt
  real_data_file = data.txt
  imaginary_data_file = data_im.txt
end
optimization
  ROLmethod = linesearch
  LSstep = Newton-Krylov
  LS_curvature_condition = null
  Max_iter_Krylov = 50
```

```

Use_FD_hessvec = false
Use_inexact_hessvec = false
opt_tolerance = 1e-10
opt_iterations = 2
end
loads
  sideset 6
    inverse_load_type = spatially_constant
    pressure=1
    function = 1
  sideset 6
    inverse_load_type = spatially_constant
    ipressure=1
    function = 2
end
function 1
  type linear
  data 1 3
  data 2 4
end
function 2
  type linear
  data 1 5
  data 2 6
end

```

Input 3.11. Modal frequency response load identification example input

Specifying `design_variable = load` applies inverse methods to determine sideset loads which best correspond with the measured displacements and/or acoustic pressures provided by the user. The material and model parameters do not change during the solution. For structures, the loads are pressures or tractions³, and for acoustics, the loads are acoustic accelerations. Note that for structures, inversion is based on the signed magnitudes of the tractions; the direction of each traction is fixed.

An example input deck is given in input 3.11. In addition to the input blocks discussed in the beginning of this section, there are several others specific to `design_variable = load`:

Loads: See Sec. 3.39.6 for a description of the inverse parameters in the `loads` block for load identification problems.

Frequency: For unknown loads, the frequencies at which the problem is solved are independent. That is, a separate load identification is performed at each frequency.

³Moments and point forces are not currently supported.

Section 3.39.7 contains a discussion of the current limitations with inverse load methods.

3.19.2. PSD Load Identification

```
solution
  modalfrf-inverse
  nmodes 100
end
inverse-problem
  design_variable = psd_load
  data_truth_table = ttable.txt
  data_file = data.txt
end
optimization
  ROLmethod = linesearch
  LSstep = Newton-Krylov
  LS_curvature_condition = null
  Max_iter_Krylov = 50
  Use_FD_hessvec = false
  Use_inexact_hessvec = false
  opt_tolerance = 1e-10
  opt_iterations = 2
end
loads
  sideset 6
    inverse_load_type = spatially_constant
    pressure=1
    function = 1
  sideset 6
    inverse_load_type = spatially_constant
    ipressure=1
    function = 2
end
function 1
  type linear
  data 1 3
  data 2 4
end
function 2
  type linear
  data 1 5
  data 2 6
end
```

Input 3.12. Modal frequency response PSD load identification example input

Specifying `design_variable = psd_load` applies inverse methods to determine a *load PSD* (power spectral density) as an output when provided with the PSD of acoustic pressures or structural displacements and a finite element model. The input is similar to modal FRF load identification, with two exceptions. First, the design variables must be defined such that there are independent real and imaginary parts of the force, traction, or pressure. Thus, there should be twice the number of design variables as the dimension of load PSD, corresponding to real and imaginary parts of the load. These design variables must be entered into the input file in the order of the load index. Furthermore, for each load index, the design variable for the real part should be immediately followed by the design variable for the imaginary part. Second, a different data file format is required, although the truth table format is identical. See Sec. 3.39.3 for further details.

An example input deck is given in input 3.12. In addition to the input blocks discussed in the beginning of this section, there are several others specific to `design_variable = psd_load`:

Data File: Experimentally determined “target” response PSDs are read from the `psd_data_file` described in Sec. 3.39.3.

Loads: See Sec. 3.39.6 for a description of the inverse parameters in the `loads` block for load identification problems. `loads` block.

Frequency: For unknown loads, the frequencies at which the problem is solved are independent. That is, a separate load identification is performed at each frequency.

3.20. ModalRanVib Solution Case

Parameter	Type	Default	Description
nmodes	<i>Integer</i>	10	Number of modes to extract. See Section 3.10.1
shift	<i>Real</i>	-1.0e6	Shift to apply to matrix system to allow solving singular systems. See Section 3.10.2
untilfreq	<i>Real</i>	Inf	Target frequency to reach. See Section 3.10.2.1
ModalFilter	<i>string</i>	none	Modal filter to define modes to retain. See Section 3.10.2.2
noSVD			Selection of method for von Mises stress computation. Less expensive, and provides additional stress moments.
lfcutoff	<i>Real</i>	0.1	Exclude any modes below this frequency from the modal computation. Often used to exclude rigid body modes.
truncation-Method	<i>none/ displacement/ acceleration</i>	none	Truncates modes with low activity
keepModes	<i>integer</i>	Inf	Keep the specified number of modes, kept modes are selected based on the highest modal activity.
checkSMatrix	<i>true/false</i>	true	Evaluate the correlation matrix S . This matrix is generated at each frequency, and must be positive semi-definite. A correlation matrix that is not positive semi-definite results in negative PSD results, which are not physically meaningful. Note, matrix evaluations are enabled if PSD output is requested.

Table 3-61. – ModalRanVib Solution Case Parameters.

Option **modalranvib** is used to perform a modal-superposition-based random vibration analysis in the frequency domain. Root-mean-square (RMS) outputs, including von Mises stress, are computed for a given input random force function. The resulting power spectral density functions may also be output at locations specified in the **frequency** section. The forcing functions (one for each input) must be explicitly specified in the **ranload** section (6.3.20). It must reference a **matrix-function** definition (see Section 2.8.17).

modalranvib accepts **eigen** solution parameters such as **nmodes**. These are used to compute an eigen solution if one does not yet exist. The **eigen** parameters should not be used in a multicas solution.

The *optional* keyword **noSVD** determines the method used to compute the RMS von Mises stress output. If **noSVD** is specified, then the simpler method which does not use a singular value decomposition is used. Additionally, that simpler method causes the second and fourth moments associated with von Mises stress to be computed and to be written to Exodus output. (The RMS von Mises stress and these two moments, along with the appropriate material properties, can be used in a manner suggested in [74] and discussed in [75] to estimate fatigue life in broad-band random excitation. Also, see the **fatigue** section (3.14). However, this method provides no information about the statistics of the stress. Only the RMS value and moments are reported.

The *optional* keyword **lfcutoff** provides a low frequency cutoff for random vibration processing. Usually, rigid body modes are *not* included in this type of calculation if RMS stress is computed. The **lfcutoff** provides a frequency below which the modes are ignored. The default for this value is 0.1 Hz. Thus, by default rigid body modes are not included in random vibration analysis. A large negative value will include all the modes.

The *optional* keyword **TruncationMethod** provides control over selection of the retained modes. By default, modes are retained if they have any contribution to the stress. As stresses are proportional to displacement, the default method is **displacement**. It is possible to not truncate at all (**none**), or to truncate based on accelerations (**acceleration**). Acceleration contributions are weighted to higher frequencies. Often zero energy modes contribute to a bad truncation, and a preferred means of controlling the truncation is to use the **lfcutoff** parameter and to ensure the integration does not go to zero frequency.

The *optional* keyword **keepmodes** is a method of truncating modes. By default, its value is **nmodes**. If a value is provided, the modes with the lowest modal activity will be truncated until **keepmodes** remain. Note that this procedure is much different from truncating the higher-frequency modes. Modal truncation is important because all the operations compute responses that require $O(N^2)$ operations. Even if **keepmodes** is not entered, modes with modal activity less than 1 millionth of the highest active mode will be truncated.

The parameters **freq_step**, **freq_min**, and **freq_max** are used to define the frequencies for computing the random vibration spectra. They are identified in the **frequency** section along with an optional application region (see Section 7.4). The range of the computed

frequency spectra is controlled by **freq_min** and **freq_max**, while **freq_step** controls the resolution. The accuracy of the computed spectra *does* depend on the magnitude of **freq_step** since it is used in the frequency domain integration. Examples are presented in inputs 3.9, 3.10 and 3.16.

In random vibration, the **frequency** block serves two purposes. First, it is used for the integration information for the entire model. Thus, Γ_{qq} for the referenced papers^{59,50} is integrated over frequency and used for all output. In addition, if an output region is specified in the **frequency** block, output acceleration and displacement power spectra may be computed for the given region at the required frequency points. At this time, **acceleration** and/or **displacement** may be specified in the frequency block for random vibration analysis. This output is described in more detail below.

Random vibration analysis is trickier than most input. A number of blocks must be specified.

1. The **Solution** block must have the required input for **eigen** analysis, and the keyword **modalranvib**.
2. The **RanLoads** block contains a definition of the spectral loading input matrix and other input. Note that the input, S_{FF} is separated into frequency and spatial components. The spatial component is specified here using **load** keywords. See Section 6.3.20. The spectral component is referred to here, but details are provided in the matrix-function section.
3. The **matrix-function** section contains the spectral information on the loading. It references functions for the details of the load. The real and imaginary function identifiers for this input are specified here (2.8.17).
4. There must be a **function** definition for each referenced spectral function. Functions of time or frequency are further described in Section 2.8.
5. There must be a **frequency** block that is used for integration and optionally also for output of displacement and acceleration output. See Section 7.4.
6. As an undamped system is singular, some type of **damping** is required. Modal damping terms are required.⁴ See Section 4.8.
7. **boundary** conditions are supplied in the usual way, but the standard **loads** block is replaced by the input in the ranloads section. The loads block will be quietly ignored in random vibration analysis.
8. The **outputs** and **echo** sections will require the keyword **vrms** for output of RMS von Mises stress. If the **stress** keyword is also found, then the natural stresses for solid elements will be output.⁵ The keywords **rotational_displacement** and **rotational_acceleration** in the outputs and frequency sections will output the

⁴Proportional damping, such as is applied with the **alpha** and **beta** terms, will *not* work in **modalranvib**.

⁵These stresses are linear functions of the displacement.

corresponding RMS and PSD quantities respectively. Quantities output are listed in Table 3-62.

All other input should remain unchanged.

Keyword	Output Variable	Description
Vrms	vrms	Root Mean Squared Von Mises Stress
	D1...D5	von Mises Stress SVD moments. Details in. ⁶⁰
	M0, M2, M4	von Mises Stress moments. noSVD
	Xrms	X component of RMS displacement
	Yrms	Y component of RMS displacement
	Zrms	Z component of RMS displacement
	Axrms	X component of RMS acceleration
	Ayrms	Y component of RMS acceleration
	Azrms	Z component of RMS acceleration
	RotXrms	Rotational X component of RMS displacement
rotational_displacement	RotYrms	Rotational Y component of RMS displacement
	RotZrms	Rotational Z component of RMS displacement
rotational_acceleration	RotAxrms	Rotational X component of RMS acceleration
	RotAyrms	Rotational Y component of RMS acceleration
	RotAzrms	Rotational Z component of RMS acceleration

Table 3-62. – ModalRanVib Output to Exodus File. The stress spectral moments are neither computed nor output if **noSVD** is selected. The stress moments are available if **noSVD** is selected, and may be used for fatigue. The RMS values of displacements and acceleration are components of a Hermitian tensor. See Section 3.20.0.1 for details.

3.20.0.1. Power spectral density When requested in the frequency block, one output from the random vibration analysis is a power spectral density or PSD (for displacement or acceleration). The power spectral density is a measure of the output content over a frequency band, and usually measured in units of cm^2/Hz or some similar unit. Acceleration PSDs are often measured in units of g^2/Hz .⁶

Like the input cross spectral forces, the output quantities are Hermitian, with 9 independent translational quantities per frequency, at each output node for each type of output. The method for transforming these quantities in alternate coordinate systems are in subsection Modal Frequency Response Methods section Solution Procedures of the Theory Manual. The rotational terms can be requested using the keywords **rotational_displacement** or **rotational_acceleration**. Note that the cross

⁶Power spectral density output is requested in the **frequency** block. A collection of nodes is indicated and the **displacement** or **acceleration** keyword is entered. PSDs of displacement or acceleration are available.

correlation terms for rotation are not output.

$$\begin{bmatrix} A_{xx} & A_{xy} + iA_{xyi} & A_{xz} + iA_{xzi} & - & - & - \\ A_{xy} - iA_{xyi} & A_{yy} & A_{yz} + iA_{yzi} & - & - & - \\ A_{xz} - iA_{xzi} & A_{yz} - iA_{yzi} & A_{zz} & - & - & - \\ - & - & & A_{Rot_x Rot_x} & - & - \\ - & - & & - & A_{Rot_y Rot_y} & - \\ - & - & & - & - & A_{Rot_z Rot_z} \end{bmatrix}$$

Because the inputs are specified in terms of force cross-correlation functions, the standard procedure for applying loads often involves application of a large concentrated mass at the input location. The force may then be applied to the mass and the acceleration determined from $a = f/m$, where we assume that m is much larger (100 to 1000 times larger) than the mass of the remainder of the structure. Some confusion can arise in the scaling of the force.

The output PSD for acceleration is defined as follows.

$$G_{ij} = H_{ki}^\dagger S_{kl} H_{lj}$$

where H_{lj} is the transfer function giving a_j/f_l , and S_{kl} is a power spectral density input (with units of force²/Hz).

Consider a single input, i.e. $k = l$, and with $f_k = m_k a_k$.

$$G_{ij} = H_{ki}^\dagger \langle m_k a_k, a_k m_k \rangle H_{lj} \quad (3.22)$$

$$= (m_k^2) H_{ki} \langle a_k, a_k \rangle H_{kj} \quad (3.23)$$

Thus, the acceleration PSD must be multiplied by the square of the mass to get the force PSD. **Sierra/SD** applies the scale factor to the spatial force distribution (which is also squared), so the scale factor in **Sierra/SD** should be m_k .

3.21. ModalShock Solution Case

Parameter	Type	Default	Description
nmodes	<i>Integer</i>	10	Number of modes to extract. See Section 3.10.1
shift	<i>Real</i>	-1.0e6	Shift to apply to matrix system to allow solving singular systems. See Section 3.10.2
untilfreq	<i>Real</i>	Inf	Target frequency to reach. See Section 3.10.2.1
ModalFilter	<i>string</i>	none	Modal filter to define modes to retain. See Section 3.10.2.2
srs_damp	<i>Real</i>	0.03	Damping coefficient used for the shock response spectra calculation

Table 3-63. – ModalShock Solution Case Parameters.

The **modals shock** solution method is used to perform a modal-superposition-based implicit transient analysis followed by computation of the shock response spectra for the degrees of freedom in a specified node set. More information about the about shock response spectra solution cases is given in Section [3.31](#) describing implicit-transient-based SRS.

A **frequency** block must also be included in the input file for **modals shock** solution cases to define the frequencies and nodesets for computing the shock response spectra (see Section [7.4](#)). The parameters **freq_step**, **freq_min**, and **freq_max** are used to define the frequencies for computing the shock response spectra. They are identified in the **frequency** section along with an application region (see Section [7.4](#)). The range of the computed frequency spectra is controlled by **freq_min** and **freq_max**, while **freq_step** controls the resolution. The accuracy of the computed spectra is not dependent on the magnitude of **freq_step**. This parameter controls the quantity of output. Examples are presented in inputs [3.9](#), [3.10](#) and [3.16](#).

3.22. *ModalTransient Solution Case*

Parameter	Type	Default	Description
time_step	<i>Real</i>		Time step size. See Section 3.32.0.1
nsteps	<i>Integer</i>		Number of time steps to take. See Section 3.32.0.2
start_time	<i>Real</i>	0.0	Solution case start time. See Section 3.32.0.3
nskip	<i>Integer</i>	1	Results output frequency. See Section 3.32.0.4
rho	<i>Real</i>	1	Select time integrator. See Section 3.32.1.1
load	<i>Integer</i>		Load to apply during solution case. See Section 3.32.0.5
write_files	<i>all none output history</i>	all	Controls which result files are written during this solution.
lfcutoff	<i>Real</i>	-Inf	Exclude any modes below this frequency from the modal computation. Often used to exclude rigid body modes.
flush	<i>integer</i>	50	Defines how often results are written to the exodus results file. See Section 2.4.1

Table 3-64. – ModalTransient Solution Case Parameters.

Option **modaltransient** is used to perform a modal-superposition-based implicit transient analysis. Damping for the model is defined in Section [4.8](#).

The parallel solution of modal transient may be slower than expected: while the eigen solution parallelizes well, there is not enough computation to parallelize the modal calculation. Only those DOFs requested for output are included in the solution. Thus, requesting output on only a small subset of the model can improve solution speed.

*If output is required at few locations, we recommend setting **write_files** to **history**, or specifying an empty **outputs** definition. Note that an empty **outputs** definition applies to all solution cases, whereas **write_files** applies only to the current solution case. Forces and displacements on modal degrees of freedom are also available via the keyword **modalvars** in the **echo** definition.*

modaltransient supports restart in the **eigen** part of the analysis, the **modaltransient** part, or both. In the latter case, two things would happen first. Any modes from the modal restart file are read, and the time history data from any previous transient restart files (direct or modal) is read. Afterwards stepping in time continues. A general discussion of restart is given in Section 2.4.2, which also discusses restart test coverage (Table 2-11).

An example of restart with the **modaltransient** solution is given below. In this case, the eigen solution is restarted prior to the modaltransient solution. The eigen solution would proceed as follows

```
Solution
  case eigen
    eigen
      nmodes 10
      restart=write
  end
```

and, subsequently, the **eigen** restart and **modaltransient** would be:

```
Solution
  case eigen
    eigen
      nmodes 20
      restart=read
  case modaltrans
    modaltransient
      nsteps 100
      time_step 1.0e-3
      restart=write
  end
```

The previous **modaltransient** cannot restart from the last computed time step when computing additional time steps. To restart it is necessary to add option **write** in the **modaltransient** case. For example, one could then do the following:

```
Solution
  case 'eigen'
    eigen
    nmodes 20
    restart=read
  case 'modaltrans'
    modaltransient
    nsteps 100
    time_step 1.0e-3
    restart=write
  case 'modaltrans'
    modaltransient
    nsteps 200
    time_step 1.0e-3
    restart=read
end
```


3.23. *ModalTransient-Inverse Solution Case*

Parameter	Type	Default	Description
nmodes	<i>Integer</i>	10	Number of modes to extract. See Section 3.10.1
shift	<i>Real</i>	-1.0e6	Shift to apply to matrix system to allow solving singular systems. See Section 3.10.2
untilfreq	<i>Real</i>	Inf	Target frequency to reach. See Section 3.10.2.1
ModalFilter	<i>string</i>	none	Modal filter to define modes to retain. See Section 3.10.2.2
lfcutoff	<i>Real</i>	-Inf	Exclude any modes below this frequency from the modal computation. Often used to exclude rigid body modes.
time_step	<i>Real</i>		Time step size. See Section 3.32.0.1
nsteps	<i>Integer</i>		Number of time steps to take. See Section 3.32.0.2
start_time	<i>Real</i>	0.0	Solution case start time. See Section 3.32.0.3
nskip	<i>Integer</i>	1	Results output frequency. See Section 3.32.0.4
rho	<i>Real</i>	1	Select time integrator. See Section 3.32.1.1
load	<i>Integer</i>		Load to apply during solution case. See Section 3.32.0.5
write_files	<i>all/ none/ output/ history</i>	all	Controls which result files are written during this solution.

Table 3-65. – ModalTransient-Inverse Solution Case Parameters.

```

solution
  modaltransient-inverse
  nsteps = 100
  time_step = 1e-3
  nskip = 1
  nmodes = 100

```

```

end
inverse-problem
  design_variable = load
  data_truth_table = ttable.txt
  data_file = dataReal.txt
  tikhonovParameter 1.0e-5
end
optimization
% optimization_package = ROL_lib
  ROLmethod = linesearch
  LSstep = Newton-Krylov
  LS_curvature_condition = null
  Max_iter_Krylov = 20
  Use_FD_hessvec = false
  Use_inexact_hessvec = false
  opt_tolerance = 1e-12
  opt_iterations = 2
end
loads
  sideset 301
    inverse_load_type = spatially_constant
    pressure=10
    function = 1
end

```

Input 3.13. Modal transient load identification example input

The `modaltransient-inverse` solution method is used to solve an inverse problem (see Sec. 3.39) for a modal transient analysis. In this solution method, only load identification is supported. Specifying `design_variable = load` applies inverse methods to determine sideset loads with best correspond with measured displacements and/or acoustic pressures provided by the user in the modal time domain. This capability differs from load identification in a `transient-inverse` problem (Sec. 3.33.1) only in that modal superposition is used to reduce computation time. As with forward analysis, the `modaltransient` solution will converge to the direct solution as the number of modes increases. See the SierraSD verification manual for an example of this convergence[54].

The parameters for load identification in a direct `transient-inverse` problem also apply in the `modaltransient-inverse` case. The latter also requires the parameter `nmodes`, the number of eigen modes calculated in the forward solve, as well as any additional parameters needed for the eigen solution case. The eigen modes need only be calculated once, and then can be re-used for each inverse iteration. Please refer to Sec. 3.39 for a general description of inverse methods. Note that the Tikhonov parameter can be used to mollify instability in the early time history.

An example is shown in input 3.13. The following input blocks are needed for `modaltransient-inverse` with `design_variable = load`:

Optimization: Control over the optimization problem is specified in the **optimization** block. See Sec. 3.39.1 for further details.

Inverse-Problem: The **inverse-problem** block provides the connection to the measurement data. It is also where `design_variable = load` must be specified. See Sec. 3.39 for details.

Truth Table: The truth table (`data_truth_table` from the **inverse-problem** block is a list of the indices of the global node numbers (a.k.a. target nodes) where displacements or acoustic pressures are measured. See Sec. 3.39.3 for file format details.

Data File: Experimentally determined “target” displacements are read from `data_file` specified in the **inverse-problem** block. See Sec. 3.39.3 for file format details.

Loads: See Sec. 3.39.6 for a description of the inverse parameters in the **loads** block for load identification problems.

Section 3.39.7 contains a discussion of the current limitations with inverse load methods.

Sierra/SD uses the Rapid Optimization Library (ROL) for solving optimization problems. During the optimization solution ROL writes an output file, *ROL_Messages.txt* that contains convergence information. Section 3.39.8 contains a discussion of the output file that is written by ROL.

3.24. Q EVP Solution Case

Parameter	Type	Default	Description

Table 3-66. – Q EVP Solution Case Parameters.

Note: Options for the **qevp** solution case depend on the type of quadratic eigen solver being used and are documented below.

3.24.0.1. Quadratic Eigenvalue Methods Comparison The quadratic eigenvalue problem is defined as,

$$(K + D\lambda + M\lambda^2)u = 0 \quad (3.24)$$

The solution of the quadratic eigenvalue problem (3.24), has applications in a variety of physics solutions including coupled structural acoustics, general eigenvalue systems with damping, and gyroscopic systems for rotating structures. Various methods have been developed to address the solution to these problems. The solution to the problem is

difficult, and knowledge of the types of systems encountered can help significantly in addressing the robustness of each of the methods. The methods are listed and described in the following paragraphs. Table 3-67 lists recommended procedures for different problem sets.

Anasazi: It is possible to use the Anasazi method, although more testing is needed. It can be used to address two problem areas, 1) the coupled structural acoustics problem, and 2) gyroscopic systems from rotating frames. Currently, it requires that both the mass and the stiffness matrix be non-singular. Previous versions of **Sierra/SD** used the solution case **qevp** with no method keyword to denote the Anasazi method, and it is the default method to keep consistent with this syntax.

A couple of the parameters for the Anasazi solver for quadratic eigenvalue problems are described in the Table 3-50. Here shifts are not supported, and a warning message may be avoided by setting **shift** to zero. Restarts are not supported either; set **implicit_restarts** to 1. Restarts make the capability easier to use. Without restarts the user is required to set the subspace size, **subspace_size**, to be sufficiently large. The way that the algorithm works is to compute all the modes, and then compare a relative residual to **eig_tol**. If the residuals are large, the modes are not returned to the user. It can be helpful to use a larger value of **eig_tol** than the default, say 10^{-8} . In this situation, it is helpful to set **anverbosity** to a large value, say And then examine other diagnostic information, including the data written to the screen, to ascertain the accuracy of the modes.

CEIGEN: The **ceigen** method uses methods in ARPACK to solve the quadratic eigenvalue problem. Of methods in **Sierra/SD**, it is the oldest, and probably the least robust.

SA_EIGEN: The **sa_eigen** method solves a coupled structural acoustics problem by solving a linear, uncoupled eigen problem on each of the domains, and using them as a basis to reduce the coupled equations to a dense system. The dense system is solved using LAPACK routines. The method is applicable to structural/acoustic systems. It is robust. Modal truncation can introduce significant errors. Some solutions can fail (or convergence may be slow) because the decomposition tools know nothing about the two domains.

PROJECTION_EIGEN: The **projection_eigen** method solves the quadratic eigenvalue problem by projecting the problem into a subspace corresponding to the real-valued modes. This smaller subspace is constructed by neglecting the damping matrix, symmetrizing the stiffness matrix, and solving the eigenvalue problem,

$$Ku = \lambda Mu. \quad (3.25)$$

This smaller problem is then used as a basis for solving the original quadratic eigenvalue problem, which takes the form

$$Ku + \lambda Cu + \lambda^2 Mu = 0 \quad (3.26)$$

The original quadratic eigenvalue problem is then pre and post multiplied by the eigenvectors obtained from the subspace eigenvalue problem. This results in a small quadratic eigenvalue problem which is then solved with a LAPACK method. Finally, the modes from the reduced space are projected to the space corresponding to the original quadratic eigenvalue problem.

As with the **sa_eigen** method, truncation error is a concern with the **projection_eigen** method. The more modes one takes, the smaller the truncation error.

Problem	Ceigen	SA_eigen	Anasazi	Projection_eigen
Damped Systems	Acceptable	Acceptable	Fails	Acceptable
structural acoustics	Fails	Acceptable	Acceptable	Acceptable
Rotational systems	N/A	N/A	Acceptable	Acceptable
Damped str/acoustics	Fails	Acceptable	Fails	Acceptable

Table 3-67. – A 2005 Comparison of Quadratic Eigenvalue Problem Methods. Although Ceigen and Anasazi have changed substantially since 2005, this table has not been updated to reflect those changes.

Use the keyword **method** followed by the name of the method (Anasazi, ceigen, sa_eigen, Projection_eigen) to select a **qevp** method. Below is a more detailed description of each **qevp** method, their parameters, and examples of how to use them.

3.24.1. Anasazi

The **Sierra/SD** interface to the Trilinos package Anasazi solves the quadratic eigenvalue problem defined as

$$(K + D\lambda + M\lambda^2)u = 0 \quad (3.27)$$

See Section 3.24.0.1 for a comparison of these methods for this problem. As currently implemented, the Anasazi method applies to systems with a non-singular mass and stiffness matrix. The damping matrix may be asymmetric. Options for input are described in Table 3-68. An example is given below.

Table 3-68. – Options for **qevp** Anasazi Solutions.

Option	Argument	Default	Comment
nmodes	<i>Integer</i>	10	number of modes
shift	<i>Real</i>	0	ignored
reorthogonalize	Yes/No/Full	“Yes”	Reorthogonalize vectors
check_diag	Yes/No/Full	“Yes”	Check that vectors diagonalize linearized system
ANverbosity	<i>Integer</i>	17	Anasazi verbosity
ANblocksize	<i>Integer</i>	1	Anasazi Block Size

Although modal truncation methods are an industry standard, the results sometimes are inaccurate. A solution may not to change when the number of modes increases, and still be inaccurate. Verification is entirely up the user in every case with modal truncation methods.

Users have never prioritized enhancing the usability of Anasazi or CEigen. In addition to the shift, users must select **eig_tol** . But discerning an effective value depends on the problem, especially on the shift and the linear solver accuracy. Too large a threshold ($> 1.e - 4$) degrades solution accuracy. Too small a threshold ($< 1.e - 13$) leads to divergence. However, with these methods, diagnostic information is provided (written to the standard output stream) to guide users with problem configuration and solution verification.

```
Solution
  case qevp
    qevp
    method=Anasazi
    nmodes=14
    anverbosity=27
end
```

3.24.2. Damped Eigenvalue Problems

The **qevp** solution with “method=ceigen” is used to select complex eigen analysis using the ARPACK package. This computes the solution to the quadratic eigenvalue problem,

$$(K + D\lambda + M\lambda^2)u = 0 \quad (3.28)$$

Note that two other solution methods may also be used to evaluate the quadratic eigenvalue problem. Each of these methods has its strengths and weaknesses. A comparison of these methods is provided in Section 3.24.0.1.

The following table gives the parameters needed for complex eigen analysis. Additionally, the eigenvalue tolerance can be set with the **eig_tol** parameter in the parameters block (Section 2.3).

Parameter	Type	Default	Description
nmodes	<i>Integer</i>	100	Number of modes to compute, reported as complex conjugate pairs
viscofreq	<i>Real</i>	1e-6	Frequency at which to evaluate material damping

Table 3-69. – ceigen Solution Case Parameters.

Table 3-70. – Ceigen Tests.

Name	Description
ceig	stiffness proportional damping
ceig_visco	viscoelastic damping
ceig_dash	dashpot damping
steel_in_foam	complex mixed materials

The optional **viscofreq** keyword indicates the frequency at which the damping properties of viscoelastic materials will be computed. It must be non-negative. The **viscofreq** parameter can be confusing. In particular, viscoelastic materials typically have high damping at lower frequencies, and lower damping at high frequencies. The **viscofreq** parameter sets a frequency from which we estimate all the viscoelastic damping. Thus, if **viscofreq** is small, the damping is large. In particular, if **viscofreq** is below the glass transition frequency, then damping appropriate to the low frequency modes will be used. *This high value of damping is applied to the entire spectrum.* It is generally better to over-estimate **viscofreq** than to underestimate it.

The reason for this difficulty is that even linear viscoelastic materials generate a more complex equation than that shown in equation 3.28. With a single term in the Prony series, the equation of motion for a damped viscoelastic structure can be written in the frequency domain.

$$\left(K + D \frac{s}{s + \omega_g} + Ms^2 \right) u = f(s) \quad (3.29)$$

Here s is the *Laplace* transform variable and $\omega_g = 1/\tau$ is the reciprocal of the relaxation constant. This system is not a simple quadratic in s . Effectively, **viscofreq** approximates this system with the linearized system below.

$$\left(K + D \frac{s}{2\pi \cdot \text{viscofreq} + \omega_g} + Ms^2 \right) u = f(s) \quad (3.30)$$

Computation of quadratic eigen modes is much more difficult than real eigen analysis. The system of equations is more difficult, and more “tricks” must be used to resolve issues that are generated, such as decreasing **eig_tol**. Even the post processing can be complicated. Like real eigen analysis, one must request displacement output in the **output** section (see 7.1.13). The output file contains 12 fields (six real and six imaginary). Few post processing tools handle complex mode shapes. Also, see subsection Complex Eigen Analysis — Modal Analysis of Damped Structures section Solution Procedures of the Theory Manual.

Because of the difficulties with complex eigen analysis, it is important to understand the problems for which we have evaluated and tested it. The tests in the test suite are listed in Table 3-70.

3.24.3. SA_eigen

The **qevp** procedure with method **SA_eigen** provides a means of computing the modal response of a coupled structural acoustic system, using a modal truncation basis. The quadratic eigenvalue problem describing this system can be written as follows.

$$\left(\begin{bmatrix} K_s & 0 \\ 0 & K_a \end{bmatrix} + \lambda \begin{bmatrix} C_s & L \\ -\rho_a L^T & C_a \end{bmatrix} + \lambda^2 \begin{bmatrix} M_s & 0 \\ 0 & M_a \end{bmatrix} \right) \begin{bmatrix} \phi_s \\ \phi_a \end{bmatrix} = 0 \quad (3.31)$$

Here the subscripts refer to structural or acoustic domains, ρ_a is the density of the fluid and L is a coupling matrix. Note that for this formulation, ϕ_a represents the acoustic velocity potential, which relates to the time derivative of the acoustic pressure, $\phi_a = \nabla \dot{u}_a$. It helps to understand⁵² the capabilities and limitations of this analysis.

The **sa_eigen** method solves this system by solving for the uncoupled eigen modes in the two domains, using them as a basis to reduce the coupled equations to a dense system, and solving the dense system. Thus, it uses a modal reduction technique similar to the Craig-Bampton methods (section 3.5) to generate a dense system of equations that are solved and results propagated to the physical space. More details are available in the theory manual.

Options of the analysis are provided in Table 3-71, and an example is provided in input 3.14. Boundary conditions are applied exactly as for the generalized eigenvalue problem. Exterior, non-reflecting boundary conditions may be applied, but modal convergence is poorer. Loads are irrelevant. Output is complex, as for the **ceigen** case (3.24.2).

```
Solution
case saeig
  qevp
    method=sa_eigen
    nmodes=20
    nmodes_acoustic=50
    nmodes_structure=26
    acoustic_lfcutoff=-1
    structural_lfcutoff=-1
    sort method = frequency
end
```

Input 3.14. SA_Eigen Example

Limitations: This is a modal superposition method. The **Sierra/SD** interface to the Trilinos package Anasazi is a more complete but less robust method which does not rely on modal truncation. The **SA_eigen** method is accurate for many structural acoustic environments. Damping may be provided, but does tend to slow

Option	Args	Description
nmodes	<i>int</i>	Number of requested eigen modes
nmodes_acoustic	<i>int</i>	Number of free-free acoustic modes in the reduction. Defaults to 2·(nmodes).
nmodes_structure	<i>int</i>	Number of free-free structural modes in the reduction. Defaults to 2·(nmodes).
acoustic_lfcutoff	<i>Real</i>	Low frequency cutoff to filter acoustic modes. By default, all modes are retained
structural_lfcutoff	<i>Real</i>	Low frequency cutoff to filter structural modes. By default, all modes are retained.
shift	<i>Real</i>	Used to eliminate negative modes Eigen shift used in computation of the subregion modes. See 3.10.
sort method	<i>string</i>	magnitude: complex magnitude of λ frequency: Sort by frequency and then damping. damping: Sort by damping and then frequency. truefreq: Sort by frequency... avoiding zero energy round off. none: Multicase requires None
linearization	<i>int</i>	1 $A = [0 \ I; -K \ -C]; B = [\ I \ 0; 0 \ M]$ 2 $A = [\ -K \ 0; 0 \ M]; B = [\ C \ M; M \ 0];$ 4 $A = [\ 0 \ -K; M \ 0]; B = [\ M \ C; 0 \ M];$ These follow the linearizations in Tisseur
reorthogonalize	<i>string</i>	no: no reorthogonalization yes: reorthogonalize all modes full: check all modes
check_diag	<i>string</i>	no: no check for orthogonalization yes: check redundant modes full: check all modes

Table 3-71. – SA_Eigen Options.

convergence. The method also depends on the solution to separate structural and acoustic subregion eigen problems. These solutions are not as robust as full system eigen analysis. Please see the notes in the verification manual for convergence details. Table 3-72 summarizes the status of this procedure.

Low Frequency Cutoff: The parameters `acoustic_lfcutoff` and `structural_lfcutoff` remove low frequency modes before initiating the `qevp`. This will reduce the number of modes (`nmodes_acoustic` and `nmodes_structure`) in the analysis. Negative cutoff frequencies are allowed.

Analytic Reference	Verification Section	Tested	Parallel Test	User Test
39	54	Y	Y	some

Table 3-72. – Verification Summary for SA_Eigen.

Specialized Output: There are a few items that are output specifically for the *sa_eigen* procedures that can be helpful in assessing the solutions.

StructuralFraction It is useful to know which modes participate in which regions. This is computed as follows.

Let ϕ be the eigenvector computed on the reduced space. We divide ϕ into its structural and acoustic components. i.e.,

$$\phi = \begin{bmatrix} \phi_s \\ \phi_a \end{bmatrix}$$

We compute,

$$F_{structure} = \frac{\phi_s^\dagger \cdot \phi_s}{\phi_s^\dagger \cdot \phi_s + \phi_a^\dagger \cdot \phi_a} \quad (3.32)$$

where $\phi^\dagger$ represents the transpose and complex conjugate of ϕ . Note that these products are computed in the reduced space which has coordinates associated with each structural or acoustic eigen mode. In the reduced space, the mass matrix is identity, and the vector product, $\phi^\dagger \cdot \phi$ represents an energy norm.

AcousticFraction The acoustic fraction is the analogue of the structural fraction (eq. 3.32) applied the acoustic domain. It represents the portion of the system level complex eigen mode that is associated with the acoustic domain.

ErrorNorm We define a normalized modal energy residual.

$$E_{resid}^n = \frac{|\phi^\dagger(k + \lambda c + \lambda^2 M)\phi|}{\phi^\dagger K \phi} \quad (3.33)$$

Here ϕ and λ are the estimates of the eigenpairs computed using the modal approximation technique. The matrices, k , c and m are the fully assembled stiffness, coupling and mass matrices. This residual norm is a measure of the

relative accuracy of the eigenvalue solution. It is available in both the text results files and the output **Exodus** files, and should be consulted to determine the convergence.

3.24.4. Projection_eigen

The *Projection_Eigen* method is the most robust of all the solvers available for quadratic eigenvalue problems. Options of the *Projection_Eigen* solver are provided in Table [3-73](#). These parameters are identical to those for the *sa_eigen* method.

Option	Args	Description
nmodes	<i>int</i>	Number of requested eigen modes
shift	<i>Real</i>	Eigen shift used in computation of the subregion modes. See 3.10.
reorthogonalize	<i>string</i>	no: no reorthogonalization yes: reorthogonalize all modes full: check all modes
check_diag	<i>string</i>	no: no check for orthogonalization yes: check redundant modes full: check all modes
sort method	<i>string</i>	magnitude: complex magnitude of λ frequency: Sort by frequency and then damping. damping: Sort by damping and then frequency. truefreq: Sort by frequency... avoiding zero energy round off. none:

Table 3-73. – Projection_Eigen Options.

3.25. NLStatics Solution Case

Parameter	Type	Default	Description
tolerance	<i>Real</i>	1e-6	Controls completion of the Newton iteration.
max_newton_ iterations	<i>Integer</i>	100	If the iteration count exceeds this value before reaching tolerance , the Newton loop is considered to have failed.
update_ tangent	<i>Integer</i>	101	How often the tangent stiffness matrix is rebuilt during the Newton iterations.
num_newton_ load_steps	<i>Integer</i>	1	Number of load steps used to incrementally step up to the final equilibrium position

Table 3-74. – NLStatics Solution Case Parameters.

If stiffness matrix K is a function of the displacement u in

$$K(u) = f,$$

then use **NLstatics** for nonlinear statics. Newton's method applied to the residual force equations to drive the residual $r = p - f$ to zero. The residual vector r is the difference between the internal force vector p and the external force vector f . The internal force vector is a function of the structural displacements (and possibly velocities). External forces can also be a function of the structural displacements in the case of follower loads such as surface pressure loads.

The **tolerance** keyword provides control over the completion of the Newton iteration. Once the change in the L^2 -norm of the *displacement* decreases below **tolerance**, the loop completes successfully.

The **num_newton_load_steps** keyword controls the number of load steps used to incrementally step up to the final equilibrium position. Large loads may cause the Newton algorithm to diverge. If this occurs, increase the number of load steps applied. Displacements will be output after each load step which may be animated similar to transient dynamics simulations.

The **update_tangent** keyword controls how often the tangent stiffness matrix is rebuilt during the Newton iterations. The default is set to update the tangent stiffness matrix at the beginning of a load step. Setting **update_tangent** to 1 is equivalent to using a full-Newton algorithm where the tangent stiffness matrix is rebuilt after each Newton iteration. For nonlinear (difficult) problems, this option may be optimal, but for most problems the extra cost of assembling a preconditioning the tangent stiffness matrix should be amortized over several solves. Note that for this option to improve Newton's method, the element types in the model must have the tangent stiffness method implemented.

An example **Solution** section is shown below.

```
Solution
  title 'Example of a nonlinear statics solution'
  nlstatics
  tolerance           = 1e-6
  max_newton_iterations = 100
  num_newton_load_steps = 10 // split load into 10 increments
  update_tangent       = 1   // full-newton algorithm
end
```

3.26. *NLTransient Solution Case*

Parameter	Type	Default	Description
time_step	<i>Real</i>		Time step size. See Section 3.32.0.1
nsteps	<i>Integer</i>		Number of time steps to take. See Section 3.32.0.2
start_time	<i>Real</i>	0.0	Solution case start time. See Section 3.32.0.3
nskip	<i>Integer</i>	1	Results output frequency. See Section 3.32.0.4
rho	<i>Real</i>	1	Select time integrator. See Section 3.32.1.1
load	<i>Integer</i>		Load to apply during solution case. See Section 3.32.0.5
write_files	<i>all/ none/ output/ history</i>	all	Controls which result files are written during this solution.
tolerance	<i>Real</i>	1e-6	Controls completion of the Newton iteration.
max_newton__iterations	<i>Integer</i>	100	If the iteration count exceeds this value before reaching tolerance , the Newton loop is considered to have failed.
update__tangent	<i>Integer</i>	101	How often the tangent stiffness matrix is rebuilt during the Newton iterations.
flush	<i>integer</i>	50	Defines how often results are written to the exodus results file. See Section 2.4.1

Table 3-75. – NLTransient Solution Case Parameters.

The **NLtransient** solution method is used to perform a³⁴ direct implicit nonlinear transient analysis. A projector-corrector step is used. Note that for a linear system, the **NLtransient** analysis will require two solves per time step.

The **tolerance** keyword provides control over the completion of the Newton iteration. Once the change in the L^2 -norm of the *acceleration* decreases below **tolerance** , the loop completes successfully. Note the difference viz a viz **NLstatics** ([3.25](#)), where **tolerance** instead refers to the *displacement*.

If the iteration count in a given time step exceeds **max_newton_iterations** , the Newton loop is considered to have failed. Thus, note that **max_newton_iterations** is not the limit for the total number of Newton iterations, as it is in the **NLstatics** (3.25) case, but is instead the limit on the number of iterations per time step.

In a **NLstatics** (3.25) analysis, load stepping can be used to help the convergence of the Newton loop by cutting the total load into a series of incremental steps. This is controlled with the **num_newton_load_steps** keyword. However, in **NLtransient** analysis, load stepping makes no sense because the dynamic response of a structure subjected to a total load is different from the response to a series of incremental loads. In effect, the load stepping is replaced by time stepping in the case of nonlinear transient analysis. Thus, the keyword **num_newton_load_steps** is inactive for nonlinear transient analysis.

For **NLtransient** problems, if Newton’s method diverges, either the tangent stiffness matrix has to be updated more often (see **update_tangent**) or the time-step should be decreased.

The **update_tangent** controls how often the dynamic tangent stiffness matrix is rebuilt during the Newton iterations. The default is set to 101, and thus unless a given Newton loop takes more than 101 iterations, the tangent matrix will not be updated by default. Setting **update_tangent** to 1 is equivalent to using a full-Newton algorithm where the dynamic tangent stiffness matrix is rebuilt after each Newton iteration. Note that currently there is no option for forcing a tangent update at the beginning of each time step, unless the **update_tangent** keyword is set to exactly the number of Newton iterations taken per time step. For non-linear problems, some control of this option is recommended. Note, for this option to improve Newton’s method, the element types in the model must have the dynamic tangent stiffness method implemented.

3.27. *ReceiveSierraData Solution Case*

Parameter	Type	Default	Description
include__internal_force	<i>off/on</i>	on	Include or disable computation of initial internal force
no_geom_stiff			Turn off the geometric stiffness term

Table 3-76. – ReceiveSierraData Solution Case Parameters.

Solution case **receive_sierra_data** inputs a deformed or preloaded model state. Calculations in SIERRA codes such as **Sierra/SM** may be transferred to **Sierra/SD**. This provides the ability to compute nonlinear responses in a separate code, followed by a supported solution case in **Sierra/SD**. For SM/SD **hand-off** also see⁵⁴ chapter

“**Sierra/SM** to **Sierra/SD** Coupling” and.⁵³ In particular see the “**Sierra/SM** output parameters for **Sierra/SD** modal analysis using file transfer.” section.

Solution method `receive_sierra_data` helps with the following.

1. Update the initial geometry from the previously computed displacements.
2. Update the element stiffness matrices due to preload stresses. This is the geometric stiffness correction. The geometric stiffness correction is supported for volumetric and membrane elements at this time. Stress transfer for shells and beams is not enabled. The geometric stiffness calculation may be disabled with the `no_geom_stiff` keyword.
3. Compute and apply an initial force associated with the input stress state. Disable this by setting `include_internal_force=off` command.
4. Update of tangent stiffness of materials based the preloaded material state.

`receive_sierra_data` solutions require a multicase solution. An example follows. In this example preload data is received from the input **Exodus** file, that preload alters the stiffness matrix, and eigenvalues are computed using this updated stiffness matrix.

```
Solution
  case xfer
    receive_sierra_data
  case eig
    eigen
    nmodes=40
    shift=-3e6
end
```

The `receive_sierra_data` solution is specifically designed to read initial conditions and tangent stiffness state computed in a **Sierra/SM** analysis. Note that connections between different parts of the model specified in the **Sierra/SM** input deck are not transferred, e.g., contact conditions, MPCs, and joints. These conditions may be equivalently specified in **Sierra/SD**. The following options are available to `receive_sierra_data`.

include_internal_force When set to the default value (`on`), the imported stress state is integrated to generate an internal force body load. This load is included in the right-hand side of subsequent static or transient analysis. If the **Sierra/SM** preload analysis is in static equilibrium, the externally applied forces from boundary conditions will be in balance with the internal forces generated by the elements. By default, the internal forces will be computed again in **Sierra/SD**. Thus, if the same forces are included in the **Sierra/SD** input deck that causes the preload in **Sierra/SM**, then the internal force should be included to keep the model in static equilibrium. An alternative is to exclude the boundary conditions causing the preload from **Sierra/SD** and set `include_internal_force=off`. In this case, **Sierra/SD** will compute no initial internal force, and the initial state of the model will be

treated as perfect equilibrium. See the verification manual chapter “**Sierra/SM to Sierra/SD Coupling**” for more detail on this topic.

no_geom_stiff Can be used to ignore the preloaded stress contribution to the geometric stiffness matrix. Generally, elements in tension have higher effective stiffness, and elements in compression have lower effective stiffness. At sufficiently large compression, element stiffness can go negative, which causes severe problems for solver stability. Turning off geometric stiffness can be used for debugging purposes, or to evaluate what effect the geometric stiffness term has on model response. See the verification manual chapter “**Sierra/SM to Sierra/SD Coupling**” for more detail on this topic.

Data Transfer The purpose of `receive_sierra_data` is transfer of data from a previous **Sierra/SM** analysis to **Sierra/SD**. Data relevant to the load is mostly transferred automatically based on expected naming conventions. For the Lamé hyperfoam model (section 4.1.4), and orthotropic layers in shells (section 4.1.2 and table 3-79), the option `from_transfer` must be used to turn on the material tangent update for specific material parameters. The fields relevant to `receive_sierra_data` are given in tables 3-77 to 3-79.

Exodus Label	Description	Effect
displ_x displ_y displ_z rotational_displ_x rotational_displ_y rotational_displ_z	Nodal Displacement and Rotation	Added to mesh coordinates and used as initial mesh configuration. See below for details
veloc_x veloc_y veloc_z rotational_veloc_x rotational_veloc_y rotational_veloc_z	Nodal Velocity	Used to set initial conditions for transient analysis
force_internal_x force_internal_y force_internal_z	Force from stress state	Used for equilibrium diagnostics

Table 3-77. – Nodal data used in `receive_sierra_data`.

Only accurately labeled data will be transferred from the **Exodus** file. An exact match is required. The only flexibility is in the underscore position. Either `stress_xx` or `stressxx` are allowed, for example. The data, typically written from **Sierra/SM**, may require special output requests (in the **Sierra/SM** input file) for proper naming. See the How To or Verification manual for examples.

Exodus Label	Description	Effect
stress_x stress_y stress_z stress_xy stress_yz stress_zx	volumetric stress components	Used for geometric stiffness by default. Can turn off with <code>no_geom_stiff</code> option. Also, used to compute internal force by default which can be turned off with <code>include_internal_force=off</code>
left_stretch_xx left_stretch_yy left_stretch_zz left_stretch_xy left_stretch_yz left_stretch_zx	Element Stretch	Used for material tangent stiffness calculations for LAME material models
rotation_xx rotation_yy rotation_zz rotation_xy rotation_yx rotation_yz rotation_zy rotation_zx rotation_xz	Element Rotation	Used for material tangent stiffness calculations for LAME material models
lame_state_neo_hookean	Material State	Used for material tangent stiffness for Neo-Hookean
lame_state_hyperfoam	Material State	Used for material tangent stiffness for hyperfoam

Table 3-78. – Element data used in `receive_sierra_data`. Only the Neo-Hookean and hyperfoam Lamé models are supported; see Section 4.1.4.

Exodus Label	Description	Effect
memb_str_xx memb_str_yy memb_str_zz memb_str_xy memb_str_yz memb_str_zx	fiber membrane stress components	Geometric stiffness and internal force
bottom_stress_xx bottom_stress_yy bottom_stress_zz	fiber membrane stress components	Geometric stiffness and internal force
fibermodulus	fiber modulus	Tangent stiffness with from_transfer
fiberdir	primary fiber direction	Tangent stiffness with from_transfer
fiberdir2	secondary fiber direction	Tangent stiffness with from_transfer
fiberdens	fiber density	Mass with from_transfer
fiberthickness	thickness	Tangent stiffness with from_transfer
E1	material modulus	Tangent stiffness with from_transfer
E2	material modulus	Tangent stiffness with from_transfer
nu12	Poisson ratio	Tangent stiffness with from_transfer
G12	material modulus	Tangent stiffness with from_transfer
density	material density	Mass with from_transfer

Table 3-79. – Data Transferred in `receive_sierra_data` specific to `orthotropic_layer` Material. Some of this data is transferred automatically. Other data requires the **from_transfer** keyword in the material definition.

Alternatively, **Sierra/SD** also supports user-defined label mappings using `initialize variable name`, `read variable`, and `variable type` options in the `FILE` section. An example for displacements is shown below, where nodal displacements are stored as “dx”, “dy”, and “dz” on the input geometry file `inputMesh.g`.

```
FILE
  geometry_file = inputMesh.g

  initialize variable name = displacement(x) # x-component of displacement
  variable type = node                      # is stored in a nodal field
  read variable = dx                        # on inputMesh.g named "dx"

  initialize variable name = displacement(y) # y-component of displacement
  variable type = node                      # is stored in a nodal field
  read variable = dy                        # on inputMesh.g named "dy"

  initialize variable name = displacement(z) # z-component of displacement
  variable type = node                      # is stored in a nodal field
  read variable = dz                        # on inputMesh.g named "dz"
END
```

A shorthand notation of the user-defined label mapping is also supported and shown below for the same example of displacement. For vector fields, ‘x’, ‘y’ and ‘z’ is appended to the variable name, and for tensor fields, “xx”, “yy”, “zz”, “xy”, “yz” and “zx” is appended. For tensor fields, the order of off-diagonal indices is irrelevant (“xy” is the same as “yx”).

```
FILE
  geometry_file = inputMesh.g

  initialize variable name = displacement # displacement is stored
  variable type = node                  # in nodal fields on inputMesh.g
  read variable = d                     # named "dx", "dy" and "dz"
END
```

A complete list of valid label keys for (`initialize variable name`) is included in table 3-80.

Coordinate Update The `receive_sierra_data` method uses **Sierra/SM** displacement to update the initial coordinates of the mesh. Mass matrices are computed in the undeformed configuration, to better ensure conservation of mass in this **hand-off**. However, the initial coordinates for **Sierra/SD** are updated with the transferred displacements.

$$\vec{x} = \vec{x}_o + \vec{u}$$

displacement(x)	velocity(x)	rotational_velocity(x)
displacement(y)	velocity(y)	rotational_velocity(y)
displacement(z)	velocity(z)	rotational_velocity(z)
force_internal(x)	acceleration(x)	rotational_acceleration(x)
force_internal(y)	acceleration(y)	rotational_acceleration(y)
force_internal(z)	acceleration(z)	rotational_acceleration(z)
stress(xx)	memb_stress(xx)	bottom_stress(xx)
stress(yy)	memb_stress(yy)	bottom_stress(yy)
stress(zz)	memb_stress(zz)	
stress(xy)	memb_stress(xy)	bottom_stress(xy)
stress(yz)	memb_stress(yz)	
stress(zx)	memb_stress(zx)	
fibermodule	fiberdir	fiberdir2
fiberdens	fiberthickness	matrixmodulus
matrixpoisson	cordpoisson	epl
compendmodulus	Q1zterm	Q2zterm
corddiameter	cordaxialforce	cordaxialstretch
time	timestep	kinetic

Table 3-80. – List of valid label keys for `receive_sierra_data`

Here x_o is the location of the undeformed coordinates. Stiffness matrices, forces, and tied MPCs are then computed in the updated frame. Tied MPCs can be defined via tied surfaces (Section 8.1) or contact (Section 8.2).

3.28. Statics Solution Case

Parameter	Type	Default	Description
load	<i>Integer</i>		Load to apply during solution case. See Section 3.32.0.5

Table 3-81. – Statics Solution Case Parameters.

The **statics** keyword is required if a static solution is needed, i.e. the solution to the system of equations $[K]\{u\} = \{f\}$. An example **Solution** section is shown below.

```
Solution
  title "Example of a statics solution"
  statics
end
```

3.29. Superposition Solution Case

Parameter	Type	Default	Description
reduced_file	<i>String</i>		Ignored?

Table 3-82. – Superposition Solution Case Parameters.



Superposition is currently BETA release.
Enable with the “**--beta**” command-line option.

The **superposition** provides superelement recovery capability. This recovers physical space solutions from generalized degrees of freedom.

A CB model generates a transformation matrix consisting of a combined set of fixed interface and constraint modes. See section 3.5. This modal basis may be stored in an exodus file. A netcdf file containing the reduced order model is also created at this time. Subsequently, this reduced model is inserted into a residual model for superelement analysis, say a transient analysis. That analysis outputs the standard exodus results, and may also generate output on the netcdf file. These data may be post-process using linear superposition to determine output quantities on the original interior degrees of freedom of the superelement. This is illustrated in Figure 3-22.




Figure 3-22. – Superposition Data Flow Diagram.

The **superposition** method requires a single argument, the file name of the reduced order model containing the output results. In addition, the **geometry_file** specified in the FILE section must contain the modal basis for the Craig-Bampton reduction. See the example in input 3.15.

```

solution
  case sp
    superposition
  
```

```

        reduced_file=rom-out.ncf
end

file
    geometry_file basis-out.exo
end

```

Input 3.15. Superposition Example. Output will be to “basis-out-sp.exo”

Limitations The superposition method is under development and contains the following limitations.

- Solutions are supported in serial at this time.
- Only displacement, velocity and acceleration may be output.
- Superposition is tested for Eigen and Transient solutions. There is no support for frequency domain solutions at this time.
- Data recovery is possible for a single superelement in each run.

3.30. *Tangent Solution Case*

Parameter	Type	Default	Description

Table 3-83. – Tangent Solution Case Parameters.

The **tangent** solution step requires a multcase solution (see paragraph 3.2). It forces an update of the tangent stiffness matrix. It is typically used following a nonlinear solution step to ensure that the following step begins using the tangent stiffness matrices computed from the previous result. However, it may also be used following a linear solution step, in which case the stiffness matrix is computed again based on the current value of displacement.

During tangent the stiffness is computed again based on element configuration in the deformed state. Tangent also adds any stress stiffening effects from the preceding load case. Stress resultants from the preload step are not added to stress resultants of subsequent cases after the tangent solution, but do contribute to increasing the effective stiffness computed in the tangent stiffness. For example, the stress stiffening that occurs from tension of a guitar string would be successfully captured using the tangent solution

case. As another example the stiffness of a warped plate may be significantly different from the stiffness of a flat plate, tangent will take into account this change in stiffness.

The tangent stiffness matrix is assembled at the subdomain level from computations at the element level. It represents the partial derivative of the force with respect to the displacement, i.e.

$$K_{tangent} = \frac{\partial f}{\partial u} \quad (3.34)$$

In eigen analysis, the tangent stiffness matrix replaces the linear stiffness matrix in the eigenvalue equation. This permits computation of modal response following a preload.

3.31. *TranShock Solution Case*

Parameter	Type	Default	Description
time_step	<i>Real</i>		Time step size. See Section 3.32.0.1
nsteps	<i>Integer</i>		Number of time steps to take. See Section 3.32.0.2
start_time	<i>Real</i>	0.0	Solution case start time. See Section 3.32.0.3
nskip	<i>Integer</i>	1	Results output frequency. See Section 3.32.0.4
rho	<i>Real</i>	1	Select time integrator. See Section 3.32.1.1
load	<i>Integer</i>		Load to apply during solution case. See Section 3.32.0.5
write_files	<i>all/ none/ output/ history</i>	all	Controls which result files are written during this solution.
SRS_damp	<i>Real</i>	0.03	Damping coefficient used for the shock response spectra calculation

Table 3-84. – TranShock Solution Case Parameters.

The **transhock** solution method is used to perform a direct implicit transient analysis followed by computation of the shock response spectra for the degrees of freedom in a specified node set. The options for configuring a **transient** solution case described in Section [3.32](#) are all applicable in the **transhock** solution case too. An example of a **transhock** solution case and **frequency** block are shown in input [3.16](#). Damping for the

implicit transient portion of the simulation is defined in Section 4.8. The `srs_damp` is not used in the modal transient portion of the simulation.

A `frequency` block must also be included in the input file for **transhock** solution cases to define the frequencies and nodesets for computing the shock response spectra (see Section 7.4). The frequencies for computing the shock response spectra are defined by setting `freq_step`, `freq_min`, and `freq_max`. They are identified in the `frequency` section along with an application region (see Section 7.4). The range of the computed frequency spectra is controlled by `freq_min` and `freq_max`, while `freq_step` controls the resolution. The accuracy of the computed spectra is independent of the magnitude of `freq_step`. This parameter controls the output resolution. Examples are presented in inputs 3.9, 3.10 and 3.16.

The shock spectrum procedure will compute acceleration results. The options specified in the `outputs` and `echo` blocks are used in the transient portion of the analysis, but are ignored for the post processing of the transient results into shock spectra. Thus, if displacement, velocity, and/or acceleration is selected in the `outputs` and/or `echo` sections for a shock spectra analysis, the results echoed to the output listing or the Exodus output file will be time history results as requested, but the shock spectra results will be for acceleration response for the nodes in the specified node set. *The calculated shock spectra are written to the frequency file (*.frq); they are not output to the Exodus results file.*

```
Solution
  transhock
    time_step .00005
    nsteps 500
    nskip 1
    srs_damp .03
end

Frequency
  freq_min 100.
  freq_max 10000.
  freq_step 100.
  nodeset 3
  acceleration
end
```

Input 3.16. Transhock Example Input

3.32. Transient Solution Case

Parameter	Type	Default	Description
time_step	<i>Real</i>		Time step size. See Section 3.32.0.1
nsteps	<i>Integer</i>		Number of time steps to take. See Section 3.32.0.2
start_time	<i>Real</i>	0.0	Solution case start time. See Section 3.32.0.3
nskip	<i>Integer</i>	1	Results output frequency. See Section 3.32.0.4
rho	<i>Real</i>	1	Select time integrator. See Section 3.32.1.1
load	<i>Integer</i>		Load to apply during solution case. See Section 3.32.0.5
write_files	<i>all none output history</i>	all	Controls which result files are written during this solution.
PredictorCorrector	<i>Integer</i>	-1	predictor-corrector implementation
ConstraintCorrectionFrequency	<i>Integer</i>	1	time-step frequency for constraint correction
ConstraintErrorDiagnostics	<i>yes/no</i>	no	prints constraint errors
flush	<i>integer</i>	50	Defines how often results are written to the exodus results file. See Section 2.4.1
nUpdateConstraints	<i>Integer</i>	0	Frequency to update constraints

Table 3-85. – Transient Solution Case Parameters.

The **transient** solution method is used to perform a direct implicit transient analysis.

3.32.0.1. Time_step Option

The **time_step** defines the time step size. The time step size is constant for a time period. Multiple time periods can be given as shown in example [3.32.0.6](#).

3.32.0.2. Nsteps Option

The `nsteps` defines the number of time steps to take. As shown in example [3.32.0.6](#), multiple time periods can be defined with different numbers of steps.

3.32.0.3. Start_time Option

The `start_time` sets the initial time of a transient analysis. By default, a transient case begins where a previous transient solution case ended, or the time transferred from a `receive_sierra_data` or other preload solution case. Otherwise, the default start time is 0.0. To set the start time in the input deck specify `start_time`.

3.32.0.4. Nskip Option

The `nskip` controls how many integration steps to take between outputting results. It defaults to 1, which is equivalent to outputting all time steps. Because transient analysis often takes little computational time per step, the overall runtime can be significantly reduced by choosing not to output the results at every step, i.e., setting `nskip` greater than 1. The `nskip` in the solution section can be overridden for history (Section [7.3](#)) and linesample (Section [7.5](#)) output by specifying `nskip` in those sections.

3.32.0.5. Load Option

The `load` defines the name of the “load” block to apply during the solution case. See Section [6.3.1](#) for details.

3.32.0.6. Defining Multiple time Periods

We note that multiple time step values, along with the corresponding number of steps, can be specified for transient analysis. This can be useful for separating the simulation into a section of tiny time steps followed by a section of larger time steps, or vice versa. The following provides an example of the use of multiple time steps.

```
solution
  time_step    1.0e-5    1.0e-3
  nsteps       100      500
  nskip        10       1
end
```

In this case, the user requested 100 time steps of $\Delta t = 10^{-5}$, followed by 500 time steps with $\Delta t = 10^{-3}$. There is no practical limit on the number of such regions that may be specified.

3.32.0.7. PredictorCorrector Option

PredictorCorrector indicates whether predictor-corrector implementation is to be used (see theory manual). It defaults to 1. If it is 1, predictor-corrector is always used, while 0 indicates that it is never used. **ConstraintCorrectionFrequency** and **ConstraintErrorDiagnostics** are related to the constraint errors arising in the predictor-corrector implementation, with the first controlling how often the correction is applied, while the second can be used to examine the evolution of constraint errors of displacement, velocity and acceleration.

3.32.1. nUpdateConstraints Option



nUpdateConstraints is currently BETA release.

Enable with the “--beta” command-line option.

nUpdateConstraints defines the frequency to update constraints (e.g. MPCs and contact). When it is undefined or set equal to 0, the constraints will never be updated. A value of 1 results in the constraints being updated every time-step. In addition to re-building the constraints, the nodal coordinates will also be updated at the corresponding time-steps based on the interpolation of any existing displacements on the mesh (if applicable).

The name of the displacement nodal variables can be defined by the **initialize variable name**, **read variable**, and **variable type** options in the FILE section, as shown in section 3.27.

This option is related to the **nodesets_with_disp** option (section 2.3), which enables specifying displacements on a subset of all nodes via nodeset output. However, nodeset displacements do not currently support user-defined variable names as with nodal displacements.

3.32.1.1. Numerical damping

Two time integrator schemes are available for direct time integration. The method and the configuration of the integrator are selected using the keyword **rho**. If this keyword is not found, the time integrator defaults to a standard Newmark beta^{22,34} integration scheme. With **rho** the Generalized Alpha method²⁷²¹ is used, and the value of the numerical damping is controlled by **rho**.

Important

Due to the inexactness of linear solvers, the Newmark beta integrator is conditionally unstable. Without damping, a solution may gradually diverge. Also, see subsection Linear transient analysis section Solution Procedures of the Theory Manual. Either proportional damping or numerical damping is strongly recommended in all cases.

The option **rho** defines the numerical damping of the Generalized Alpha method. **rho** varies from 0 (maximal damping case) to 1 (minimal damping case). **If rho is not specified in the input file, the integrator defaults to the Newmark beta method.** Otherwise, the code uses the value of **rho** given by the user to configure the Generalized Alpha method. Therefore, there is no value default for **rho**, as shown in the table above, since if it is not specified the code uses the Newmark beta method instead. If **rho** is specified to be greater than 1 or negative an error message is printed. **rho** determines **Newmark beta**, α_f , and α_m of the Generalized Alpha method. More detailed information on the implementation, and references can be found in the description of the method in the **Sierra/SD** Theory Manual.

The following conditions suffice to achieve second order accuracy and unconditional stability:

$$\begin{aligned}\alpha_m &< \alpha_f \leq \frac{1}{2} \\ \gamma_n &= \frac{1}{2} - \alpha_m + \alpha_f \\ \beta_n &\geq \frac{1}{4} + \frac{1}{2}(\alpha_f - \alpha_m)\end{aligned}$$

The parameters are determined to satisfy the conditions. Specifically,

$$\begin{aligned}\alpha_f &= \rho / (1 + \rho) \\ \alpha_m &= (2\rho - 1) / (1 + \rho) \\ \beta_n &= (1 - \alpha_m + \alpha_f) \cdot (1 - \alpha_m + \alpha_f) / 4 \\ \gamma_n &= 1/2 - \alpha_m + \alpha_f\end{aligned}$$

We note some special cases of interest. If $\rho = 0$, we have that $\alpha_f = 0$ and $\alpha_m = -1$. This is the maximum damping case. If $\rho = 1$, we have that $\alpha_f = \alpha_m = \frac{1}{2}$, which yields $\beta_n = \frac{1}{4}$, and $\gamma_n = \frac{1}{2}$. This is similar to the classical undamped Newmark beta method, although we note that it is a different algorithm since $\alpha_f = \alpha_m = \frac{1}{2}$ implies some lagging in the time-stepping procedure. The classical undamped Newmark beta method has $\alpha_f = \alpha_m = 0$.

Unlike proportional damping, there is no direct relation between **rho** and an equivalent modal damping term. A value of **rho=0.9** is recommended for most analyses. The Generalized Alpha integrator imparts numerical damping to the solution that most strongly affects *high* frequency content. Users must check that the damping in the frequency range of interest is physical. For example with a time step size of $1e-5$, damping has the most effect at frequencies above the Nyquist frequency $.5e+5$.

3.32.1.2. Initial Acceleration

Determining the initial acceleration is necessary³⁵ for quadratic convergence. If force, velocity or displacement is specified, an initial solve may be required to determine a consistent acceleration. Remember that determination of the displacement at step $n+1$

depends on values at the previous step. Specifically, the acceleration at the previous step is given by the solution to the equation,

$$A_n = M^{-1}(F_n + Kd_n + Cv_n)$$

where M , K , and C are the corresponding mass, stiffness and damping matrices.

If this mass solve is not performed, it is possible to introduce a spike in acceleration which can oscillate through time. Initialization can be a somewhat tricky process. An example set of use cases is provided in Appendix 6.5. By default, GDSW is used. Some combinations of MPCs can lead to a singular mass matrix that will cause solver errors. For these cases the initial mass solve is deactivated with the following command in the **parameters** block (see Section 2.3),

```
Parameters
  DoInitialMassSolve = false
end
```

3.33. Transient-Inverse Solution Case

Parameter	Type	Default	Description
time_step	<i>Real</i>		Time step size. See Section 3.32.0.1
nsteps	<i>Integer</i>		Number of time steps to take. See Section 3.32.0.2
start_time	<i>Real</i>	0.0	Solution case start time. See Section 3.32.0.3
nskip	<i>Integer</i>	1	Results output frequency. See Section 3.32.0.4
rho	<i>Real</i>	1	Select time integrator. See Section 3.32.1.1
load	<i>Integer</i>		Load to apply during solution case. See Section 3.32.0.5
write_files	<i>all none output history</i>	all	Controls which result files are written during this solution.

Table 3-86. – Transient-Inverse Solution Case Parameters.

The **transient-inverse** solution method is used to solve in inverse problem (see Sec. 3.39) for a time domain analysis. With a few exceptions, the parameters for the forward **transient** solution case discussed in Sec. 3.32 apply to this solution method as

well. The user provides a time series of displacements and/or pressures at a set of nodes in the model, and the solution to the inverse problem is a set of loads, materials, etc. that best correspond with the user's input.

The basic requirements for a **transient-inverse** simulation are as follows:

Optimization: Control over the optimization problem is specified in the **optimization** block. See Sec. 3.39.1 for further details.

Inverse-Problem: The **inverse-problem** block provides the connection to the measurement data. It is also where **design_variable** is specified (e.g., load, material, etc.). See Sec. 3.39 for details.

Truth Table: The **data_truth_table** from the **inverse-problem** block is a list of the indices of the global node numbers (a.k.a. target nodes) where displacements or acoustic pressures are measured. See Sec. 3.39.3 for file format details.

Data File: Experimentally determined “target” displacements are read from **data_file** specified in the **inverse-problem** block. See Sec. 3.39.3 for file format details.

Sierra/SD uses the Rapid Optimization Library (ROL) for solving optimization problems. During the optimization solution ROL writes an output file, *ROL_Messages.txt* that contains convergence information. Section 3.39.8 contains a discussion of the output file that is written by ROL.

3.33.1. Load Identification

```
solution
  transient-inverse
  nsteps = 100
  time_step = 1e-3
  nskip = 1
end
inverse-problem
  design_variable = load
  data_truth_table = ttable.txt
  data_file = dataReal.txt
end
optimization
% optimization_package = ROL_lib
ROLmethod = linesearch
LSstep = Newton-Krylov
LS_curvature_condition = null
Max_iter_Krylov = 20
Use_FD_hessvec = false
Use_inexact_hessvec = false
opt_tolerance = 1e-12
```

```

    opt_iterations = 2
end
loads
  sideset 301
    inverse_load_type = spatially_constant
    pressure=10
    function = 1
end

```

Input 3.17. Transient Load Identification Example

Specifying `design_variable = load` applies inverse methods to determine sideset loads which best correspond with the measured displacements and/or acoustic pressures provided by the user. The material and model parameters do not change during the solution. For structures, the loads are pressures or tractions¹, and for acoustics, the loads are acoustic accelerations. Note that for structures, inversion is based on the signed magnitudes of the tractions; the direction of each traction is fixed.

An example input deck is given in input 3.17. In addition to the input blocks discussed in the beginning of this section, there is another that is specific to `design_variable = load`:

Loads: See Sec. 3.39.6 for a description of the inverse parameters in the `loads` block for load identification problems.

Section 3.39.7 contains a discussion of the current limitations with inverse load methods.

3.33.2. Material Identification

Specifying `design_variable = material` with the `transient-inverse` method applies inverse methods to determine material parameters when provided with both loads and structural displacements and/or acoustic pressures in a given finite element model². The load parameters do not change during the solution, which provides the material parameters for elements in the model that are specified to have unknown materials.

In addition to the input blocks discussed in the beginning of this section, there are several others specific to `design_variable = material`:

Block: See Sec. 3.39.4 for a description of the `block` specifications for material inverse problems.

Material: See Sec. 3.39.5 for a description of the provides `material` specifications for material inverse problems.

¹Moments and point forces are not currently supported.

²As the system matrices (and consequently the modes) change at every inverse iteration, `design_variable` cannot be set to `material` for `modaltransient-inverse` problems.

3.34. *TSRpreload Solution Case*

Parameter	Type	Default	Description
linedata_only	<i>True/False</i>	False	Indicates that no system matrices should be computed, but the linedata specified in the linesample file should be computed for verification of data transfer.

Table 3-87. – TSRpreload Solution Case Parameters.

The **tsr_preload** solution method reads an **Exodus** file with a previously computed Thermal Structural Response (TSR) into **Sierra/SD** for a subsequent statics or transient dynamics analysis. This is not a fully coupled calculation. Stress results are read from the geometry file, an equivalent internal force is computed, and that internal force is combined with the applied force throughout the transient run. If temperature data is also included in the file, it will be read and used to compute temperature dependent material properties. A **tsr_preload** requires a multcase solution, and it must be followed by a transient dynamics or statics solution (see paragraphs 3.2 and 3.32 respectively).

Note that since the stresses are converted into a force, and since there is no immediate deformation in transient dynamics, the elastic stresses output by **Sierra/SD** will be small initially, i.e. they will not contain a contribution from the thermal stress. However, at large times, the deformation from the internal force will result in an elastic stress opposite to that of the thermal stress. The **linesample** method 7.5 recovers the input thermal stress as an output quantity (in either MATLAB or **Exodus** format).

The **tsr_preload** solution method is considered to be a temporary solution to a more complicated problem. In the future, TSR analysis will involve coupling to other mechanics codes. In many cases a thermal load (Section 6.3.8), may provide equivalent capability.

Data in the **Exodus** file from which TSR data is read must strictly match the following criteria. There must be one time step in the result. That time step must have a number of different element fields defined. These correspond to the six stress values of the stress tensor and the number of stress tensors defined per element must correspond exactly to the number of integration points for that element. For instance, a hex20 element requires exactly 27 stress tensor values per element. An error is produced if the number of data points read in does not match the number of integration points for that element. For instance, a fully integrated hex8 will produce an error if reading in Gauss point data produced by a fully integrated hex20 element. Shell and beam type elements are not supported in **tsr_preload**.

The labels for the stresses must be as shown in the table below where **SigXX** is interchangeable with **SigMA_XX** or **STRESS_XX**. Both sequential and ijk⁷¹ numbering are

supported for integration point data. For sequential numbering, replace %d with an integer representing the integration point value (0 to 26). Do not zero pad. For ijk numbering,⁷¹ labeling goes as SigXX_Hex20_GP%d where %d is replaced with the ijk numbering scheme and _Hex20 is replaced with the element type. Do not include _%d for labeling centroid stress data.

Name	Definition
SigXX_%d	σ_{xx} , the xx component of stress
SigYY_%d	σ_{yy} , the yy component of stress
SigZZ_%d	σ_{zz} , the zz component of stress
SigYZ_%d	σ_{yz} , the yz component of stress
SigXZ_%d	σ_{xz} , the xz component of stress
SigXY_%d	σ_{xy} , the xy component of stress

The following is an example solution section for a TSR preload followed by transient dynamics.

```
Solution
  title 'Pure bending from initial stress'
  case TSR
    tsr_preload
    load 1
  case bend
    transient
    time_step 1.e-6
    start_time 1.0e-3
    nsteps 3
    nskip 1
    load 2
end
```

If executed on a file with `geometry_file = example.exo`, this will produce two output files, `example-tsrx.exe` and `example-bend.exe`. The first of these has little useful information. The second contains the displacements (or other variables) from the transient analysis.

3.34.0.1. Line Sample

One additional feature for thermal structural response is the ability to do line sampling [7.5](#) on the original **Exodus** file containing the element stresses. This is useful for debugging and verification. It allows the stresses along lines within the structure to be examined. Sampling occurs for data stored on integration points using the variables names described above. Line sample is used for energy deposition (see the *Two Element Exponential Decay Variation Hex20* problem⁵⁴). Energy deposition is interchangeable with supplying an applied temperature.

In `tsr_preload`, the input **Exodus** file is required to contain at least one of the following fields: stress, temperature or energy deposition. Any field that is not found in the input **Exodus** file is reported as a zero field in the output line sample output file.

3.35. *Residual Vectors Solution Case*

Parameter	Type	Default	Description
<code>node_list_file</code>	<i>File</i>		Unknown

Table 3-88. – Residual Vectors Solution Case Parameters.



Residual Vectors is currently BETA release.
Enable with the “`--beta`” command-line option.

The **residual_vectors** solution method Modal truncation augmentation (MTA)²⁴ provides a method to represent the modes not retained in an eigen solution. It is particularly useful in component mode synthesis approaches where multiple models are joined together. In NASTRAN, MTA vectors are referred to as ‘residual vectors’. The theory of MTA²⁴ is established. We use the following terminology:

N	Number of degrees of freedom
nev	Number of retained eigenvalues/eigenvectors from the eigen solution
nf	Number of applied forces and/or moments
M	Mass matrix of size NxN
K	Stiffness matrix of size NxN
Φ	Matrix with eigenvectors as columns, size Nx (nev)
Ω²	Diagonal matrix of eigenvalues, size (nev)x(nev)
R₀	Applied spatial load vector, size Nx(nf)
R_s	Modally represented spatial load vector, size Nx(nf)
R_t	Force truncation vector, size Nx(nf)
X	Static displacements due to applied loads, size Nx(nf)
ω̄²	Diagonal matrix of reduced eigenvalues, size (nf)x(nf)
Q̄	Matrix of reduced eigenvectors, size (nf)x(nf)
P	Matrix of modal truncation (residual) vectors, size Nx(nf)

The algorithm for computing MTA vectors is:

1. Solve the generalized eigenvalue problem

$$\mathbf{K}\Phi = \mathbf{M}\Phi\Omega^2$$

for **nev** eigenvalues and eigenvectors. This is done by first specifying **eigen** in a multcase solution procedure.

2. Compute the force truncation vector

$$\mathbf{R}_t = \mathbf{R}_0 - \mathbf{R}_s = \mathbf{R}_0 - \mathbf{M}\Phi\Phi^T\mathbf{R}_0.$$

3. Compute the static displacements $\mathbf{X}$ due to the force truncation vector $\mathbf{R}_t$ by solving $\mathbf{K}\mathbf{X} = \mathbf{R}_t$.
4. If rigid body modes are present, orthogonalize $\mathbf{X}$ to them. The optional input `nrbms` allows the user to specify the number of rigid body modes present.
5. Form the reduced matrices $nf \times nf$,

$$\bar{\mathbf{K}} = \mathbf{X}^T\mathbf{K}\mathbf{X}, \quad \bar{\mathbf{M}} = \mathbf{X}^T\mathbf{M}\mathbf{X}.$$

6. Solve the reduced generalized eigenvalue problem $\bar{\mathbf{K}}\bar{\mathbf{Q}} = \bar{\mathbf{M}}\bar{\mathbf{Q}}\bar{\omega}^2$
7. Form the modal truncation (residual) vectors: $\mathbf{P} = \mathbf{X}\bar{\mathbf{Q}}$
8. Construct the pseudo modal set: $\tilde{\Phi} = [\Phi|\mathbf{P}]$.

In Sierra-SD, the multi-case solution strategy is:

1. Solve the eigenvalue problem and specify the keyword `resvecs`
2. For each column of $\mathbf{R}_0$, solve a statics problem
3. Solve a `residual_vectors` problem to form the pseudo modal set

An example `solution` block is given here for a case with six rigid body modes.

```
Solution
title 'Sample MTA solution procedure'
case 'eigen'
  eigen
    nmodes=30
    shift=-1000.    // needed for floating
    solver=gdsw
case 'ResVecs'
  residual_vectors
    node_list_file node1
end
```

3.36. *GeometricRigidBodyModes Solution Case*

Parameter	Type	Default	Description

Table 3-89. – GeometricRigidBodyModes Solution Case Parameters.

Nominal rigid body modes may be determined from the coordinates. No attempt is made to account for boundary conditions. This solution method requires that the GDSW linear solver be used.

The intent of the examples is to first introduce rigid modes, next use the modes to solve an eigenvalue problem, and then demonstrate the Modal Transient capability 3.22. The third example uses the modes in a modal transient simulation to deflate out the rotations. This section depends on Section 6.3.19.

The geometry rigid body mode capability will always generate six modes and these modes are explicitly ordered +X, +Y, +Z, +RX, +RY, +RZ.

Rigid body modes are requested in the SOLUTION block.

```
SOLUTION
    geometric_rigid_body_modes
END
PARAMETERS
    num_rigid_mode 6
END
```

Rigid body modes can be incorporated into the modes computed in a modal analysis, and then used for other purposes.

```
SOLUTION
    case out
        geometric_rigid_body_modes
    case flexibleModes
        eigen
        nmodes 10
        shift -1e6
    END
PARAMETERS
    num_rigid_mode 6
END
```

Rigid body modes are the 6 lowest frequency eigenvectors. In this case 4 more modes are computed, for a total of 10.

In this example a modal transient simulation uses the geometric rigid body modes to deflate out the (infinitesimal) rotation, while retaining the translational rigid body modes. This is equivalent to use of the `FilterRbmLoad` for direct transient solutions (though accomplished differently).

SOLUTION

```
case out
  geometric_rigid_body_modes
case vibration
  eigen
  nmodes 10
case filter
  modalfiltercase
  modalfilter rotation
case transient
  modaltransient
  time_step 1.e-5
  nsteps 62
  load 42
```

END

PARAMETERS

```
num_rigid_mode 6
```

END

```
modalfilter rotation
```

```
add all
```

```
remove 4:6
```

END

3.37. Waterline Solution Case

Parameter	Type	Default	Description
max_iterations	<i>Integer</i>	100	Maximum number of solution iterations
tolerance_force	<i>Real</i>	1.0e-6	Target force balance accuracy
point_a	<i>Real(3)</i>		Coordinates of point on estimated water surface
point_b	<i>Real(3)</i>		Coordinates of point on estimated water surface
point_c	<i>Real(3)</i>		Coordinates of point on estimated water surface
VizOption	<i>none/Enight</i>	none	Whether Enight writes a file for visualizing the waterline plane

Table 3-90. – Waterline Solution Case Parameters.

It can be advantageous to determine the waterline of a ship prior to commencing more complex analysis. The **waterline** capability solves the nonlinear geometric equations of equilibrium for a rigid ship in water. An example is shown in input 3.18.

```

Solution
  case 'waterline'
    waterline
      max_iterations 100
      tolerance_force 1.0e-6 // absolute tolerance on force convergence
      point_a 0 0 0 // coords of point 'A' on estimated water surface
      point_b 1 0 0 // coords of point 'B' on estimated water surface
      point_c 1 1 0 // coords of point 'C' on estimated water surface
      load 1

  case 'transient'
    ...
end

LOAD 1
  sideset 1 // wetted sideset
  pressure = 1
  function = 1 // this defines rho g h

```

```

    body
      gravity = 0 0 9.8
    end

// this assumes rho=1000, g=9.8
Function 1
type Linear
data 0.0 0.0
data 1.0e6 9.8e9
end

```

Input 3.18. Waterline Example

The arguments `point_a`, `point_b` and `point_c` indicate the Cartesian coordinates of three points A , B , C on the estimated water surface. These three points define a plane, which serves as the initial guess of the waterline. The waterline normal is determined using the right-hand rule with these points, as shown in Figure 3-23. The Newton's method implementation then uses this plane as the initial guess, and begins iterations towards force and moment equilibrium. On completion, we write out the coordinates of three points on the final (converged) waterline surface, along with the Cartesian coordinate system defined by these points. This output appears in the result file in text format. A **grepos** script for moving the body may also be written.

The optimization is configured as follows.

max_iterations sets the maximum number of iterations.

tolerance_force is a normalized force residual. The norm is computed from the residual vector,

$$F_{residual} = [F_z/W, M_{\theta_1}/(LW), M_{\theta_2}/(LW)]$$

where $W = Mg$ is the total weight of the ship, and L is a characteristic length of the model.

VizOption may be `none` or `Ensign` to generate a visualization script.

In addition to the entries in the **Solution** section of the input, this method requires two **load** entries and a **function**. The **load** entries define the sideset for the wetted surface and the gravity load.³ The **function** defines the pressure as a function of depth. In the example of input 3.18, the argument to the function is the depth, h . The **function** returns $P = \rho gh$.

³Gravity is specified using the standard load keywords of a body load with a gravity vector. However, for the **waterline** solution, the magnitude of the gravity vector is relevant. The gravity direction is *always* directed opposite the normal to the surface for this solution type.




Figure 3-23. – Waterline Coordinate Definition. The plane of the surface is defined by three points: A,B, and C. The θ_1 rotation is about the line from A to B, while the normal is defined using the right-hand rule.

The waterline iteration may output nodal data during the iteration. Select **force** to output the buoyancy force. Select **npressure** to output the nodal pressure. See the Outputs section, 7, for details.

3.37.0.1. Limitations There are a number of limitations to this method.

gradient-based optimization: These powerful algorithms are based on nonlinear gradient-based optimization and have subtle limitations. Limitations are listed below.

1. Singular tangent matrices are generated in various conditions, which cause the solution to terminate. A common condition causing a singular tangent matrix is a body completely submerged in a constant-density fluid. For simplicity, consider an unrotated cube⁴ of edge length S . The net force on the cube is,

$$F_{\text{net}} = (A_{\text{bottom}}P_{\text{bottom}} - A_{\text{top}}P_{\text{top}}) - mg \quad (3.35)$$

$$= \rho_f g S^2 (h_{\text{bottom}} - h_{\text{top}}) - mg \quad (3.36)$$

$$= \rho_f g S^3 - \rho_s g S^3 \quad (3.37)$$

⁴The arithmetic is easier for a cube, but the arguments can be shown to be completely valid for any rigid body.

where ρ_f and ρ_s are the densities of the fluid and solid respectively. Significantly, the net force does not depend on the average depth. Thus,

$$K_t = \frac{\partial F_{net}}{\partial z} = 0.$$

$K_t = 0$ for a ship that is completely out of the water too.

Real seawater is not constant density. An optimal solution may be found in this case. However, because the pressure is usually expressed as a piecewise linear function, the same problem occurs. Use of a runtime function may allow computation of higher-order derivatives, but this has not been evaluated.

Figure 3-24 plots net force versus depth for a body. Only the partially submerged region has a nonzero tangent matrix that can be determined by a gradient-based optimization scheme.




Figure 3-24. – Net Force vs depth for a Rigid Body. Only the unshaded region, where the body is partially submerged, has a non-singular tangent matrix.

2. Gradient-based solution methods often have trouble with local minima. These can occur in the case of unstable systems, such as a light, tall cylinder floating on a dense fluid. A local minimum occurs for the cylinder standing vertically. A global minimum is achieved when the cylinder is perturbed and falls to the side.
3. Gradients may also go to zero for symmetry reasons. A perfect cylinder floating on the water has no sensitivity to roll.

sideset orientation: The wetted surface defines the pressure surface. It does not need to be closed. However, there can be no contribution to the net force from portions of the model that are submerged, but not part of the sideset.

One and only one sideset defines the wetted surface. Its outward direction should point into the water. There is no check for a reversal of the normal vectors on the sideset. This must be evaluated by the analyst.

Z-orientation Current design requires that the initial configuration has gravity approximately aligned with the global *Z* coordinate.

3.38. *Gap Removal Solution Case*

Parameter	Type	Default	Description
ignore_gap_inversion	<i>true/false</i>	false	suppress fatal error, gap output behavior if true

Table 3-91. – gapremoval Solution Case Parameters.

If two meshes are tied using either **tied data** or **contact definition**, then along the interface opposite elements may initially overlap or leave gaps. Gap removal, which is done by default, attempts to remove these gaps and overlaps. Gap removal is the same as **initial overlap** removal.

The `gap_removal` solution case is used to debug contact setup prior to submitting a full run. The `gap_removal` solution case runs quickly and uses low memory. This solution method enables visualization of the constraints created and the gap removed from **tied data** 8.1 and **contact definition** 8.2 blocks. The `gap_removal` solution method reads in the mesh, applies the contact search and gap removal algorithm and writes out the output mesh with gap removed. An example input is given below.

```
solution
  gap_removal
end

tied data
  surface 1,2
  name "tied_1-2"
  search tolerance 1.0e-3
end
```

Removal of contact gaps is essential to maintaining rigid body invariance. This is illustrated in Appendix 8.3. However, the removal of gaps in tied surfaces can occasionally result in distorted elements which may make it difficult or impossible for the solver to converge. Thus, it can be advantageous to investigate the results of gap removal before committing to a full and expensive solution case.

Gap removal output will include two element variables. The variable `elementInversionFlag` is set to one on any inverted element. The variable `elementQuality` gives a non-dimensional element quality metric (one is ideal higher is worse) for the deformed configuration of every element. The same quality metric is used as described in Section 7.1.12.

If gap removal inverts any element the filename extension “-gap” will be appended to the output exodus file and a fatal error will be given. One of the parameters described in Section 2.3 influences the Gap Removal solution case. The parameter is `ignore_gap_inversion`. It is false by default. Set it to true to suppress both the fatal error and the “-gap” output behavior.

Gap removal for lofted surfaces is discussed in Section 8.3.

In addition to the element shape, information diagnostics regarding the contact constraints are available. The `rslt` file will list basic information on the number of constraints found. Detailed visual information on constraint locations is obtained by requesting `constraint_info` in the `outputs` block as described in Section 7.1.48. Furthermore, adding MPC to the `echo` block prints every contact created MPC as described in Section 7.7.2.

Coupled Electro-Mechanical Analysis Piezoelectricity is the production of electrical charges on a surface by the imposition of mechanical stress. **Sierra/SD** supports coupled electro-mechanical physics in order to model piezoelectric materials subjected to electrical and mechanical forces. This support includes static, transient, eigen, and direct frequency response solution methods, piezoelectric and dielectric material models (4.5, 4.6), and voltage measurement based inverse methods such as material and source identification. Electrical boundary conditions such as charge-based Neumann (6.3.12) and voltage-based Dirichlet(6.1.2, 6.1.3) boundary conditions are also supported.⁵²

3.39. Inverse Problems

Inverse problems optimize parameters to reproduce experimental results. Inverse methods include transient and direct frequency response load identification, direct frequency response material identification, and material identification from eigenvalues. The methods are based on solving optimization problems, with the goal to minimize the norm of the difference between measured and predicted data. More detail is provided in the references found in the theory notes. Inverse methods for identifying an unknown material (3.11, 3.8.2) or an unknown load (3.8.1, 3.33.1) require solution block input. There may be additional input required, such as the specification of test data results 3.39.3 and the specification of the parameters in the Optimization 3.39.1 and Inverse Problem 3.39.2 sections. Input 3.19 illustrates a *partial* input for a “directfrf-materialid” problem. Highlighted portions of the input are outlined below.

Sierra/SD uses the Rapid Optimization Library (ROL) as an optimization engine. Portions of the ROL documentation can be found on the Trilinos website.⁵

```
solution
  directfrf-inverse
end
optimization
  optimization_package = ROL_lib
  ROLmethod = linesearch
  LSstep = secant
  opt_tolerance = 1e-10
end
inverse-problem
  design_variable = material
  data_truth_table = ttable.txt
  real_data_file = data.txt
  imaginary_data_file = data_im.txt
end
block 1
  inverse_material_type=homogeneous
  material 1
end
block 2
  inverse_material_type=known
  material 2
end
material 1
  isotropic
  density 10
  G 1
  K 1
end
material 2
  isotropic
  density 1
  G 2
  K 2
end
```

Input 3.19. Sample “directfrf-inverse” input for material identification.

⁵  <https://trilinos.org/packages/rol>

Portions of the input that are specific to inverse methods are emphasized.

3.39.1. Optimization

The **optimization** section provides options to control the optimization strategy as part of an inverse method such as material identification. Parameters for the optimization section are listed in Table 3-92, and an example is shown in input 3.21.

Sierra/SD uses the Rapid Optimization Library (ROL) [40], which is a Trilinos [33] package for large-scale optimization. ROL is particularly well suited for the solution of optimal design, optimal control and inverse problems in large-scale engineering applications. The currently supported methods are trust region and line search, and the corresponding parameters for these methods are listed in Tables 3-92, 3-93 and 3-94. These are only a subset of the parameters available in ROL. We note that in tables 3-93 and 3-94, the abbreviations *tr* and *ls* stand for *trust region* and *line search*, respectively. We use these abbreviations to keep the parameter names succinct.

We note that for *source inversion* problems that involve **directfrf-inverse** , **transient-inverse** or **modaltransient-inverse** solution cases, we recommend using Krylov-based methods such as *line search* with the *newton-krylov* option, or *trust region* with *truncatedcg* option. An example for that case is given in 3.20.

```
optimization
  ROLmethod = linesearch
  LSstep = newton-krylov
  Max_iter_Krylov = 50
  opt_tolerance = 1e-8
  opt_iterations = 5
  scaleDesignVars = yes | no
END
```

Input 3.20. Optimization Section Example for Source Inversion

In the case of *material inversion* problems, the best algorithm is problem-dependent, and may require some experimentation to arrive at the optimal parameters. We note that for material inverse problems, the parameter **scaleDesignVars** has been shown to significantly help convergence. This parameter defaults to **no**, but can be set to **yes** to facilitate convergence.

Sierra/SD currently uses ROL methods for unconstrained and bound-constrained optimization with line searches and trust regions (ROLmethod, boundConstraints, Table 3-92). To provide the context for the parameter tables, we review some notation and provide references to optimization textbooks.

Suppose $\mathcal{X}$ is a Hilbert space of functions mapping Ξ to $\mathbb{R}$. For example, $\Xi \subset \mathbb{R}^n$ and $\mathcal{X} = L^2(\Xi)$ or $\Xi = \{1, \dots, n\}$ and $\mathcal{X} = \mathbb{R}^n$. We assume that the objective function $f : \mathcal{X} \rightarrow \mathbb{R}$ is twice continuously Fréchet differentiable and that the bound constraints $a, b \in \mathcal{X}$ are given with $a \leq b$ almost everywhere in Ξ . We focus on methods for solving unconstrained and bound-constrained optimization problems,

$$\underset{x}{\text{minimize}} \quad f(x) \quad \text{and} \quad \underset{x}{\text{minimize}} \quad f(x) \quad \text{subject to} \quad a \leq x \leq b,$$

respectively. The methods implemented in ROL utilize derivative information and two strategies for guaranteed (global) convergence from remote starting points, line searches and trust regions. We use the notation $\nabla f(x)$ to denote the gradient of f at x and $\nabla^2 f(x)$ to denote the Hessian of f at x .

Line-search methods. Let x_k be the k -th optimization iterate. For unconstrained problems, line search methods compute an update to x_k in the form of

$$x_{k+1} = x_k + \alpha_k s_k,$$

where s_k is a descent vector, and $\alpha_k > 0$ is a scalar. The vector s_k can be computed using a variety of methods, including steepest descent, nonlinear conjugate gradients, quasi-Newton (secant) methods and Newton-Krylov methods, see [45, Ch. 3, 5, 8 and 6, resp.] and [23, Ch. 6.3]. Table 3-93 lists the parameters corresponding to the choice of the method to compute the descent vector s_k (LSstep, Table 3-93). To compute the scalar α_k , a line search approximately minimizes the one-dimensional function $\phi_k(\alpha) := f(x_k + \alpha s_k)$, i.e., it approximately solves the optimization problem

$$\underset{\alpha}{\text{minimize}} \quad \phi_k(\alpha) := f(x_k + \alpha s_k).$$

In general, the approximate minimizer α_k must satisfy sufficient decrease and curvature conditions to guarantee global convergence [45, Ch. 3]. **Sierra/SD** uses the *cubic interpolation* line search from ROL, which includes a backtracking procedure that satisfies the Armijo sufficient decrease condition

$$\phi_k(\alpha_k) \leq \phi_k(0) + c_1 \alpha_k \phi'_k(0) \quad \Longleftrightarrow \quad f(x_k + \alpha_k s_k) \leq f(x_k) + c_1 \alpha_k \langle \nabla f(x_k), s_k \rangle_{\mathcal{X}},$$

where $0 < c_1 < 1$, and does not require a curvature condition. An initial guess for the line-search parameter can be specified if steepest descent or nonlinear conjugate gradient methods are used for the computation of the descent vector s_k (initial_LS_Par, Table 3-93). For bound constrained problems, the line search is a *projected* search. That is, the line search approximately minimizes the one-dimensional objective function

$$\phi_k(\alpha) = f(P_{[a,b]}(x_k + \alpha s_k)),$$

where $P_{[a,b]}$ denotes the projection onto the upper and lower bounds. Such line-search algorithms result in projected gradient, projected quasi-Newton and projected Newton algorithms (see below).

Trust-region methods. For unconstrained problems, given the k -th iterate x_k trust-region methods compute the trial step s_k by approximately solving the trust-region subproblem

$$\underset{s}{\text{minimize}} \quad \frac{1}{2} \langle B_k s, s \rangle_{\mathcal{X}} + \langle g_k, s \rangle_{\mathcal{X}} \quad \text{subject to} \quad \|s\|_{\mathcal{X}} \leq \Delta_k,$$

where $B_k \in L(\mathcal{X}, \mathcal{X})$ is an approximation of $\nabla^2 f(x_k)$, g_k approximates $\nabla f(x_k)$, and $\Delta_k > 0$ is the trust-region radius. The approximate minimizer s_k must satisfy the fraction of Cauchy decrease condition

$$-\frac{1}{2} \langle B_k s, s \rangle_{\mathcal{X}} - \langle g_k, s \rangle_{\mathcal{X}} \geq \kappa_0 \|g_k\|_{\mathcal{X}} \min \left\{ \Delta_k, \frac{\|g_k\|_{\mathcal{X}}}{1 + \|B_k\|_{L(\mathcal{X}, \mathcal{X})}} \right\}$$

for some $\kappa_0 > 0$ independent of k . ROL implements several trust-region methods, including Cauchy point, dogleg, double dogleg, and truncated conjugate gradient methods, see [45, Ch. 4] and [23, Ch. 6.4]. Table 3-94 lists the parameters corresponding to the choice of the method (TRstep, Table 3-94). Additionally, the user can specify the initial trust-region radius, Δ_0 (initial_TR_Radius, Table 3-94). For bound constrained problems, ROL employs projected gradient, projected secant, and projected Newton-type methods. These methods prune variables based on the binding set (see below) and run standard trust-region subproblem solvers on the remaining variables. To ensure sufficient decrease, ROL then performs a modified projected line search.

Bound Constraints. The bound constraint methods in ROL require the active set of an iterate x_k ,

$$\mathcal{A}_k = \{ \xi \in \Xi : x_k(\xi) = a(\xi) \} \cap \{ \xi \in \Xi : x_k(\xi) = b(\xi) \}.$$

The active set is the subset of Ξ corresponding to points in which x_k is equal to the upper or lower bound. The complement of the active set (called the inactive set), $\mathcal{I}_k = \mathcal{A}_k^c = \Xi \setminus \mathcal{A}_k$, is the subset of Ξ corresponding to points in which x_k is strictly between a and b . Given $\mathcal{A}_k$ and the gradient $g_k = \nabla J(x_k)$, we define the binding set as

$$\mathcal{B}_k = \{ \xi \in \Xi : x_k(\xi) = a(\xi), -g_k(\xi) < 0 \} \cap \{ \xi \in \Xi : x_k(\xi) = b(\xi), -g_k(\xi) > 0 \}.$$

The binding set contains the values of $\xi \in \Xi$ such that if $x_k(\xi)$ is equal to either the upper and lower bound, then $(x_k - g_k)(\xi)$ will violate bound. For both projected line-search and trust-region methods, the step is computed by fixing the variables in the active set and only optimizing over the inactive variables. That is,

$$s_k|_{\mathcal{A}_k} = 0 \quad \text{and} \quad s_k|_{\mathcal{I}_k} \text{ are free.}$$

Considering active, inactive and binding variables results in poor overall performance. To circumvent this behavior, ROL employs ϵ variations of this set. Namely, if $\epsilon > 0$, then

$$\begin{aligned} \mathcal{A}_k^\epsilon &= \{ \xi \in \Xi : x_k(\xi) \leq a(\xi) + \epsilon \} \cap \{ \xi \in \Xi : x_k(\xi) \geq b(\xi) - \epsilon \} \\ \mathcal{B}_k^\epsilon &= \{ \xi \in \Xi : x_k(\xi) \leq a(\xi) + \epsilon, -g_k(\xi) < 0 \} \cap \{ \xi \in \Xi : x_k(\xi) \geq b(\xi) - \epsilon, -g_k(\xi) > 0 \}. \end{aligned}$$

The ϵ inactive set is similarly defined as the complement of the ϵ active set. ROL dynamically controls ϵ so that as an algorithm approaches the optimal solution, ϵ decreases to zero. [38, 13, 42, 18, 19].

Krylov methods. Both line-search and trust-region methods may involve iterative methods of the Krylov type in the step computation. If such methods are requested, the stopping conditions for these sub-solvers can be defined through the parameters Absolute_Krylov_tol, Relative_Krylov_tol and Max_iter_Krylov described in Table 3-92. Scenarios requiring Krylov methods are triggered, for instance, if TruncatedCG is selected for a trust-region algorithm or if newtonkrylov is selected for a line-search algorithm.

Parameter	type	default	Description
ROLmethod	select	required	LineSearch, TrustRegion
boundConstraints	Yes/No	Yes	Bound constraints on design variables
scaleDesignVars	Yes/No	no	Controls scaling of design variables by initial guess to make them nondimensional
opt_tolerance	real	1e-10	Gradient tolerance
opt_iterations	int	20	Maximum iterations; an iteration may include one solve of the Newton system
objective_tolerance	real	1e-20	Objective function tolerance
Absolute_Krylov_tol	real	1e-6	Krylov absolute tolerance
Relative_Krylov_tol	real	1e-3	Krylov relative tolerance
Max_iter_Krylov	int	20	Krylov iteration limit

Table 3-92. – Optimization Section Parameters

Parameter	type	default	Description
LSstep	select	secant	nonlinearcg, steepest, secant, newtonkrylov, newton (require useTransferMatrix=on 3-95)
initial_LS_Par	real	1.0	only applicable for nonlinearcg and steepest.

Table 3-93. – Optimization Section Parameters for Line Search

Parameter	type	default	Description
TRstep	select	CauchyPoint	TruncatedCG, CauchyPoint, Dogleg, DoubleDogleg
initial_TR_Radius	real	-1	initial radius; if -1, ROL computes an initial value

Table 3-94. – Optimization Section Parameters for Trust Region

```

optimization
  ROLmethod = linesearch
  LSstep = nonlinearcg
  Max_iter_Krylov = 50
  opt_tolerance = 1e-8
  opt_iterations = 5
END

```

Input 3.21. Optimization Section Example

3.39.2. Inverse-Problem

The **inverse-problem** block of the input deck connects externally defined data describing the test. The format for the files is described in the Inverse Data Files section (3.39.3). This section also controls parameters of the optimization (such as regularization). Options for the **inverse-problem** section are included in Table 3-95.

Inverse Problem Parameters	Load Id		Material Id	
	Transient	FRF	FRF	Eigen
<i>General Parameters</i>				
design_variable	R	R	R	R
data_file	R	NA	NA	NA
data_truth_table	R	R	R	NA
data_weight_table	NA	NA	NA	-
real_data_file	NA	R	R	NA
imaginary_data_file	NA	R	R	NA
modal_data_file	NA	NA	NA	R
modal_weight_table	NA	NA	NA	R
data_type	-	-	-	-
useTransferMatrix	NA	-	-	NA
<i>Regularization Parameters</i>				
gradientSurfaceRegParameter	-	-	NA	NA
tikhonovParameter	-	-	-	-
gradientTikhonovRegularizationParameter	NA	NA	-	-
MECE_penalty	NA	NA	-	NA
<i>Multi-Experiment Parameters</i>				
nresponses	NA	NA	NA	NA
loadID	NA	NA	-	NA

Table 3-95. – Inverse-Problem parameters. Items marked with “-” are optional. Items marked ‘R’ are required, and items marked NA are not applicable.

3.39.2.1. Regularization Parameters

gradientSurfaceRegParameter a penalty term, penalizing jumps in load input parameters between neighboring patches. Increasing the regularization parameter decreases checker-boarding on multi-patch outputs.

tikhonovParameter a penalty term penalizing large values of the design variables. A larger term forces smaller design variables.

MECE_penalty a penalty term, penalizes the size of the misfit error between the measured and predicted data. Unlike least squares methods, MECE methods do not strive to exactly reproduce the measured data. Increasing the penalty term decreases the misfit error.

3.39.2.2. Multi-Experiment Parameters

nresponses the number of experiments in a multi-experiment inverse problem. If not specified by the user, this parameter defaults to 1. This parameter must match the number of data files listed for 'real_data_file' and 'imaginary_data_file'.

loadID list of load identifiers, each corresponding to a separate experiment in a multi-experiment inverse problem. The length of this list must match the value specified for 'nresponses'.

3.39.2.3. Transfer Matrix Option The transfer matrix is a linear map between the design space and the state space for linear inverse problems such as source inversion. When the **useTransferMatrix** is set to true, the transfer matrix is internally computed and stored, and explicitly used to compute the gradient and the full hessian matrix. This option may significantly decrease computational time at the expense of in core memory for select optimization problems with some or all of the following characteristics: 1. small number of design variables, 2. linear inverse problems with uncertainty, and 3. PSD source inversion problems (future release capability). In addition to increased computational performance, this option will allow the user to solve linear inverse problems using the full newton method (see Table 3-92).

Currently, the useTransferMatrix is only applicable to solution type **inverse-directfrf** problems, and must be flagged with **beta** when executing **Sierra/SD**.

3.39.3. Inverse Data Files

The interface for measured data involves several data files. The specific files needed depends on the solution method, as summarized in Table 3-95. This section describes the data files and data file formats. Input 3.22 provides an example of a frequency domain inverse problem.

```

inverse-problem
  design_variable = load | material
  data_truth_table = truthTable.txt
  real_data_file = dataReal.txt
  imaginary_data_file = dataImag.txt
  data_type = disp | accel | moduli | voltage
end

```

Input 3.22. The inverse problem section for direct frequency response sets the names of the files that contain user specified data.

In the case of a structural-only problem, the `data_type` can be set to either `disp` or `accel`, indicating that the incoming data is in the form of displacements or accelerations, respectively. In addition to displacement and acceleration data, the `data_type` can also be set to `moduli` for viscoelastic material identification problems. For acoustics-only problems, `data_type` is not applicable as the incoming data always corresponds to acoustic pressure. For models incorporating piezoelectric materials, setting `data_type` to `voltage` indicates that the experimental data is in the form of voltages. The parameter `data_type` defaults to `disp` if not specified.

Data Truth Table

The `data_truth_table` file contains the **global** node numbers (a.k.a. target nodes) where the experimental data measurements are given. The first line in the file contains the number of points where measurements are given, and the remaining lines contain the global node numbers where the experimental data is specified. If `data_type` is set to `moduli`, the `data_truth_table` file contains the node numbers where the displacements are computed.

Figure 3-25 provides a simple example of the truth table format for an acoustics-only problem, Figure 3-26 provides an example for a structural-only problem, and Figure 3-27 provides an example for a structural acoustics problem. Note that for acoustics-only, the file consists of the list of nodes where the acoustic measurements were taken. For voltage based problem, the file looks identical to an acoustic-only problem, where the listed nodes represent the locations where voltage measurements were taken. For structural-only problems, each line in the truth table contains 4 columns that indicate the node numbers where measurements were taken, and then 3 columns of binary (0/1) input that indicate which dofs (i.e. x , y , and z) are active in the optimization. For structural acoustics problems, Fig. 3-27 shows that each line contains 5 columns, with the first column again indicating the node number, and the remaining four columns contain either 1 or 0, which allows to turn on/off the x , y , z , and *pressure* degrees of freedom in the optimization.

The data truth table identifies the location of measurement data. For example, if the experimental data is collected at three microphones, which correspond to nodes 10, 120, and 3004, then the `data_truth_table` file is as follows.

```

3
10
120
3004

```

Thus there are a total of 4 lines in the file, even though the first line specifies three nodes for the measurement data. The global node numbers correspond to the global ID in the exodus input to **Sierra/SD**. The order of the nodes in this list corresponds to the order of the data in the corresponding real or imaginary data files. Other than that, there is no restriction on the ordering of the node numbers (i.e. they do not need to be in ascending or descending order).

Figure 3-25. – Sample Data Truth Table Input for Acoustic Problem

```

4
25 1 1 1
96 1 1 1
13 1 1 1
17 1 1 1

```

Figure 3-26. – Example Data Truth Table for Structures. There are 4 nodes, with numbers 25, 96, 13, 17. The second through fourth columns of 1, 1, 1 mean that all structural degrees of freedom are active, the fifth column is acoustic and implied to be zero.

```

4
5 0 1 0 0
6 0 1 0 0
7 0 0 0 1
8 0 0 0 1

```

Figure 3-27. – Example data truth table for Structural Acoustics. There are 4 nodes, with numbers 5 to 8. Nodes 5 and 6 are structural where only the y component of displacement is measured. Nodes 7 and 8 are acoustic, where only the pressure is measured.

Real Data File

For inverse problems in the frequency domain, the `real_data_file` contains the real component of the measurement data at each frequency, corresponding to the nodes that are specified in the `data_truth_table` file. For a multi-experiment inverse problem, several files must be specified.

An example of the `real_data_file` for an acoustics-only problem is shown in Figure 3-28. The first line of the file contains the number of nodes where measurement data is provided, followed by the number of frequencies of data. Starting on the second line, the real part of the data at the first node is given for all frequencies. In particular, starting on the second line the data corresponds to the first node in the truth table list, not the node with the lowest ID number. Similarly, subsequent lines contain the real part of the data, at all frequencies, for the remaining nodes. Note that, since this is an acoustics-only example, only one line of data is needed for each node in the truth table. The format for the data files describing voltages exactly matches the format of the acoustics-only data file.

We build on the small example given in Figure 3-25 that has measurements at nodes 10, 120, and 3004, and consider the case where there are 2 frequencies in the data set. The `real_data_file` file for an acoustics-only problem could look as follows

3	2
1.1	2.4
0.7	3.3
2.1	1.4

The actual values in the above table were chosen arbitrarily, but observe that there is one “header row” followed by 3 data rows. There are 2 columns, corresponding to the two frequencies of the measured data. The units of the measurements must correspond to appropriate units in the analysis. One row of data is required for each DOF in the truth table. Data is required even if the value in the truth table is zero. If the truth table value for a DOF is zero, the data is not used in the analysis.

Figure 3-28. – Sample Real Data File Input for an acoustics-only problem

An example of the `real_data_file` for a structural-only problem is shown in Figure 3-29. Starting on the second line, the real part of the data at the first node is given for all frequencies. In particular, the first node is the node first in the truth table list, not the node with the lowest ID number. The x displacements for all nodes in the truth table is listed first, followed by the y displacements for all nodes, followed by the z displacements for all nodes.

For viscoelastic material identification problems in the frequency domain with `data_type` set to `moduli`, the `real_data_file` contains the real part of the shear and bulk modulus. An example of the `real_data_file` is shown in Fig. 3-30. The first line of the file is always

We build on the small example given in Figure 3-26 that has measurements at nodes 25, 96, 13, and 17, and consider the case where there is only 1 frequency in the data set. The `real_data_file` file for a structural-only problem could look as follows

```

12      1
1.1 // x-displacement for node 25
0.7 // x-displacement for node 96
2.1 // x-displacement for node 13
1.1 // x-displacement for node 17
0.7 // y-displacement for node 25
2.1 // y-displacement for node 96
1.1 // y-displacement for node 13
0.7 // y-displacement for node 17
2.1 // z-displacement for node 25
1.1 // z-displacement for node 96
0.7 // z-displacement for node 13
2.1 // z-displacement for node 17

```

The actual values in the above table were chosen arbitrarily, but observe that there is one “header row” followed by 12 data rows. There is only one column, corresponding to the single frequency of the measured data. The units of the measurements must correspond to appropriate units in the analysis. For structures these are units of displacement. Rows 2 – 5 correspond to the x components of the displacements at the nodes from the truth table, rows 6 – 9 correspond to the y components of displacement, and rows 10 – 13 correspond to the z components. Note that the x components of displacements for all nodes are listed first, followed by the y components for all nodes, followed by the z components for all nodes. One row of data is required for each DOF in the truth table. Data is required even if the value in the truth table is zero. If the truth table value for a DOF is zero, the data is not used in the analysis.

Figure 3-29. – Sample Real Data File Input for a structural-only problem

2, followed by the number of frequencies of data. Currently, this feature only supports one frequency of data, hence the first row should read 2 1. The second and third lines are respectively the shear modulus and bulk modulus. There is only one column, corresponding to a single frequency of measured data.

The `real_data_file` for a viscoelastic material identification problem with `data_type = moduli` could look as follows

```

2          1
5.7 // Shear modulus
3.1 // Bulk modulus

```

The actual values in the above table were chosen arbitrarily, but observe that there is one “header row” followed by 2 data rows.

Figure 3-30. – Sample Real Data File Input for a data type moduli problem

The frequencies of the measured data are specified in the **frequency** section. The frequencies given by **frequency** section must correspond to the frequencies where the experimental data was measured. These frequencies can be either uniformly or non-uniformly spaced, as specified in the **frequency** section.

Imaginary Data File

The `imaginary_data_file` has the exact same format as the `real_data_file` except that it contains the imaginary part of the data rather than the real part.

PSD Data File

In the context of PSD inversion, instead of specifying the real and imaginary data files, a single `psd_data_file` is needed, which contains the complex (Hermitian) PSD matrices for all the frequencies. The first two numbers in the `psd_data_file` must be the number of frequencies and the number of (measured) response degrees of freedom. The remainder of the file contains the PSD matrices, one matrix each for each frequency. The format of the PSD matrix should be in the natural row-oriented format, with each line containing a row of the matrix. The complex values should be entered in (*real part, imaginary part*) format, with space between the numbers. Sierra-SD automatically checks if each of the PSD matrices is Hermitian and positive definite, as each PSD matrix should be.

Data File

Transient time history data is stored in the `data_file`. The format is identical to the `real_data_file` except for the addition of the time step as the third column in the first

row. Each column of data now corresponds to a time step of analysis. An example is shown in Fig. 3-31. No interpolation is performed, and the measured data must exactly match the time steps of the analysis. As in the `real_data_file`, the rows of data are grouped as all of the x components of displacements at the measured nodes, followed by all of the y displacements, followed by the z displacements.

We note that for time-domain inverse problems, the data in the `data_file` must be padded with zeros at the beginning, since the sensor data is typically started at time $t = 0$. The maximum time-of-flight from the input loads to the sensors can be estimated easily (an upper bound is fine). Then, dividing the maximum time-of-flight by the time step gives the number of zeros to be added to the beginning of each time history. Without these zeros, the forward problem would not be able to come up with loads that match the sensor data in the early time response, due to the finite wave propagation speed.

6	4	0.1	
0.0	0.0	2.1	2.3
0.0	0.0	2.3	3.5
0.0	0.0	3.6	4.1
0.0	0.0	1.5	1.8
0.0	0.0	0.9	1.4
0.0	0.0	3.4	9.5

Figure 3-31. – Sample Transient Data File Input for a structural-only problem

Modal Data File

The `modal_data_file` contains measured modal results for inversion with the `eigen` solution. The first line of the file is the number of eigenvalues. Each subsequent line contains only the eigen frequency of the measured mode. It is important that the simulation modes are in the same order as the test modes.

Modal Weight Table

Not all computed eigenvalues may correspond to a test (or measured) mode. Further, even those modes identified may be more or less important. The `modal_weight_table` contains the weights applied to each computed mode.

The first line contains the number of weights, which must match the number of modes computed in the analysis, and each subsequent line contains the weight for the corresponding computed mode. All weights must be non-negative, and computed modes with no corresponding measured data should have zero weight.

Data Weight Table

In the case of eigenvector inversion (described in Sec. 3.11.2), the `data_weight_table` contains weights applied to each computed eigenvector error. This parameter allows for independent weighting of the eigenvalues and eigenvectors, but is not required; by default, it is the same as the `modal_weight_table`.

As above, the first line contains the number of weights, which must match the number of modes computed in the analysis, and each subsequent line contains the weight for the corresponding computed mode. All weights must be non-negative, and computed modes with no corresponding measured data should have zero weight.

3.39.4. Block section for Material Identification

For material inverse problems, the **block** section provides an additional option to control the optimization strategy. In particular, blocks can be specified as **known**, **homogeneous**, or **heterogeneous**, as shown in Table 3-96, and an example is shown in input 3.23.

The default behavior for the **inverse_material_type** keyword is **known**, which implies that the material parameters given in the corresponding material are fixed, and are not modified in the optimization process. In the case where all blocks are known, there is no need to solve the material inverse problem at all.

The remaining two options for the **inverse_material_type** keyword are:

- **homogeneous**. In this case, the material parameters are unknown and will be optimized during the inversion process, but are treated as constant over the entire block. This option is best used in the case when material properties are unknown in a block, but not expected to vary much over the block.
- **heterogeneous**. In this case, the material parameters are also unknown, will be optimized in the inversion process, but each element in the block will be given its own material properties. This is typically referred to as *spatially varying material properties*.

In input 3.23 Block 1 is **heterogeneous**, Block 2 is **homogeneous**, and Block 3 is **known**.

Parameter	type	Description
<code>inverse_material_type</code>	string	block inverse type

Table 3-96. – Block Section Parameters for Material Inversion

```
block 1
  material 1
  inverse_material_type heterogeneous
```

```

END
block 2
  material 2
  inverse_material_type homogeneous
END
block 3
  material 3
  inverse_material_type known
END

```

Input 3.23. Block Section Example for Material Inversion

3.39.5. Material section for Material Identification

For material inverse problems, the **material** section provides additional options to control the optimization strategy. Currently, for 3D elements, only **isotropic**, **isotropic_viscoelastic**, **orthotropic**, and **acoustic** material types are supported for material inverse problems.

We note that the parameters **G_bounds** and **K_bounds** only apply to elastic materials, whereas **Greal_bounds**, **Kreal_bounds**, **Gim_bounds**, and **Kim_bounds** only apply for viscoelastic materials. The parameters **Eij_bounds**, **Gij_bounds** and **Aij_bounds** apply only to orthotropic elastic materials, and the parameter **c0_bounds** applies only to acoustic materials.

For elastic, orthotropic elastic, and acoustic materials, the set of unknown material parameters may be customized by explicitly specifying the parameter **num_material_parameters** and a list of textttts following **material_parameters**. These textttts may include **bulk** and/or **shear** for elastic materials, **orthotropic** for orthotropic elastic materials, **sound_speed** for acoustic materials, and **rho** for all of the above.

For elastic materials, specification of **E_bounds** and **Nu_bounds** is permitted instead of specifying bounds for G and K . When specifying bounds for the elastic modulus and Poisson ratio, the material parameters must also be input as E and ν , and inversion must be requested over both elastic parameters.

For acoustic materials, the parameter **impedance_match** may also be specified. This has the effect of optimizing for sound speed c_0 and density ρ_0 under the condition where impedance $Z = \rho_0 c_0$ is constant.

Orthotropic Material Inversion: Unlike isotropic inversion, orthotropic inversion requires special care with respect to inadvertent material instabilities, i.e., during the inversion iterations, the elasticity tensor may not satisfy positive definiteness, potentially leading to the failure of even the forward solution. To avoid this, we parametrize the material tensor using the standard normal moduli E_{ii} , shear moduli G_{ij} , and special dimensionless parameters $A_{ij} = E_{ij} / \sqrt{E_{ii} E_{jj}}$. This is referred to as *Alpha* parametrization.

Parameter	type	Description
G_bounds	real real	lower and upper bounds on shear modulus
K_bounds	real real	lower and upper bounds on bulk modulus
E_bounds	real real	lower and upper bounds on Young's modulus
Nu_bounds	real real	lower and upper bounds on Poisson's ratio
Greal_bounds	real real	lower and upper bounds on real part of shear modulus
Kreal_bounds	real real	lower and upper bounds on real part of bulk modulus
Gimag_bounds	real real	lower and upper bounds on imag part of shear modulus
Kimag_bounds	real real	lower and upper bounds on imag part of bulk modulus
Eij_bounds	6 reals	lower and upper bounds on the three normal moduli
Gij_bounds	6 reals	lower and upper bounds on the three shear moduli
Aij_bounds	6 reals	lower and upper bounds on the three α parameters
density_bounds	real real	lower and upper bounds on density
c0_bounds	real real	lower and upper bounds on sound speed
impedance_match	real	value of $\rho_0 c_0$ to match for acoustic materials
num_material_parameters	int	(optional) number of material parameters
material_parameters	list of strings	(optional) bulk , shear , rho , orthotropic , and sound_speed

Table 3-97. – Material Section Parameters for Material Inversion

Such parametrization lends to easier imposition of material stability constraint which is done through a single inequality constraint and several bound constraints. Further details will be provided in an upcoming SAND report. Notwithstanding the theoretical details, the user is asked to pay special attention to the definition of A_{ij} , when applying bound constraints on these parameters. Bound constraints can be specified as shown in Table 3-97, and an example is shown in input 3.26. These parameters constrain the optimizer to work in a restricted space of possible design variable values, and prevent convergence to physically unrealistic values of the parameters. An alternative to Alpha parametrization is *Cholesky* parametrization, where the modulus matrix is parametrized through Cholesky factorization, avoiding the need of inequality constraints. This can be utilized by setting the flag `alphaparametrization` to `no`.

Transverse Isotropic Material Inversion is performed essentially through orthotropic inversion, by setting the flag `transverselyisotropic` to `yes`, followed by the plane of isotropy represented by a two-digit number, i.e., 12, 23, or 13, where 1,2,3 represent the three axes of material symmetry consistent with the coordinate system within the element block. Input 3.26 contains an example for transversely isotropic inversion.

Additional material parameters that can be optimized include `joint2g` elements and block proportional stiffness damping `blkbeta`. Since these parameters live in the **block** section rather than in the **material** section, their bound constraints are specified in the former. However, we specify their input syntax here since they are material parameters. An example of input syntax for `joint2g` and `blkbeta` inputs for material optimization are given in Table 3-98 and input 3.24.

The `joint2g` elements involve 6 parameters that can be optimized. The

`joint_truth_table` specifies which of these parameters will be included in the optimization solution. For example, in input 3.24, only the first (i.e. *x-component*) of the `joint2g` parameters will be optimized. The remaining parameters will stay fixed in the optimization, including those that are specified as *NULL*. In the case of `blkbeta`, the optimal *blkbeta* for the specified block will be computed and written to the result file.

Parameter	type	Description
<code>joint_elastic_bounds</code>	real real	bounds on joint2g spring stiffnesses
<code>joint_damper_bounds</code>	real real	bounds on joint2g dashpot (damper) parameters
<code>blkbeta_bounds</code>	real real	bounds on block stiffness proportional damping

Table 3-98. – Block Section Parameters for Material Inversion

```

block 50
  joint2g
  kx = elastic 2.474e5
  ky = elastic 2.e8
  kz = elastic 2.e8
  krx = NULL
  kry = NULL
  krz = NULL
  joint_elastic_bounds = 2.0e5 1e9
  joint_truth_table = yes no no no no no
end

block 51
  inverse_material_type = homogeneous
  material 1
  blkbeta 2.0e-4
end

```

Input 3.24. Block Section Example for Material Inversion

Initial guesses for the material parameters are also given in the **material** block. In the case of elastic materials, the initial guess is specified as the values for G and K that are input in the **material** block. In the example shown in input 3.25, the shear and bulk moduli are given initial guesses of 100 and 200, respectively.

For viscoelastic materials, the real and imaginary components of G and K are frequency-dependent. Thus the initial guesses are specified in functions that define the values of these parameters as a function of frequency. As shown in input 3.25, in this example functions 1 – 4 specify the initial guesses for the real and imaginary components of G and K .

```

material 1
  isotropic
  G_bounds 0 1e4
  K_bounds 0 1e4
  G 100.0
  K 200.0
  density 10.0
end
material 2
  isotropic_viscoelastic_complex
  Greal_bounds 0 1e4
  Kreal_bounds 0 1e4
  Gim_bounds 0 1e2
  Kim_bounds 0 1e2
  Greal = function 1
  Gim = function 2
  Kreal = function 3
  Kim = function 4
  density = 10.0
end
material 3
  isotropic
  E_bounds 1e-10 1e40
  Nu_bounds -0.9999 0.49999
  E 1.0e6
  Nu 0.25
  density 10.0
end
end

```

Input 3.25. Material Section Example for Material Inversion

```

material 4 // Full orthotropy
  density = 1.0
  orthotropic
  Cij = 4.0 1.0 2.0
  5.0 3.0
  6.0
  2.0
  1.0
  3.0

  Eii_bounds = 0.01 20.0 0.01 20.0 0.01 20.0

```

```

    Gij_bounds = 0.01 10.0 0.01 10.0 0.01 10.0
    Aij_bounds = 0.0 0.707 0.0 0.707 0.0 0.707
end

material 5 //Transeverse isotropy
    density = 1.0
    orthotropic
    alphaparametrization no
    inequalityconstraints no
    transverselyisotropic yes 23
    Cij = 5 1 1
        16 6
    16
    5
    1
    1
    Eii_bounds = 1 100 1 10 1 100
    Gij_bounds = 0.1 5 0.1 5 0.1 5
end

```

Input 3.26. Material Section Example for Orthotropic Material Inversion

The `directfrf-inverse` method supports viscoelastic material identification using homogenized moduli data by setting `data_type` to `moduli` in the **inverse-problem** block (see section 3.39.3). This feature is only available for the `inverse_material_type` `homogenous` option. In addition, all material blocks, including known materials, must be specified as `isotropic_viscoelastic_complex`.

Figure 3-99 shows which material parameters can be optimized in the different solution procedures. In general, parameters associated with damping can only be optimized in **directfrf** solutions, and those associated with stiffness can be optimized in either **directfrf** or **eigen**.

Damage Identification: Damage identification is a design variable implemented for `directfrf-inverse` problems as a special case of elastic material identification. In damage identification, a damage phase field variable is used to interpolate between damaged (weak) and full-strength material, ideally converging to near-binary values to indicate presence of material damage. Damage ID is activated by `design variable = damage` in the **inverse-problem** block.

Damage identification employs techniques originated for topology optimization problems but may be used for a variety of applications, including identification of weakened material regions, determination of contact area in a thin layer of elements at an interface, or two-phase material design. The elastic damage model uses a Solid Isotropic Material with Penalization (SIMP) model [10] that interpolates the isotropic elastic and mass density

Parameter	directfrf-inverse	eigen-inverse
K	yes	yes
G	yes	yes
E	yes	yes
Nu	yes	yes
K_real	yes	no
K_imag	yes	no
G_real	yes	no
G_imag	yes	no
Cij	yes	no
c0	yes	no
density	yes	no
joint2g elastic	yes	yes
joint2g damper	yes	no
blkbeta	yes	no

Table 3-99. – Table of Supported Material Parameters for Inverse Methods

properties in the unknown material between two phases using the scalar phase field variable $\beta \in [0, 1]$, expressed as

$$G(\beta) = G_0 + (G_u - G_l)\beta^p \quad (3.38)$$

$$K(\beta) = K_0 + (K_u - K_l)\beta^p \quad (3.39)$$

$$\rho(\beta) = \rho_0 + (\rho_u - \rho_l)\beta^q, \quad (3.40)$$

where $\{G_u, K_u, \rho_u\}$ and $\{G_l, K_l, \rho_l\}$ are the upper and lower bounds for the bulk modulus, shear modulus, and density, respectively. Definition of different powers $p \geq q \geq 1$ for the elastic and mass density components renders elastic properties relatively weaker, for a given mass density, thus disincentivizing intermediate density values. Powers p and q are specified with `penalizationElasticity` and `penalizationMass`, respectively.

Bounds of the material interpolation are specified by setting limits for shear modulus `G_bounds`, bulk modulus `K_bounds`, and mass density `density_bounds` in the **material** block. Meanwhile, the initial damage phase field value is controlled using the `G`, with respect to `G_bounds`, as

$$\beta^{(0)} = \frac{(G - G_l)}{(G_u - G_l)}. \quad (3.41)$$

Initial `density` and `K` values must still be given, though are not used to determine the initial phase field value. The solution for β is not available for output, but rather the material parameters are output as evaluated within the penalized elastic and mass density models using the phase field solution.

Damage identification is enabled for both homogeneous and heterogeneous unknown material blocks. In the homogeneous unknown material block case, only the penalization

powers for the mass density and elastic properties must be specified. In heterogeneous unknown material blocks, additional filtering and projection operations are employed to control solution length scale and to encourage binary quality. Kernel filtering prevents development of mesh dependent or checkerboard (i.e. alternating 0-1 phase field density) patterns. In this strategy, a weighted kernel is convolved with the density field, producing a locally-averaged filtered field. The filter kernel radius is defined in the **inverse-problem** block as `filter_radius` (default = 0.1), which should be set no smaller than the minimum distance between element centroids in the unknown material block.

A Heaviside-approximating function is then used to map filtered values closer to 0 or 1 bounds and recover a more-binary valued field. The smooth Heaviside function is defined

$$\tilde{\beta}(\beta) := \frac{\tanh(\zeta\eta) + \tanh(\zeta(\beta - \eta))}{\tanh(\zeta\eta) + \tanh(\zeta(1 - \eta))}. \quad (3.42)$$

Here, the slope of the smooth Heaviside $\zeta > 1$ is specified by `smooth_heaviside_slope`, while its inflection threshold $\eta \in (0, 1)$ is specified by `smooth_heaviside_threshold`. Typically, modest values for `smooth_heaviside_slope` (5-10) can produce high contrast fields without creating an excessively severe projection, which can impede optimization convergence.

Below, we summarize the necessary parameters in the **material** and **inverse-problem** sections for damage identification problems.

Parameter	type	Description
G	real	initial value of shear modulus (determines initial phase field value)
K	real	initial value of bulk modulus
density	real	initial value for mass density
G_bounds	real real	lower and upper bounds on shear modulus interpolation
K_bounds	real real	lower and upper bounds on bulk modulus interpolation
density_bounds	real real	lower and upper bounds on mass density interpolation

Table 3-100. – Parameters in **material** section for Damage Identification

Parameter	type	Description
design_variable	string	damage
penalizationElasticity	integer	elasticity penalty exponent (default = 3)
penalizationMass	integer	elasticity penalty exponent (default = 1)
filter_radius	real	filter kernel radius (default = 0.1)
smooth_heaviside_slope	real	smooth heaviside projection slope (default = 1)
smooth_heaviside_threshold	real	smooth heaviside inflection threshold (default = 0.5)

Table 3-101. – Parameters in **inverse-problem** section for Damage Identification

3.39.6. Loads section for Load Identification

For inverse load problems, the **loads** section provides additional options to control the optimization strategy. Inverse loads are currently supported for acoustics and structures. Inverse acoustic loads are only supported for sidesets. Inverse structural loads are supported on both sidesets and nodesets. On sidesets, both pressures and traction loads may be optimized in the inverse problem. For nodesets, both forces and moments can be optimized, the latter only being applicable in the case when the nodeset in question has rotational degrees of freedom (e.g. a concentrated mass). Table 3-102 summarizes the load options which apply to inverse methods. The **inverse_load_type** options are detailed in Table 3-103.

The default behavior for the **inverse_load_type** keyword is **known**, which implies that the loads on that sideset or nodeset are fixed, and are not modified in the optimization process. In the case where all blocks were known, there would be no need to solve the inverse loads problem.

In input 3.27 sideset 1 is **known**, sideset 2 is a real-valued, unknown acoustic load that is **spatially_constant**, and sideset 3 is an imaginary-valued, unknown acoustic load that is **spatially_constant**. We note that in the case of a transient problem, the **loads** block in an inverse loads problem would look the same except that there would be no imaginary loads in that case.

For acoustic problems currently inverse acoustic loads are limited to the **acoustic_accel** option. The **acoustic_vel** keyword is not supported for acoustic loads.

Input 3.28 shows a similar example for a structural pressure, traction and force load case. This example contains a known pressure load, and unknown pressure, traction, and force loads. Note that in the case of traction and force loads, the direction of the traction load (in this case 111) is fixed, and only the function amplitudes are calculated in the inverse problem.

In input 3.28, functions 1 – 4 contain the initial guesses for the load amplitudes for sidesets 1 – 3 and forces/moments on nodeset 4. These load amplitudes are then refined during the optimization process. The resulting loads are written to a text file called *force_function_data.txt*, which could then be included in a subsequent forward or inverse loads case in a restart analysis.

Parameter	type	Description
inverse_load_type	string	load inverse type

Table 3-102. – Loads Section Parameters for Force Inversion

```
loads
  sideset 1
    acoustic_accel = 1
    function = 1
  sideset 2
```

Parameter	Description
spatially_constant	Load amplitude is unknown and will be optimized during the inversion process, but is treated as constant over the entire sideset or nodeset.
spatially_variable	Load amplitude is unknown and will be optimized during the inversion process. Each dof on the sideset or nodeset is optimized.
known	default. No optimization performed.

Table 3-103. – Inverse Load Type Options

```

    acoustic_accel = 1
    function = 2
    inverse_load_type = spatially_constant
sideset 3
    iacoustic_accel = 1
    function = 3
    inverse_load_type = spatially_constant
end

```

Input 3.27. Loads Section Example for Acoustic Force Inversion

```

loads
  sideset 1
    pressure = 1
    function = 1
  sideset 2
    pressure = 1
    function = 2
    inverse_load_type = spatially_constant
  sideset 3
    traction = 1 1 1
    function = 3
    inverse_load_type = spatially_constant
  nodeset 4
    force = 0 1 1
    function = 4
    inverse_load_type = spatially_constant
  nodeset 4
    moment = 0 1 1
    function = 5
    inverse_load_type = spatially_constant
end

```

Input 3.28. Loads Section Example for Structural Force Inversion

3.39.7. Limitations for Inverse Load Problems

Limitations: There are a number of limitations which apply to transient load identification. These include the following.

1. Structural, acoustic and structural-acoustic domains may be addressed.
2. Only the simple Newmark integrator should be used, i.e. do not use generalized alpha integration and do not use the `rho` keyword.
3. Pressure is always applied along the surface normal, and tractions are applied along the direction specified in the `loads` block. Inverse methods do not support follower pressures.
4. Load identification applies to `acoustic_accel` loadings on acoustic sidesets, pressures/tractions on structural sidesets, and forces and moments on structural nodesets.
5. In force identification problems, a pressure, traction or force may be applied on a shell. The measured displacement fields in the truth table for shells can be applied to nodes with rotational DOFs; however, only displacement DOFs can be specified in the data files. Specifying rotational DOFs as measured data at nodes is not supported.

3.39.8. ROL Output for Inverse Problems

Sierra/SD uses the Rapid Optimization Library (ROL), which is part of Trilinos [33] for solving optimization problems. During the optimization process, ROL writes out a text file called *ROL_Messages.txt* that contains information about the convergence of the optimization solution. It is important to examine this file to assure that the solution is adequately converged.

An example of a *ROL_Messages.txt* is given in Figure 3-32. The first 2 lines show the optimization method that was used by ROL, and the following lines contain convergence information. Each line corresponds to a single optimization iteration. The first column shows the iteration number under the *iter* heading. The second column shows the value of objective function at that iteration, under the *value* heading. The third column shows the absolute norm of the gradient (i.e. the derivative of the objective function with respect to the optimization variables). In Figure 3-32 we only show the first three columns of output as these will typically be of most interest to the user, but the remaining columns contain information about step size, number of function evaluations, etc... (Denis, Drew, any input here would be great).

For a typical **Sierra/SD** user, the first three columns in the *ROL_Messages.txt* file will typically be of the most important to pay attention to. As the goal is to minimize the objective function, a substantial decrease in the second column should be observed. Also, the desired minimum of the optimization corresponds to a zero gradient, and thus the third column should be observed to be as close to zero as possible.

Newton-Krylov with Cubic Interpolation Linesearch satisfying Null Curvature Condition Krylov Type Conjugate Gradients		
iter	value	gnorm
0	5.171112e-03	7.337197e-03
1	1.160506e-10	1.245979e-06
2	1.453148e-17	4.667996e-10

Figure 3-32. – Example of *ROL_Messages.txt* file for Inverse Problem Solution

4. MATERIALS

The material section has a (unique) material identifier (an integer or a string name). The material identifier is used in assigning material properties to element blocks. Material types and their parameters are summarized in Table 4-104.

Table 4-104. – Material Stiffness Parameters.

material type	parameters
isotropic	any two of K , G , E or ν
orthotropic	nine C_{ij} entries
orthotropic_prop	$E1$, $E2$, $E3$, $\nu23$, $\nu13$, $\nu12$, $G23$, $G13$, $G12$
anisotropic	21 C_{ij} entries

For example,

```
Material steel
  isotropic
  E 3e7
  nu .3
End
```

A materials may be **isotropic** , **orthotropic** , **orthotropic_prop** , **anisotropic** , or **isotropic_viscoelastic** .

The Joint2G element 5.21 has material models for joints including an elastic-plastic model.

4.1. Elastic

Elastic materials may be isotropic section 4.1.1, orthotropic section 4.1.2, or anisotropic section 4.1.3. Some material models from the Lamé library are available section 4.1.4.

4.1.1. Isotropic

Isotropic materials require specification of two of the following parameters. They can be defined directly as `parameter = <real>`, as functions of temperature (Section 4.4.6), or as spatially dependent properties (Section 4.4.7.)

Parameter	Description
E	Young's Modulus
nu	Poisson's Ratio
G	Shear Modulus
K	Bulk Modulus

Isotropic materials are the default, and the keyword **isotropic** is not required. Of the four parameters, *exactly two* must be supplied. They are related by

$$E = 3K(1 - 2\nu), \quad G = \frac{3KE}{9K - E}.$$

Internally, **Sierra/SD** stores the values of **K** and **G**.

4.1.2. Orthotropic

Orthotropic material entry is similar to the anisotropic case.

A difference is that the keyword **orthotropic** replaces **anisotropic**, and only 9 C_{ij} entries are specified. These entries correspond to C_{11} , C_{12} , C_{13} , C_{22} , C_{23} , C_{33} , C_{44} , C_{55} and C_{66} . Like the anisotropic material definition, we follow the stress/strain ordering of xx , yy , zz , zy , zx , xy .

Alternatively, an orthotropic material may be specified using **orthotropic_prop** and the material parameters **E1**, **E2**, **E3**, **nu23**, **nu13**, **nu12**, **G23**, **G13**, and **G12** as shown in the following example. As with isotropic materials, temperature-dependent parameters may be defined via a function as `parameter = function <string>` (see Section 4.4.6). Note that all elastic materials must satisfy requirements that the elasticity matrix is positive definite.

```
Material honeycomb
  orthotropic_prop
  E1 = 508.7
  E2 = 7641.0
  E3 = 14750.0
  nu12 = .2
  nu23 = .0825
  nu13 = .1
  G12 = 115
  G23 = 2320.
```

```
G13 = 450.  
density=0.5  
End
```

A single orthotropic layer may be specified using **orthotropic_layer** . An orthotropic layer must specify 4 of the above parameters (E1, E2, nu12, G12). The syntax **parameter = from_transfer** may also be used to **receive_sierra_data** solution case to obtain material parameter values from a file transfer. Here is an example:

```
Material 13  
  orthotropic_layer  
  E1 = 508.7  
  E2 = 7641.0  
  nu12 = 1.293  
  G12 = 115  
  density=0.5  
End
```

If sensitivity analysis is being performed (see Section 2.6), one indicates the parameters for analysis by following these parameters with the +/- characters. In the first entry method, a sensitivity analysis must be performed on all 9 parameters. In the second, each individual parameter must be requested individually. The concept is that the sensitivity is performed with respect to the labeled parameters, i.e. either the set of C_{ij} parameters, or each individually labeled E1 term.

4.1.3. Anisotropic

Anisotropic materials require specification of a 21 element C_{ij} matrix corresponding to the upper triangle of the 6×6 stiffness matrix. Data is input in the order C_{11} , C_{12} , C_{13} , C_{14} , C_{15} , C_{16} , C_{22} , etc. The C_{ij} must be preceded by the keyword **Cij** . The keyword **anisotropic** is also required. Materials are specified according to the stress/strain ordering xx , yy , zz , yz , zx , xy . This is generally consistent with published Materials Science data. However, NASTRAN and Abaqus use a different convention. An input deck illustrating anisotropic material input is provided in Section 9.2.

If an element block uses a coordinate system the anisotropic material is defined in the $\hat{r}$, $\hat{s}$, $\hat{t}$ local frame 7.1.6. If an element block does not use a coordinate system the anisotropic material is defined in the X, Y, Z frame. The frame used by each element may be visualized by requesting **material_direction_1**, **material_direction_2**, and **material_direction_3** in the outputs block.

For anisotropic (and orthotropic) materials, the check to make sure material properties are acceptable is skipped. A message is printed notifying the user that this check is skipped.

4.1.4. Lamé Material



Lamé material is currently BETA release.
Enable with the “`--beta`” command-line option.

Sierra/SD provides a limited capability to use linearized versions of the non-linear Lamé material models in **Sierra/SD**. This can be used in conjunction with the capability to **hand-off** a nonlinear preload to a linear **Sierra/SD** analysis. This hand-off is accomplished by reading in the element variables **stress** , **left_stretch** , **rotation** , and Lamé state (e.g. **lame_state_hyperfoam**) at the last time-step of your **Exodus** input file. One example where this would be useful is computing the tangent stiffness of a compressed foam.

Lamé materials are defined in a material section. As in input 4.1. A Lamé material definition begins with **begin-lame-material** and ends with **end-lame-material**. Currently, only Neo-Hookean and Hyperfoam materials are supported.

```
Material 1
  begin-lame-material
    begin parameters for model Hyperfoam
      bulk modulus      = 1.e6
      Poissons ratio    = 0.1
      n = 3
      shear = 3.74e6, -3.17e6, 1.18e4
      alpha = 2.536, 2.090, -8.807
      Poisson = 0.5630, 0.5507 0.3662
    end
  end-lame-material
  density = 5.0
End
```

Input 4.1. Example material section for a Lamé Hyperfoam material model.

Note: unlike the rest of a **Sierra/SD** input file, the material model definition (emphasized text in input 4.1) must strictly follow syntax rules including newlines. Details of the allowed syntax for each material model are given below. **Lamé Neo-Hookean model**
Acceptable syntax for Neo-Hookean models is given below.

```
BEGIN PARAMETERS FOR MODEL NEO_HOOKEAN
#
# Elastic constants
#
YOUNGS MODULUS = <real> E
POISSONS RATIO = <real> ν
SHEAR MODULUS  = <real> G
```



```

BULK MODULUS    = <real>  $K$ 
LAMBDA          = <real>  $\lambda$ 
TWO MU          = <real>  $2\mu$ 
END [PARAMETERS FOR MODEL NEO_HOOKEAN]

```

A detailed discussion of the theory of the Neo-Hookean model can be found in the Sierra Solid Mechanics User Manual.⁶⁹

Lamé Hyperfoam Model Acceptable syntax for Hyperfoam models is given below.

```

BEGIN PARAMETERS FOR MODEL HYPERFOAM
#
# Elastic constants
#
YOUNGS MODULUS = <real>  $E$ 
POISSONS RATIO = <real>  $\nu$ 
SHEAR MODULUS  = <real>  $G$ 
BULK MODULUS   = <real>  $K$ 
LAMBDA         = <real>  $\lambda$ 
TWO MU         = <real>  $2\mu$ 
#
# Strain energy density
#
N = <integer>  $N$ 
SHEAR    = <real_list>  $\mu_i$ 
ALPHA    = <real_list>  $\alpha_i$ 
POISSON  = <real_list>  $\nu_i$ 
END [PARAMETERS FOR HYPERFOAM]

```

As with Neo-Hookean models, a detailed discussion of the theory of the Hyperfoam model can be found in the Sierra Solid Mechanics User Manual.⁶⁹

4.2. *Acoustic*

Linear acoustic materials require the specification of the fluid density, and the linear speed of sound. In addition, the keyword **acoustic** must be in the material block.

```

Material air
  acoustic
  density 1.293
  c0 332.0
  cavitating
  pvapor 0.0
End

```

Nonlinearity can be activated by the keyword **nonlinear**. Nonlinear acoustic materials require one additional parameter, **B_over_A**, which is a measure of fluid nonlinearity. For air, **B_over_A**=0.4. Tables of **B_over_A** for various fluids can be found in.³²

Cavitation can be activated by the keyword **cavitating**. This requires an additional parameter, **pvapor**, which is the vapor pressure with a default value equal to 0. Cavitating elements also require the definition of vectors for **hydrostatic_gravity** and **free_surface_point** in the **block** section as shown in input 4.2. **hydrostatic_gravity** and **free_surface_point** are vectors used to compute the external pressure at the block.

```
BLOCK 1
  material 1
  hydrostatic_gravity = 0 0 -32.2
  free_surface_point  = 0 0 34.0
End
```

Input 4.2. This is an example of a **block** section for a cavitating material.

For computational acoustics see Section 2.4.4.

4.3. *Linear Viscoelastic*

Linear viscoelastic materials require the specification of the density, and the limiting moduli E_g , E_{inf} , G_g , G_{inf} . The subscript 'g' refers to the glassy modulus, which occurs at $t = 0$, or $\omega = \infty$. The subscript 'inf' refers to the rubbery modulus, which occurs at $t = \infty$, or $\omega = 0$. In addition the Prony series for the viscoelastic materials have to be specified using keywords **K_coeff**, **K_relax**, **G_coeff**, and **G_relax**. All of these parameters are required.

For the bulk modulus K , the Prony series parameters are defined by the following equation:

$$K(t) = K_{inf} + (K_g - K_{inf}) \sum_i K_{coeff}[i] * e^{-\frac{t}{K_{relax}[i]}} \quad (4.1)$$

A similar equation holds for the shear modulus. Note that, the **K_coeff** and **G_coeff** *must* sum to 1.0 (individually). Otherwise, the formulation is inconsistent. That is,

$$\sum_i K_{coeff}[i] = \sum_i G_{coeff}[i] = 1.0. \quad (4.2)$$

Note that the number of terms in **K_coeff** and **K_relax** must be the same, and the number of terms in the **G_coeff** and **G_relax** must be the same. However, the number of terms in the **K** series does not have to equal the number of terms in the **G** series. Thus, one could simulate a case where the material shear modulus G is viscoelastic, but the bulk modulus is not. In this case, the latter would have no terms in its series.

E_g , E_inf , G_g , G_inf may be constant or may depend on temperature. Temperature functions can specify the value for the limiting moduli, for a given value of temperature. For example, if the limiting moduli typically depend linearly on temperature, a linear function can be specified for the values of E_g , E_inf , G_g , G_inf . We refer to the example given below for the specifics on how to set this up.

Optional parameters for viscoelastic materials include reference (T_0), glassy (T_g). In the event that none of the moduli are specified as functions, the values specified for these parameters determine which model is used for temperature-dependent behavior. The T_g parameter may be given as a constant or as a spatially dependent property (Section 4.4.7.)

Two such models are available: the WLF (Williams-Landel-Ferry) model,⁷³ and the Hinnerich's model. The WLF model^{1,29} is used when shifting temperatures above T_g , while the Hinnerich's model is used shifting temperatures below T_g . Both models incorporate the reference temperature T_0 (the reference temperature at which the input viscoelastic constants are defined), which may differ from T_g . The WLF and Hinnerich's model use model-specific constants. We note that any units of temperature can be used, as long as they are consistent with the values of the constants. The shift factors computed from the Hinnerich's or WLF equations are used to scale the coefficients in the Prony series.

The WLF model is

$$\log_{10}(a_T) = -\frac{C_1(T_{elem} - T_0)}{C_2 + T_{elem} - T_0}, \quad (4.3)$$

where T_{elem} is the current temperature in the element, and T_0 , C_1 , and C_2 are material parameters that are determined experimentally. Typically, T_0 is the glass transition temperature of the material of interest. The shift factors computed from the WLF equation are a strong function of temperature.

The Hinnerich's model, provided by Terry Hinnerich's, accurately characterizes many viscoelastic materials below the glassy transition temperature. Its form is

$$\log_{10}(a_T) = a_{T1} * (1 - e^{a_{T2} * (T_{elem} - T_0)}), \quad (4.4)$$

where a_{T1} and a_{T2} are user-specified constants. This equation used to determine an approximate set of shift factors when experimental data for a particular material is not at hand.

Note if a model is shifted through T_g a composite shift is used. For example if T_0 is less than T_g and T_{elem} is above T_g first a Hinnerich's shift is used to shift parameters from T_0 to T_g then a WLF shift is used to transition those constants from T_g to T_{elem} . These shifts are automatically computed given T_0 , T_g , block temperature, C_1 , C_2 , a_{T1} , and a_{T2} . Note that if these shifting parameters are not specified in the input file, then no shifting will be done and the relaxation times as specified in the input deck will be used. If T_g is unspecified then it defaults to the provided value of T_0 . Either specify all the shifting coefficients, or specify none of them. Partial specification rarely succeeds.

After computing the shift factors using one of the two approaches given above, the relaxation times are shifted. This occurs before computations begin using the relations,

$$K_{relax}[i] = a_T K_{relax}[i] \quad (4.5)$$

$$G_{relax}[i] = a_T G_{relax}[i]. \quad (4.6)$$

Example 4.3 demonstrates how Hinnerich’s model is used to set a linear viscoelastic material.

```
Material foam
  isotropic_viscoelastic
  T_0= 10
  T_g = exo_var scalar t_g_input
  C_1=15.
  C_2=35.
  aT_1=6.
  aT_2=.0614
  K_g = function 1
  K_inf 1.e7
  G_g 1.e2
  G_inf 12.
  K_coeff .5 .5
  K_relax 3. 2
  G_coeff .5 .5
  G_relax 1 3
  density 0.288
End
```

Input 4.3. Hinnerich’s viscoelastic material specification

Note that the coefficients of both K and G sum to 1.0. This is necessary for a consistent formulation. Also, in this case we specify a temperature function for K_g . Thus, the value of K_g used in the simulations is the value of function 1, at the particular element temperature T_{elem} . The T_g value is shown reading from an input mesh exodus field named ‘t_g_input’. T_g can also be specified as a constant like the other parameters.

4.3.1. Limitations of Viscoelastic Use.

Linear viscoelastic materials are “linear” in the sense that (linear) transient dynamics accommodates them exactly. There are limitations for the use of these materials in **Sierra/SD**.

1. When using viscoelastic materials in a **nonlinear transient** simulation, it is necessary to specify “**nonlinear=no**” in the BLOCK section of the viscoelastic block. This is because different internal force mechanisms are called for linear and nonlinear cases, and viscoelastic materials in **Sierra/SD** only support a linear constitutive model and small deformation.
2. When viscoelastic materials are used in a **statics** simulation, the material is assigned the properties G_{inf} and K_{inf} . This is because in a slow (static) loading, the material responds with these material properties since they are the long-time or slow response properties.
3. Likewise, **Eigen** solutions and the modal based solutions derived from them apply only the first terms of the Prony series for G_{inf} and K_{inf} , which are used to define the elastic constants of an isotropic elastic material. This is because the real modal solution must use a constant mass and stiffness matrix, and has no damping contribution.
4. It is possible to evaluate the eigenvalues expanded about a given frequency, **viscofreq**. See the discussion of **ceigen** in Section 3.24.2. The damping matrix, C , is taken into account in the eigenvalue problem. Few modal solutions are adapted to use these complex modes.

Complex Viscoelastic Complex isotropic viscoelastic materials may be defined for directFRF solution cases with the **isotropic_viscoelastic_complex** keyword. This material type has 4 required parameters, representing the real and imaginary shear and bulk moduli: **Greal**, **Gim**, **Kreal**, and **Kim**. (The real and imaginary components are commonly known as the storage and loss moduli, respectively.) Each of these complex viscoelastic parameters must be defined by a frequency-varying function, e.g. **parameter = function <string>**.

Example 4.4 demonstrates how a complex linear viscoelastic material is set.

```
Material 99
  isotropic_viscoelastic_complex
  Kreal = function 1
  Kim = function 2
  Greal = function 3
  Gim = function 4
  density 1.0
End
```

Input 4.4. Viscoelastic material specification

4.4. Properties

4.4.1. Density

For solutions requiring a mass matrix, all material specifications must define density. This can be set via the keyword **density** followed by a scalar value. Alternatively, the density can be defined as a temperature dependent property (Section 4.4.6), as a spatially dependent property (Section 4.4.7.), or be defined via a file transfer to `receive_sierra_data` solution case with (`density = from_transfer`).

4.4.2. High Cycle Fatigue

Material parameters for high cycle fatigue (Section 3.14) may be provided to define the statistical failure behavior of materials. These properties are summarized in Table 4-105.

Parameter	Type	Default	Description
Fatigue_A1	<i>Real</i>	0.0	S-N curve constant
Fatigue_A2	<i>Real</i>	-3.0	S-N curve slope
Fatigue_A3	<i>Real</i>	0.0	S-N curve translation, $S_{eq} = (1 - R)^{A_3}$
Fatigue_A4	<i>Real</i>	0.0	S-N curve endurance limit. This parameter is not used in Sierra/SD fatigue computations.
Stress_Ratio	<i>Real</i>	-1.0	R , the ratio of max/min stress. The default -1 indicates oscillation about a zero-mean stress state
Fatigue_A	<i>Real</i>	1.0	S-N curve constant
Fatigue_m	<i>Real</i>	3.0	S-N curve coefficient
Fatigue_Stress	<i>Real</i>	1.0	Stress unit conversion factor
std_err	<i>Real</i>	0.0	to shift S-N curve by material uncertainty
t_dist	<i>Real</i>	0.0	to shift S-N curve by material uncertainty

Table 4-105. – Material Section Parameters for Fatigue Parameters.

4.4.3. S-N curve Definitions

There are two competing S-N curve definitions in literature, which are equivalent for S-N curves that are linear in log-log space, and are both supported by **Sierra/SD**. The first is used by Wirsching, Paez, and Ortiz in their book *Random Vibrations Theory and Practice*,⁷⁵ and is supported with the `Fatigue_A`, and `Fatigue_m` parameters:

$$NS^m = A$$

The second is adaptable to a wider variety of materials, and is defined in chapter 9 of MMPDS²⁶ (otherwise known as MIL-HDBK-5) as:

$$\log_{10}(N) = A_1 + A_2 \log_{10}(S(1-R)^{A_3} - A_4)$$

For real materials, A_1 is always positive, and A_2 is always negative. A_2 gives the S-N curve its negative slope, and A_1 represents the crossing of the S-axis on the S-N curve. For the case of an S-N curve which is linear in log-log space, $A_4 = 0$ and the analysis is assumed to operate at a constant stress ratio such that:

$$S_{eq} = S(1-R)^{A_3} = S(2^{A_3})$$

and

$$\log_{10}(N) = A_1 + A_2 \log_{10}(S_{eq})$$

These functions are equivalent, and the material properties can be mapped to each other by:

$$A_1 = \log_{10}\left(\frac{A}{(1-R)^{A_2 \cdot A_3}}\right), \quad A_2 = -m$$

Note that A_4 is equivalent to the endurance limit of the material, if it exists. Parameters A_1 , A_2 , A_3 , and A_4 can all be found in MIL-HDBK-5²⁶ as empirically derived values for most metallic materials. See the example in Section 4.4.5. Since the Narrowband approach has an inherent assumption of zero mean stress, the default is $R = -1$. While the user can specify a different R-ratio, such usage would be inconsistent with the Narrowband and Wirsching methods.

4.4.4. S-N Curve Units

While **Sierra/SD** requires only a consistent set of units, the introduction of experimental data with their own units can confuse the solution. This is particularly challenging because the units for some parameters are mixed, and the data is gathered and presented in only a single unit system.

The stress scaling parameter helps reduce that problem. Consider the scaled equation for narrow band damage.

$$D_{NB} = \frac{\nu_o^+ \tau}{A} (\sqrt{2} \sigma_s F_{ss})^m \Gamma\left(\frac{m}{2} + 1\right) \quad (4.7)$$

Here F_{ss} is the stress scaling parameter. This parameter lets the user convert from the unit system of the analysis to the unit system of the test data. Table 4-106 provides common conversions of these stresses.

Model Units	Experimental Units			
	PSI	Ksi	SI	CG
PSI (lbs/in ²)	1	0.001	6894.76	68947.6
Ksi (K lbs/in ²)	1000	1	6894757	68947573
SI (N/m ²)	.000145037738	1.4503774e-7	1	10
CG (dynes/cm ²)	.0000145037738	1.4503774e-8	0.1	1

Table 4-106. – Common Unit Scalings using **Fatigue_Stress_Scale**. The MIL-HNBK typically uses Ksi for experimental units.

4.4.5. Typical Material Data for Fatigue

Figure 4-33 shows typical fatigue data from MIL-HDBK-5.²⁶ For a range of stress ratios in this material, the number of cycles to failure is represented by the equation,

$$\log_{10} N_n = 9.65 - 2.85 \log_{10}(S_{eq} - 61.3)$$

In this range, and for this material, we have the definitions given in Table 4.4.5.

Parameter	Value	Comment
Fatigue_A1	9.65	offset in S-N curve. $A_1 = \log(A)$
Fatigue_A2	-2.85	slope of S-N curve, $-m$
Fatigue_A3		S-N curve translation, $S_{eq} = (1 - R)^{A_3}$
Fatigue_A4	61.3	endurance limit
Stress_Ratio	-0.60	R , determines equivalent stress, $S_{eq} = (1 - R)^{A_3}$
Fatigue_Stress_Scale	0.001	Stress unit conversion to Ksi

MIL-HDBK-5J
31 January 2003




Figure 2.3.1.2.8(a). Best-fit S/N curves for unnotched 4130 alloy steel sheet, normalized, longitudinal direction.

Correlative Information for Figure 2.3.1.2.8(a)

Product Form: Sheet, 0.075 inch thick

Properties: TUS, ksi TYS, ksi Temp., °F
117 99 RT

Specimen Details: Unnotched
2.88-3.00 inches gross width
0.80-1.00 inch net width
12.0 inch net section radius

Surface Condition: Electropolished

References: 3.2.3.1.8(a) and (f)

[Caution: The equivalent stress model may provide unrealistic life predictions for stress ratios beyond those represented above.]

Test Parameters:

Loading - Axial
Frequency - 1100-1800 cpm
Temperature - RT
Environment - Air

No. of Heats/Lots: Not specified

Equivalent Stress Equations:

For stress ratios of -0.60 to +0.02
 $\log N_f = 9.65 - 2.85 \log (S_{eq} - 61.3)$
 $S_{eq} = S_{max} (1-R)^{0.41}$
Std. Error of Estimate, $\log (\text{Life}) = 0.21$
Standard Deviation, $\log (\text{Life}) = 0.45$
 $R^2 = 78\%$

Sample Size = 23

For a stress ratio of -1.0
 $\log N_f = 9.27 - 3.57 \log (S_{max} - 43.3)$

4.4.6. Temperature dependence

Material properties in **Sierra/SD** can be specified to be temperature dependent. Temperature dependent material properties are supported when temperatures are read in from an **Exodus** file, or when they are specified on a block-by-block basis. In the case of **Exodus** temperatures, the material properties would vary from element to element, since the temperatures vary with each element. The temperature dependent material properties are calculated from an elemental average of the **Exodus** temperatures. When temperatures are specified on a block-by-block basis, the temperature-dependence of the material properties can be specified explicitly in the input deck. If temperatures are specified in the **Exodus** file and block-by-block in the input deck, then the input deck values take precedence.

For linear elastic materials, an example of specifying temperature dependent properties is given below.

```
Material 1
    E function=eTempFunc1
    alphasat .001
    tref 100
    nu 0.0
    density 7700.0
End

Material 2
    E function=eTempFunc2
    alphasat .001
    tref 100
    nu 0.0
    density 7700.0
End

FUNCTION eTempFunc1
    type LINEAR
    data 0.0 4.0
    data 5.0e9 4.0
End

FUNCTION eTempFunc2
    type LINEAR
    data 0.0 3.0
    data 5.0e9 3.0
End
```

In this case, the elastic modulus of material 1 is specified by function eTempFunc1, and the

elastic modulus of material 2 is specified by function `eTempFunc2`. The moduli of each element will be determined from its temperature and an interpolation on the function. In this example, the functions are trivial, and thus the moduli of materials 1 and 2 will be 4 and 3, respectively. Note that the moduli, density and any of the 4 elastic constants k , g , e , ν can be specified as temperature dependent, and can be given different functions. In this example, the Poisson's ratio is constant and only the elastic modulus is temperature dependent.

For viscoelastic materials, functions do not need to be specified in the **material** block to designate temperature dependence of the shift factors. This is accounted for automatically. See Section 4.3 on viscoelastic materials for more details.

In addition to linear elastic and linear viscoelastic material properties, the coefficient of thermal expansion **alphat** may be given temperature-dependent material properties.

If the **thermal_time_step** keyword is used the temperature that effects material properties will be based on the temperature read from the provided time step. Alternatively the **nUpdateDynamicMatrices** keyword can be given which will update the material stiffness based on the last read temperature. The last read temperature is controlled by the **nUpdateTemperature** keyword. Updating the dynamic matrices is computationally expensive and should be done only when temperature has changed significantly.

4.4.7. Spatially Variant Material Properties

Multiple methods exist to define element-to-element spatial dependence in material properties.

Some material properties in **Sierra/SD** can be read on an element-by-element basis from the **Exodus** mesh file. The syntax for this is `parameter=exo_var scalar <string>`. Here `<string>` is the provided exodus field name. This must be a scalar field (one component) defined on each element using the material.

For linear elastic materials, an example of specifying exodus based properties is given below.

```
Material 1
    E = exo_var scalar e_input
    nu = exo_var scalar my_poisson
    density = exo_var_scalar elem_density
End
```

In this case, the elastic modulus of material 1 is specified by the exodus mesh field 'e_input', the Poisson's ratio by 'my_poisson', and the density by 'elem_density'. The exodus based properties can be used to define complex spatial dependence of material properties as may be found in partially compressed foams for example. No time-variance of material properties is considered, if the input mesh exodus file has multiple time steps the material property fields should be constant over time to avoid confusion.

Additionally it is possible to define material properties through the general function `function` syntax. Syntax and requirements are detailed in table 2-26. Unless otherwise specified material property functions are evaluated as a function of temperature. The temperature can be defined via a block temperature in the input deck, an element-by-element temperature in the exodus input mesh, or a nodal temperature (which is then interpolated to the element centroid temperatures) in the exodus input mesh.

```
Material tempFunc
  E = function e_temperature_func
  nu = 0.3
  density = 1
End

Function e_temperature_func
  type linear
  data 0      1.0e+6
  data 400    1.0e+6
  data 500    0.9e+6
  data 900    0.3e+6
End
```

Alternatively specific input variables other than temperature can be explicitly defined in analytic functions. Nodal function responses are mapped to element material properties by evaluating the function at the nodes, and applying element shape functions to interpolate those values to the centroid.

```
Material readsFromNodal
  E = function e_input_func
  nu = 0.3
  density = 1
End

Function e_input_func
  type analytic
  expression variable X = nodal a_mesh_var
  expression variable Y = nodal another_mesh_var
  evaluate expression "sqrt(X) + Y"
End
```

4.4.8. Specific Heat

Conversion of energy deposited in a structure to a change in temperature may be effected by a specific heat.

$$Q = \rho V C \Delta T. \quad (4.8)$$

Here Q is the total heat energy, ρ is the density, V is the volume, C is the specific heat and ΔT is the change in temperature. It is up to the analyst to ensure that consistent units are employed. Note also that the analyst must determine under what conditions the specific heat is applied (constant pressure or constant volume).

Specific heat is used only in applying boundary conditions. Energy deposited within a structure is converted to temperature using equation 4.8. Once converted to temperatures, thermal stresses and temperature dependent material properties may be applied. A fatal error is encountered if the specific heat is not specified for each material containing an energy load. The keyword `defaultSpecificHeat` defined in the “parameters” section, can be used to specify a default specific heat for all materials.

```
Material 'Steel-SI'
  E=2e11           // Pa
  NU=0.28
  density=7850      // kg/m^3
  specific heat = 0.45 // J/(gK)
  tref = 300 // K
  alphasat = 0.001
End
```

The reference temperature is used only for temperature dependent material properties, such as in viscoelastic materials. In other words,

$$\Delta T = \frac{Q}{\rho V C} \quad (4.9)$$

$$T_{\text{elem}} = T_{\text{ref}} + \Delta T \quad (4.10)$$

$$\epsilon_{\text{thermal}} = \alpha_T (T_{\text{elem}} - T_{\text{ref}}). \quad (4.11)$$

Energy loads use the energy per unit mass or specific energy,

$$\tilde{E} = \frac{Q}{\rho V},$$

as described in Section 6.3.9.

4.4.9. Frequency dependence

For the `CJdamp` solution method (see Section 3.4), a frequency dependent damping coefficient, $\eta(f)$, may be specified.⁶ All other solution methods will ignore this keyword. The `CJetaFunction` keyword requires as a parameter the identifier of a function. Its use is specified in the following example. See Section 2.8 for details in specifying the function. If no function is specified, the block will be treated as if the function were identically zero everywhere.

⁶ η is twice the normal modal damping coefficient. Thus, if $\eta=0.02$ for all materials, the equivalent modal damping will be 1 percent.

```

Material 1
  E=1e7
  NU=0.28
  density=0.098
  CJetaFunction=1
End

function 1
  name 'function to use for material 1 eta'
  type linear
  data 0.0 0.001
  data 100 0.010
  data 200 0.030
  data 400 0
end

```

The function specifies the frequency and amplitude pairs for η . The frequencies are in Hertz. The **CJdamp** solution process interpolates the function at the eigen frequencies to determine the effective damping for any particular mode.

4.5. *Piezoelectric Material*

Sierra/SD supports two material models which possess voltage degrees of freedom in addition to displacements and rotations: dielectric and piezoelectric materials. The piezoelectric material is characterized by an electro-mechanical coupling in the stiffness matrix. In general, piezoelectric materials are defined by three constitutive tensors: a rank four orthotropic elasticity tensor, a rank two anisotropic permittivity tensor, and a rank three piezoelectric coupling tensor. See theory manual for further details on the constitutive tensors. Piezoelectric material tensors may be specified in a material block by including the keyword **orthotropic_piezoelectric** followed by the required three material tensor specifications. The nine parameters defining an orthotropic elasticity tensor are given by the keyword **Cij** followed by the upper triangle of a six by six matrix. The piezoelectric coupling tensor is given by the keyword **e_ij** followed by a six by three coupling matrix. **Sierra/SD** assumes the coupling matrix is in its stress-charge form (units *charge/Area*). The anisotropic permittivity tensor is provided by the keyword **permittivity_ij** followed by a three by three matrix. The permittivity matrix should be populated by absolute permittivity values (not normalized by the permittivity of free space).

Here is an example of a PZT5A piezoelectric material where ϵ_0 is the permittivity of free space:

```

MATERIAL 1
  ORTHOTROPIC_PIEZOELECTRIC

```

```

Cij = 12.1e10  7.5e10  7.5e10  0  0  0
          12.1e10  7.5e10  0  0  0
                11.1e10  0  0  0
                        2.1e10  0  0
                                2.1e10  0
                                        2.3e10

permittivity_ij = 916 * e0  0  0
                      0  916 * e_0  0
                      0  0  830 * e0

e_ij = 0  0 -5.4
        0  0 -5.4
        0  0 15.8
        0 12.3  0
       12.3  0  0
        0  0  0

density = 7.75e3
END

```

Input 4.5. Piezoelectric Material

4.6. *Dielectric Material*

A dielectric material, the second available material possessing a voltage degree of freedom, is used to model the electro-static behavior of materials that do not exhibit electro-mechanical coupling (i.e. non-piezoelectric). Dielectrics can be generated with the keyword **dielectric**, and are defined only by an orthotropic elasticity tensor and a permittivity tensor. The following is an example of the dielectric equivalent to the PZT5A described in 4.5.

```

Material 1
  dielectric
  Cij = 12.1e10  7.5e10  7.5e10  0  0  0
            12.1e10  7.5e10  0  0  0
                  11.1e10  0  0  0
                          2.1e10  0  0
                                  2.1e10  0
                                          2.3e10

  permittivity_ij = 916 * e0  0  0
                    0  916 * e_0  0

```

```
0 0 830 * e0  
  
density = 7.75e3  
End
```

Input 4.6. Dielectric Material

4.7. *Block*

Each element block in the **Exodus** file must have a corresponding **block** section in the input file. The converse is not true — there can be **block** entries in the input deck that do not have corresponding entries in the **Exodus** file. There are two cases where this can happen:

- Virtual blocks. These are blocks that have entries in the input deck and are intended to be part of the model, but have no corresponding entries in the **Exodus** file. At this time, only Joint2g elements (see Section 5.21) can be defined to be virtual blocks.
- Extra blocks that have entries in the input deck but are not intended to be part of the model. These blocks are silently ignored by **Sierra/SD**.

It is an error to have multiple definitions for the same block. However, **Sierra/SD** does not report the error. The behavior of Sierra/SD in this case is not defined. This section contains information about the properties of the elements within the block.

4.7.1. **Block Parameters**

There are two main types of block parameters:

1. Parameters common to most elements include:
 - Material property references are required for most elements. The material reference is of the form, **material=material_id**, where **material_id** is a string representing the material identifier (see Section 4).
 - coordinate frames - *optional*
 - nonlinear behavior - *optional*
 - block damping - *optional*
 - non-structural mass - *optional*
2. Element-specific parameters. These properties depend on the element type. A block parameter applies to all elements in the block. In contrast, a unique attribute may be specified for each element in a block.

Currently, four groups of elements have parameters. Shells have membrane and bending factor factors 5.7.5. The Infinite Element 6.1.9 and Perfectly Matched Layer 6.1.10 are configured using several unique block parameters. And the Beam2 5.9 and Nbeam 5.10 both have several parameters.

An example is provided in input 4.7. The block names can be provided in three forms as shown in the example. An integer number can be given, the number refers to the **Exodus** block id number in the mesh. Alternatively this number be provided as 'BLOCK_##' which is compatible with **Sierra/SM** syntax. Also, an option the block name can be given such as 'aft_cover_plate', again this is the name of the block in the specified in the **Exodus** input mesh. Note that the material ID specified for BLOCK 32 uses an index (material 2), whereas BLOCK aft_cover_plate uses a specified material ID string "aluminum" These refer to materials defined by blocks "MATERIAL 32 ... END" and "MATERIAL aluminum ... END" respectively (see Sec. 4 for details).

```
BLOCK 32
    material 2
    tria3
    thickness 0.01
END

BLOCK aft_cover_plate
    material aluminum
END

BLOCK block_3
    coordinate 1
    spring
    Kx=1e6
    Ky=0
    Kz=0
    BlkBeta=0.0031
END

Material aluminum
    ....
END
```

Input 4.7. Example Block input

A list of the applicable attributes⁵² for different element types is shown in Table 4-107. Each element type is outlined in Section 5.

Table 4-107. – Element Attributes.

Element Type		keyword	Description
ConMass	1	Mass	concentrated mass
	2	Ixx	xx moment of inertia
	3	Iyy	yy moment of inertia
	4	Izz	zz moment of inertia
	5	Ixy	xy moment of inertia
	6	Ixz	xz moment of inertia
	7	Iyz	yz moment of inertia
	8,9,10	offset	offset from node to CG
Beam	1	Area	Area of beam
	2	I1	First bending moment
	3	I2	Second bending moment
	4	J	Torsion moment
	5,6,7	Orientation	orientation vector. For the orthogonal direction
	8,9,10	offset	beam offset
Spring	1	Kx	spring constant in X
	2	Ky	spring constant in Y
	3	Kz	spring constant in Z
Triangle	1	thickness	thickness
	2	fiber orientation (theta)	fiber orientation
	3	offset	shell offset in normal direction
Quad	1	thickness	thickness
	2	fiber orientation (theta)	fiber orientation
	3	offset	shell offset in normal direction

4.7.2. General Block Parameters

Parameters that are generally applicable to almost all blocks are listed in Table 4-108. More detailed descriptions are available in the following paragraphs.

4.7.2.1. Nonlinear Behavior The nonlinear behavior of the block in nonlinear solutions is controlled by the **nonlinear** keyword. The global default for block-level nonlinear behavior is set in the **parameters** section (2.3). Within each block, we can override that default value. For example, to set a block to default to linear behavior, we would have the following BLOCK definition.

```
BLOCK 3
  nonlinear=no
  material 2
  tria3
  thickness 0.01
END
```

Similarly, to turn on the nonlinear behavior for the block, we would have,

```
BLOCK 3
  nonlinear=yes
  material 2
  tria3
  thickness 0.01
END
```

Note that these block-level nonlinear flags override the global **nonlinear_default** keyword that is set in the **parameters** section (2.3).

Some elements types are incapable of nonlinear behavior. This includes RBE3s, Rbars, SUPERELEMENTs, and Rrods. By default, use of these elements in nonlinear analysis will generate a fatal error. Use of the command **nonlinear=no** in these element blocks will enable overriding this fatal error and use pure linear behavior for these elements in the nonlinear analysis. Care should be taken with this option, pure linear elements have some incompatibilities in nonlinear analysis. For example use of pure-linear elements in nonlinear analysis can artificially constraint large rotations.

Limitations:

Linear element behavior in a nonlinear solution is limited to the linear range of the element. For example, rotations are stored incrementally in nonlinear solutions. This permits us to use geometrically nonlinear element formulations (such a corotational formulation). However, it limits the *linear* behavior in such solutions to rotations less than 360°.

Table 4-108. – General Block Parameters.

Keyword	Values	Description
nonlinear	yes/no	blockwise nonlinear behavior
material	<i>string</i>	material identifier
rotational_type	Eulerian or Lagrangian or none	blockwise behavior for rotational dynamics terms
coordinate	<i>string</i>	reference coordinate frame
blkalpha	<i>Real</i>	blockwise mass proportional damping
blkbeta	<i>Real</i>	blockwise stiffness proportional damping
nsm	<i>Real</i>	blockwise non-structural mass
density_scale_factor	<i>Real</i>	blockwise density scaling factor
stiffness_scale_factor	<i>Real</i>	blockwise stiffness scaling factor
t_current	<i>Real</i>	Temperature to set on every element of the block.

4.7.2.2. Rotational Loading Matrices For problems involving rotational loads, the **rotational_type** keyword allows the analyst to specify which type of rotational formulation to use for a given block of elements. The Eulerian formulation involves a fixed (non-rotating) coordinate system. The Lagrangian formulation attaches a rotating coordinate system to the block. If the None options is chosen, then rotational loads are ignored for this block. Thus, a structure with a rotating disk would only have the rotational terms applied to the spinning disk, and not the entire structure. The default is for the **rotational_type** keyword is None.

4.7.2.3. Coordinate Frame Reference The reference coordinate system may be defined in a block. This definition applies to all the elements of the block and the associated materials. At this point, the coordinate system is only recognized for a subset of the elements (solid elements and springs). Further information on coordinate systems may be found in Section 2.7.

4.7.2.4. Block Specific Damping In Section 4.8, various methods of specifying the damping parameters for a model are identified. In addition to these methods, block specific damping parameters may be applied. These apply a stiffness (or mass) proportional damping matrix on an element by element basis within the block. Thus, if a model is made of steel and foam, one could apply a 5% stiffness proportional damping term to the foam, but leave the steel undamped.

There is no physical justification for proportional damping, and there is no expectation that it will accurately represent damping mechanisms in a structure. However, it is easy to apply, and there are cases where proportional damping may reveal a need for more

accurate damping models. As with all damping models, the effects depend on the solution type. For example, both Statics and Eigen analysis ignore the damping matrix.

The damping matrix generated from block specific damping is defined as follows.

$$D = \sum_i^{nblks} \alpha_i M_i + \beta_i K_i \quad (4.12)$$

Where D is the real system damping matrix, and α_i and β_i are the proportional mass and damping coefficients for block i . These coefficients are completely analogous to the system level coefficients described in Section 4.8. The damping contributions from these block parameters are always added to the other contributions.

Block specific damping is applied using the **blkalpha** and **blkbeta** parameters. Block proportional damping generates a damping matrix that would couple modal based solutions. It is not currently available in modal solutions such as **modaltransient**. Also, see Section 4.4.9 for material modal like damping.

4.7.2.5. Non-Structural Mass Non-structural mass (NSM) is specified per element block in the input deck. It is added to the internal mass of the element. One reason to add a NSM is to use a gravity load to simulate an external load. Another reason is to stabilize solutions with mass-less nodes. The units depend on the element type as defined in Table 4-109.

4.7.2.6. Non-Structural Mass Corner Cases "Three Dimensional" is determined from the mesh topology. HexShell elements are three-dimensional. Layered shell elements add non-structural mass once per meshed element, not once per layer. Some conditions can cause non-structural mass to be silently ignored. These cases do not trigger a warning, and are shown in Table 4-110.

The following is an example of how to use non-structural mass in the input file:

```
//nsm specified in pounds per square inch
BLOCK 3
    material 2
    tria3
    thickness 0.01
    nsm 0.005
END

MATERIAL 2
    density 0.5
END
```

Table 4-109. – Non-Structural Mass Units.

Element Type	Units	Example
ConMass	Mass Per Element	lbs
Spring or Joint2g	Mass Per Element	lbs
Beam or Truss	mass/length	lbs / in
Two Dimensional	mass/area	lbs / sq-in
Three Dimensional	mass/volume	lbs / cu-in

Table 4-110. – Unhandled Corner Cases.

Element Type	State	Result
Beam or Truss	Area=0	NSM Silently Ignored
Two Dimensional	Thickness=0	NSM Silently Ignored

4.7.2.7. Blockwise Density Scaling An element block may define a scale factor to be applied to the density of the material. This can be used to calibrate the exact mass in each block and account for discretization errors, or to reuse materials that only differ in density. The interaction of this feature with non-structural mass is documented in [Table 4-111](#).

The following is an example of blockwise density scaling in the input file:

```
BLOCK 3
  material 2
  density_scale_factor = 1.0025
END

MATERIAL 2
  density 0.5
END
```

Table 4-111. – Combining NSM with Density_Scale_Factor.

Element Type	Mass Per Element
ConMass	NSM + Mass
Spring or Joint2g	NSM
Beam or Truss	NSM*Length + Density*Scale*Volume
Two Dimensional	NSM*Area + Density*Scale*Volume
Three Dimensional	NSM*Volume + Density*Scale*Volume

4.7.2.8. Blockwise Stiffness Scaling An element block may define a scale factor to be applied to the linear stiffness of the material, but the capability has been limited to isotropic material models only. This can be used to tune the precise stiffness of components without requiring separate material definitions, usually when adapting a model to match test results. The `stiffness_scale_factor` is applied to isotropic materials consistently, even if the Young's modulus is not defined explicitly by the user. Any valid combination of material constants will still be valid.

Note that `stiffness_scale_factor` is ignored on blocks without a material, such as spring elements, and attempting to use `stiffness_scale_factor` with any anisotropic material will result in an error.

The following is an example of blockwise stiffness scaling in the input file:

```
BLOCK 3
    material 2
    stiffness_scale_factor = 1.0025
END

MATERIAL 2
    isotropic
    E 1e7
    nu 0.3
    density 0.5
END
```

4.7.2.9. Piezoelectric Material Damping Only stiffness (`blkbeta`) and mass (`blkalpha`) proportional damping can be applied to electro-mechanical materials, and damping models can only be specified at the block level. Global damping is prohibited on any model containing piezoelectric or dielectric materials. Voltage degrees of freedom do not couple with mass or damping (see theory manual). Hence, a piezoelectric element's damping matrix, defined by stiffness and(or) mass proportional damping, is zero at all voltage degrees of freedom. The effects of mechanical damping will only impact the voltage degree of freedom responses due to the electro-mechanical stiffness coupling.

4.8. *Damping*

This section allows input of simple global viscous damping models, using either modal damping rates or stiffness and mass proportional damping. The various options for the DAMPING section are shown in Table 4-112.

The damping matrix or modal damping coefficient is determined by summing contributions from all damping parameters given in Table 4-112. For modal superposition-based

Table 4-112. – DAMPING Section Options.

Parameter	Description
alpha	mass proportional damping parameter (real)
beta	stiffness proportional damping parameter (real)
gamma	uniform modal damping ratio (fraction of critical) applied to all modes (real)
mode	ADDITIONAL modal damping ratio applied to individual mode(fraction of critical) (integer, real)
ratiofun	index of function to define modal damping ratios
FilterRbm	remove rigid body mode contribution to damping
maxRatioFlexibleRbm	controls check for 6 RBM with FilterRbm

transient analysis, **modaltransient** , all the given parameters are defined. For linear direct implicit transient analysis, the modal damping parameters apply only to modes for which eigenvalues and eigenvectors have previously been computed. This depends on the presence of the keyword **nmodes** in the solution section of the input file. In the case of a **modalranvib** (or **ModalFrF** analysis in the case of complex modes), modal damping is available, but the proportional damping parameters **alpha** and **beta** are currently ignored.

The effect of the mass and stiffness proportional parameters on modal damping depends on the frequencies of the modes. For modal-based analysis, the damping rate for mode i with radial frequency ω_i is given as

$$\zeta_i = \alpha/(2\omega_i) + \beta \cdot \omega_i/2 + \Gamma + \text{mode}[i] + \text{ratiofun}(i),$$

where the viscous damping term in the modal equilibrium equation is $2\zeta_i\omega_i$. For example the following damping input section could be used in a modal transient analysis. ²

```
DAMPING
  alpha  0.001    //
  beta  0.00005   //  C = .001 * M + .00005 * K
  gamma  0.005    //  0.5% critical
  mode   1  0.01  //  gamma+mode_1 = 1.5% of critical
  mode   2  0.005 //  gamma+mode_2 = 1.0% of critical
  mode   3  0.015 //  gamma+mode_3 = 2.0% of critical
END
```

It produces the following damping ratios.

²Use of block specific proportional damping is explained in Section 4.7.2.

Mode	modal damping ratio	modal viscous damping term
1	$0.015 + 0.001/(2\omega_1) + 0.00005\omega_1/2$	$0.030\omega_1 + 0.001 + 0.00005\omega_1^2$
2	$0.010 + 0.001/(2\omega_2) + 0.00005\omega_2/2$	$0.020\omega_2 + 0.001 + 0.00005\omega_2^2$
3	$0.020 + 0.001/(2\omega_3) + 0.00005\omega_3/2$	$0.040\omega_3 + 0.001 + 0.00005\omega_3^2$

In direct transient analysis ³, the full mass and stiffness matrices are integrated for the solution. Specification of a modal damping method triggers construction of a damping contribution³ from the previous modal solution. This contribution is combined with other damping terms such as the proportional damping. Thus, the same damping input section would produce the damping ratios shown above for the first three modes. Modal damping is applied to modes computed in a previous solution case. ⁴

The **ratiofun** keyword permits definition of modal damping terms based on a frequency dependent function. The associated function definition (see Section 2.8) provides a table look up for damping ratios. For example, consider a system with modes at 200 and 500 Hz. The following example will establish modal damping ratios of .03 and .06 respectively. The function describes a line defined by $ratio(f) = 0.01 + 0.1/1000f$.

```
DAMPING
  ratiofun=100
END

FUNCTION 100
  type=linear
  data 0 0.01
  data 1000 0.11
END
```

The **FilterRbm** keyword permits proportional damping without damping the rigid body response. Thus, mass proportional damping can be used with no impact on the rigid body response. The theory behind this method of damping is described in subsection Damping of Flexible Modes Only section Solution Procedures of the Theory Manual.

In order for this method of damping to work properly, the structure must have the conventional six rigid body modes of three translations and three rotations. A check of this condition is made inside of **Sierra/SD**, and a fatal error results if this condition is not satisfied. Specifically, the condition is met if

$$RatioFlexibleRbm = \frac{\|K\Phi_r\|_2}{\|K_d\|_\infty \|\Phi_r\|_2} \leq \epsilon \quad (4.13)$$

³i.e. non-modal based, but linear transient

⁴A previous modal solution case must have been specified to use modal damping, otherwise **Sierra/SD** will warn the user and abort.

where K is the stiffness matrix, Φ_r is the matrix of six rigid body modes, and $\|K_d\|_\infty$ is the largest entry on the diagonal of K . The scalar tolerance ϵ can be specified using the **maxRatioFlexibleRbm** keyword.

```
DAMPING
  alpha=0.1
  FilterRbm
  maxRatioFlexibleRbm=0.001 // default is 1e-10
END
```

4.8.1. Nonlinear transient solutions with damping

Using the stiffness proportional damping parameter **beta** in a **NLtransient** analysis will generate damping terms using the initial (or linear) stiffness matrix. The tangent stiffness matrix is not used. This reason is that the tangent matrix would be required to compute the damping terms at each iteration.

Nonlinear solutions do not support standard modal damping.

While nonlinear solutions do not currently support standard modal damping, they may be damping using the Distributed Damping method of the next section (4.8.2). Like modal damping, this is a system level damping model.

4.8.2. Nonlinear Distributed Damping using Modal Masing Formulation

The purpose of this formulation is to implement a subsystem or system level nonlinear distributed damping model into **Sierra/SD**. The theory on this method is found in the **Sierra/SD** Theory Manual.⁵² Distributed damping is a method developed to model the nonlinear damping response of a subsystem. It implements the damping in a nonlinear manner with the use of an internal force term. The damping is modeled by either an Iwan model or a linear damper, and distributed to the subsystem by a modal expansion. This method augments the internal force vector through a modal Masing formulation.²

Previous to the nonlinear transient solution which computes the distributed damping, eigenvectors must be computed. This is done in a previous solution 'case' option using "eigen" methods.

The damping section is used to define the type of damping behavior. Currently, only two types of damping behavior are defined: a linear damper, **damper**, and an Iwan model, **Iwan**.⁵² Each mode will have a keyword defined after it with an associated parameter

²Masing and Iwan models are used almost interchangeably in this document. Iwan models are a subset of more general Masing models.

number. The parameters are used to define the damping behavior. If nothing is specified for a mode, then no damping for that mode is defined. An example input is shown below.

```
SOLUTION
case 'eig'
  eigen
    nmodes 16
    shift -1e5
case 'nonlinear'
  NLtransient
    nsteps = 200
    time_step = 5.0e-3
    rho = 0.8
END
```

```
DAMPING
  mode 1 damper 1
  mode 2 damper 2
  mode 3 damper 2
  mode 4 damper 2
  mode 5 damper 2
  mode 6 damper 2
  mode 7 Iwan 4
  mode 8 Iwan 4
  mode 9 Iwan 4
  mode 10 Iwan 3
  mode 11 Iwan 3
  mode 12 Iwan 3
  mode 13 Iwan 3
  mode 14 Iwan 3
  mode 15 Iwan 3
  mode 16 Iwan 3
END
```

```
Property 1
  Mu = 0.001
  K = 0
end
```

```
Property 2
  Mu = 0.02
  K = 0
end
```

```
Property 3
  chi = -0.82139
  phi_max = 1.0325e-04
  R = 7.608594e+06
  S = 5.616950e+06
end

Property 4
  chi = -0.82139
  phi_max = 1.0325e-04
  R = 7.608594e+06
  S = 5.616950e+06
end
```

5. Element Library

Short descriptions of each of the types of elements follow. Materials are described in the materials section. Some elements have attributes. An attribute may be specified in the **Exodus** file per element or the input deck for per block or both. Input deck values override values from the **Exodus** file. Attributes are described in this section. An attribute is either required or optional.

5.1. Hex8

The **Hex8** is a standard 8 node hexahedron with three degrees of freedom per node. The **Hex8** element has 8 integration points. The shape functions are trilinear. It supports isotropic and anisotropic materials.

There are three variations of **Hex8** . The default element is a bubble hex element, which can be specified by **Hex8b** , or by no specification at all. The bubble element still has 8 nodes and 3 degrees of freedom per node, and thus from a user's perspective it is indistinguishable from the standard **Hex8** . The **Hex8b** element uses bubble functions^{67,36,43} to augment the standard element shape functions. It bends more accurately than the hex8.

The **Hex8u** specifies an under integrated Hex with properties similar to those of most commercial finite element codes. There are two versions of this element. The first is a standard under integrated element, and the second is a mean quadrature element with selective deviatoric control. In both cases, the under integration produces an element that is soft relative to a fully integrated element. Both elements are specified using by the keyword **Hex8u** , along with the parameter **sd_factor** . If **sd_factor** is specified, then the mean quadrature element with selective deviatoric is invoked with the value of **sd_factor** specified. If the parameter **sd_factor** is not specified, the standard under integrated element is invoked. For example, to use the under integrated element, one could specify

```
block
  hex8u
  material 1
end
```

On the other hand, the following block would use the mean quadrature element with a selective deviatoric parameter of 0.9

```
block
  hex8u
  material 1
  sd_factor 0.9
end
```

Note that **sd_factor** must be between 0 and 1. With a value of 0, the element is a mean quadrature element. With a value of 1, the element is again mean quadrature, but with fully integrated deviatoric component. More information is given in subsection Isoparametric Solid Elements. Selective Integration section **Sierra/SD** Elements of the Theory Manual.

The fully integrated Hex is specified by **Hex8F**. While it performs adequately when the element shape is nearly cubic, it performs poorly for larger aspect ratios. For most problems involving bending the Hex8b is recommended.

The only *required* parameter for these elements is the material specification. Any material may be applied.

For computational acoustics, see Section [2.4.4](#).

5.2. **Hex20**

The 20 node variety of Hex element provides quadratic shape functions. It is a far better element than the Hex8, and should be used if possible. The Hex20 element in **Sierra/SD** is similar to elements found in most commercial codes. A material specification is required, and any structural material may be used.

Shape Function and Gauss point locations for the Hex20 are described in table [7-143](#), and in subsubsection Shape Functions and Gauss Points subsection Quadratic Isoparametric Solid Elements section **Sierra/SD** Elements of the Theory Manual.

The stress may be output at the Gauss points as described in Section [7.1.24](#).

5.3. **Wedge6**

The Wedge6 is a compatibility element for the **Hex8**, it is not recommended that the entire mesh be built of **Wedge6** elements. They are primarily intended for applications where triangles are naturally generated in mesh generation. A material specification is required, and any structural material may be used.

5.4. *Wedge15*

The Wedge15 element adds mid-side nodes to the Wedge6. Like the Hex20 and Tet10, it has quadratic shape functions, and is recommended. A material specification is required, and any structural material may be used.

5.5. *Tet4*

This is a standard 4 node tetrahedral element with three degrees of freedom per node. The **Tet4** element has one integration point. The shape functions are linear. It is not recommended to use only Tet4 elements for the entire mesh because standard, linear tetrahedron are typically much too stiff for structural applications. The **Tet4** is provided primarily for those applications where a mesh may be partially filled with these elements. If a model is constructed of all tetrahedral elements (as by an automatic mesh generator), the **Tet10** is strongly recommended over the **Tet4**.

A material specification is required, and any structural material may be used.

5.6. *Tet10*

This is a standard 10 node tetrahedral element with three degrees of freedom per node. The **Tet10** uses 4-point integration for the stiffness matrix and 16-point integration for the mass matrix. The shape functions are quadratic. This element is recommended for use in most structural analyses.


A material specification is required, and any structural material may be used.

5.7. *Two-Dimensional Shell and Membrane Elements*

Sierra/SD supports a variety of topologically 2D elements that capture shell or membrane behavior. A specific 2D element formulation can be selected by explicitly specifying the element type in the input block with keywords such as **quadt**, **quadm**, **nquad**, etc. For three noded triangles if no specific element formulation is given then by default one of the Tria3 (isotropic materials) or TriaShell (orthotropic or layered materials) shell formulations will be used. For all other topologies (4-node quad, 8-node quad, and 6-node triangle) the default element is composed of sub-triangles using the Tria3 or TriaShell shell formulations as described in [Section 5.7.1](#)

Figure 5-34. – QuadT Element.

The element is generated by internally combining two Triangle elements.



5.7.1. QuadT, Quad8T, and Tria6

The 4-node quad **QuadT** , 8-node quad **Quad8T** , and 6-node triangle **Tria6** are all internally composed of sub-triangles. Each of these elements is the **Sierra/SD** default formulation for its respective topology. These elements have both membrane and bending stiffness. The element stiffness and mass matrices are derived by composing internally generated triangle elements, as illustrated in Figure 5-34, 5-35, and 5-36. Output quantities such as stress are the average stress over the sub-triangles. Composing elements from low-order triangles, though not optimal, is adequate for most applications.

The sub-triangles may be based on either the **Tria3** , or on the **TriaShell** element depending on the material properties. The **Tria3** is used for isotropic, single-layer elements. More complex materials require use of the **TriaShell**. The underlying triangle formulation is determined automatically by **Sierra/SD**, and cannot be selected by the user. See the description of the **Tria3** and **TriaShell** for details of the formulations of the triangle elements that compose the **QuadT** , **Quad8T** , and **Tria6** elements.

Table 5-113 lists the supported inputs.

Sierra/SD example input files that use this element can be found in

```
Salinas_rtest/patchtests/quadt/quadt-patch8_test  
Salinas_rtest/patchtests/quadt/quadt-patch9_test
```


Figure 5-35. – Quad8T Element.





Figure 5-36. – Tria6 Element.



Keyword	Description
thickness	Thickness of the shell
offset	offset of the shell midplane: see Section 5.7.9
material	Linear elastic material used by element
layer	Layer properties: see Section 5.7.8
coordinate	Base global coordinate system: see Section 5.7.7
rotate about axis	Coordinate system rotation: see Section 5.7.7
rotate about normal	Coordinate system rotation: see Section 5.7.7
membrane_factor	Stiffness scale factor: see Section 5.7.6
bending_factor	Stiffness scale factor: see Section 5.7.6

Table 5-113. – QuadT, Quad8T, Tria6 Inputs.

5.7.2. QuadM

QuadM is a 4-node quadrilateral membrane element. It has only membrane stiffness and no rotational degrees of freedom. Membranes are well suited to structures with very low bending stiffness, such as fabric. Use of shell elements for such low-bending-stiffness structures can generate a problematic near-singularity.

In the input deck a **block** section indicating **QuadM** is required.

For two-dimensional problems, the **QuadM** reduces to the standard plane elasticity element. For three-dimensional problems, it behaves like the plane elasticity element in the plane, and like a stretched balloon out-of-plane. The out-of-plane stiffness for the element comes from the preload. Without a preload, the out-of-plane stiffness vanishes and the element can be singular. The out-of-plane behavior results from an additional stiffness term that is applied to the out-of-plane degrees of freedom. The stiffness resembles the stiffness associated with Laplace’s equation. This additional stiffness is derived in classical textbooks.³⁹

Table 5-114 lists the supported inputs.

Keyword	Description
thickness	Thickness of the membrane, required
sd_factor	Selective deviatoric parameter used for numerical integration
material	Linear elastic material used by element, required
coordinate	Base global coordinate system: see Section 5.7.7
rotate about axis	Coordinate system rotation: see Section 5.7.7
rotate about normal	Coordinate system rotation: see Section 5.7.7
membrane_factor	Stiffness scale factor: see Section 5.7.6

Table 5-114. – QuadM inputs.

Both full and selective integration methods are available for the membrane. The full integration is the default. Selective deviatoric integration can be specified by using the

parameter **sd_factor** . For example, for full integrated membrane, one would specify

```
block
  QuadM
  material 1
  thickness 0.1
end
```

On the other hand, the following block would use the mean quadrature element with a selective deviatoric parameter of 0.9

```
block
  QuadM
  material 1
  sd_factor 0.9
  thickness 0.1
end
```

Note that **sd_factor** must be between 0 and 1. With a value of 0, the element is a mean quadrature element. With a value of 1, the element is again mean quadrature, but with fully integrated deviatoric component. More details on the theory behind these elements is given in the theory manual.

This element could be used in any situation where a preload is applied to the elements before the analysis of interest (e.g., a static preload followed by eigen analysis), or even in cases where no preload is applied but the membranes are sufficiently constrained (such as a hex element with a layer of membrane elements on the surface).

The QuadM element can be used in coupled simulations. In these cases, Adagio performs the preload calculation, and the preload information is passed to **Sierra/SD** for later analysis. In these cases the element is non-singular due to the preload calculation.

5.7.2.1. Known Issues

- The membrane element does not currently compute stress, strain, strain energy, or strain energy density outputs. All these outputs will be reported as zero for the element.

5.7.3. Nquad/Ntria

The **Nquad** and **Ntria** elements are isoparametric shells with membrane and bending stiffness. They are shear-deformable elements with six degrees of freedom (DOF) per node which support isotropic, orthotropic, and layered materials. The formulation of the **Nquad/Ntria** is generated by decoupling the membrane and bending DOF. These

elements currently only have linear behavior implemented. If using a non-linear solution method, these elements will not calculate a true internal force, but a linear force.

The **Nquad**/**Ntria** isotropic stiffness matrix is based on the plane elasticity and shear deformable (Mindlin) formulations as outlined in Reddy.⁴⁹ The layered shell stiffness matrix uses a composite laminate formulation.⁴⁶

In the input deck a **block** definition indicating either **Nquad** or **Ntria** is required. The block definition must also have a material keyword referencing the isotropic material properties (Section 4) or orthotropic layer properties (Section 4.1.2) with properties E_1 , E_2 , ν_{12} , and G_{12}). An example element block for a single layer isotropic material is shown below:

```
block 2
  Nquad
  thickness 0.1
  material 2
end
block 3
  Ntria
  thickness 0.4
  material 4
end
```

Inputs are given in Table 5-115.

Keyword	Description
thickness	Thickness of a single layer shell
material	Linear elastic material used by single layer element
layer	Layer properties: see Section 5.7.8
nquad_eps_max	Numerical integration parameter
coordinate	Base global coordinate system: see Section 5.7.7
rotate about axis	Coordinate system rotation: see Section 5.7.7
rotate about normal	Coordinate system rotation: see Section 5.7.7

Table 5-115. – Nquad/Ntri inputs.

The stabilization method from Belytschko⁹ is used for the **Nquad** element. Using single-point integration $K_s^{[1x1]}$ for the shear stiffness matrix leads to hourglass modes for some problems. Using full integration $K_s^{[2x2]}$ can cause shear locking in some problems. Belytschko recommends a shear stiffness matrix that is a linear combination of the reduced integration and full integration shear stiffness matrices,

$$K_s = (1 - \varepsilon)K_s^{[1x1]} + \varepsilon K_s^{[2x2]}.$$

The fraction, $\varepsilon = rt^2/A$ is a function of thickness and area. Here $r = 0.03$, t is the element thickness and A is the area of the shell. This automatic selection of ε works well for thin

plates, but can be a problem for thicker elements; ε should never exceed 1. To limit shear locking, the fraction may be capped using `nquad_eps_max`, as shown in the example below.

```
block 1
  nquad
  thickness 1
  nquad_eps_max 0.1
end
```

The value for ε is adjusted using the function $\hat{\varepsilon} = \frac{\varepsilon_{max}}{\sqrt[4]{1+\varepsilon^4}}$. This is done to address problems with “elbow functions” in the code. Figure 5-37 shows this function for `nquad_eps_max` = 1.




Figure 5-37. – Function for `nquad_eps_max`.

5.7.3.1. Known Issues

- The `Ntria` with orthotropic materials sets G_{23} and G_{13} to the input value of G_{12} .

5.7.4. TriaShell

The **TriaShell** is a 3-noded triangular element with 6 degrees of freedom (DOF) per node. The formulation of the **TriaShell** is generated by decoupling the membrane DOF and the bending DOF. Allman’s Triangular (AT) element² models the membrane DOF, while the Discrete Kirchhoff Triangle⁸ (DKT) models the bending DOF. These two elements are combined into the **TriaShell** element.

In general, the `Tria3` element is preferred to the **TriaShell** because it is less prone to shear locking behavior, and it is computationally efficient. The **TriaShell** element is required for orthotropic or layered materials.

- **TriaShells** support orthotropic or anisotropic materials.

Keyword	Description
thickness	Thickness of the shell
offset	Offset of the shell midplane: see Section 5.7.9
material	Linear elastic material used by single layer element
layer	Layer properties: see Section 5.7.8
coordinate	Base global coordinate system: see Section 5.7.7
rotate about axis	Coordinate system rotation: see Section 5.7.7
rotate about normal	Coordinate system rotation: see Section 5.7.7
membrane_factor	Stiffness scale factor: see Section 5.7.6
bending_factor	Stiffness scale factor: see Section 5.7.6

Table 5-116. – TriaShell input options.

- TriaShells support layered materials. Note however that mass lumping is not allowed with layered TriaShell elements.

5.7.5. Tria3

The **Tria3** is a three-dimensional triangular shell with membrane and bending stiffness. There are 6 degrees of freedom per node. In most respects it is similar to the **TriaShell**. It is the default element for triangular meshes. The **Tria3** was provided by Carlos Felippa of UC Boulder. It only supports only isotropic single-layer materials.

Keyword	Description
thickness	Thickness of the shell, required
offset	offset of the shell midplane see Section 5.7.9
material	must be isotropic, required
membrane_factor	Stiffness scale factor, see Section 5.7.6
bending_factor	Stiffness scale factor, see Section 5.7.6

An example element block is shown below.

```

block 3
  Tria3
  Thickness 0.01
  material 71
  membrane_factor=0 // turns off membrane stiffness
end

```

5.7.6. Stiffness Scaling

The stiffness of the 2D element bending and membrane responses are computed independently and can be independently arbitrary scaled using the commands **membrane_factor** and **bending_factor**. Each of these parameters default to 1.0. Reasons for scaling the element stiffness could include accounting for damage to the structure or tuning of model response to match experimental data.

5.7.7. Shell Coordinate Systems

For orthotropic materials a coordinate system must be defined in the element to set the local material orientation. First, a user-defined global coordinate system is referenced with the **coordinate** keyword (Section 2.7.) The global coordinate system is evaluated at the element centroid to define a local $\hat{r}, \hat{s}, \hat{t}$ system for the element.

This $\hat{r}, \hat{s}, \hat{t}$ coordinate system is projected onto the surface of the shell as shown in Figure 5-38. This projection will rotate the coordinate system such that the $\hat{t}_p$ axis aligns with the shell normal and $\hat{r}_p$ and $\hat{s}_p$ are as close as possible to the original orientations of $\hat{r}$ and $\hat{s}$. It is recommended to use a global coordinate system in which $\hat{t}$ lies as close as possible to the shell normal vectors to minimize ambiguity in this coordinate system alignment step.

Two additional inputs are available in the shell block to alter the element local coordinate system. First a rotation can be applied to the $\hat{r}, \hat{s}, \hat{t}$ vectors prior to projection onto the shell plane with the command:

```
rotate <real> about axis <int>
```

The rotation angle is given in degrees. The axis integer 1, 2, or 3, represents the $\hat{r}$, $\hat{s}$, or $\hat{t}$ coordinate axes. After rotation of the coordinate system the $\hat{r}', \hat{s}', \hat{t}'$ system is aligned to the plane of the shell. See Figure 5-39 for an example.

A second input option can rotate the the element local coordinate system after projection with the command:

```
rotate <real> about normal
```

The angle is given in degrees. This will rotate the projected system around the $\hat{t}_p$ vector as shown in Figure 5-40.

Once a shell-local system is defined, the $\hat{r}_p$ and $\hat{s}_p$ vectors will define the orientation of orthotropic materials. The **fiber_orientation** command can do one more local rotation of this orientation layer-by-layer as described in Section 5.7.8.




Figure 5-38. – Projection of global coordinate system to shell.




Figure 5-39. – Rotation of global coordinate system about axis prior to projection to shell.




Figure 5-40. – Rotation of local system about normal after to projection to shell.

5.7.8. Layered Shells

Several of the shell element formulations allow composing the element via a stack of layers of different materials. This is used to model layered composites. When using layers, the available materials are *isotropic* and *orthotropic_layer*. Each layer must specify a material and thickness. A fiber orientation, which is a rotation of the layer material coordinate system with respect to the element coordinate system, may optionally be given. Thickness and fiber orientation for a multi-layer material must be specified layer by layer in the input deck. **Exodus** attributes may not be used.

Keyword **layer** defines a new layer for the current shell. Layers of the shell are stacked from the bottom to the top based on the order of the **layer** keyword in the input deck. The **layer_ID** input is an identifier provided by the user and is not used to select stacking order. A shell may have up to 250 different layers defined. Figure 5-41 shows a simple schematic explaining how layers are stacked in **Sierra/SD**. An example element block for a four-layer orthotropic layered shell is shown below. ³

An important parameter for the layered shells with orthotropic materials is the specification of a user-defined coordinate system with the **coordinate** option, see Section 5.7.7 In the example shown here, a coordinate system is defined for the shell block and orthotropic material properties are defined via a fiber orientation rotation with respect to that base element system.

³For layered shells, the **thickness** parameter specifies the actual thickness of that layer of the shell. This is in contrast to the HexShell which specifies a relative thickness: see Section 5.8.




Figure 5-41. – Stacking arrangement for a multi layer shell element.

```

block 2
  TriaShell
  coordinate 1
  layer <layer_ID>
    material 1
    thickness 0.02
    fiber orientation 40
  layer <layer_ID>
    material 2
    thickness 0.04
    fiber orientation 44
  layer <layer_ID>
    material 3
    thickness 0.03
    fiber orientation 54
  layer <layer_ID>
    material 4
    thickness 0.01
    fiber orientation 4
end

begin rectangular coordinate system 1
  origin 0.0 1.0 1.0
  z point 2.0 1.0 1.0

```

```
xz point 0.0 1.0 10.0
end
```

Input 5.1. Layered Shell Example

Stress output for shells with more than one layer can be written to an **Exodus** file or can be obtained from the result file by specifying stress in the **echo** section. The layer stresses will be computed only at the midpoint of each layer. Thus, layer stresses at the top and bottom of each layer are not supported.

5.7.8.1. Known Issues

- The Navy Layered Shell with orthotropic materials sets G_{23} and G_{13} to the input value of G_{12} .

5.7.9. Offset Shells

By default, the meshed shell element lies at the midplane of the material volume represented by the shell. Shells may be offset from the midplane by specifying the offset input. The offset vector is the element unit normal vector scaled by the offset. An example is shown in Figure 5-42.

The resulting mass and stiffness properties are equivalent to the stiffness generated by translating the shell by the offset vector, and constraining the resulting offset nodes to the untranslated nodes using rigid links. The performance of offset shells is better than that of the constraint approach. Note that for curved surfaces there may be modeling issues with offset elements since there is no change in curvature with the change in radius.

In the `.inp` file the element offset is specified as,

```
offset=-3.14e-2
```

Offsets may also be specified in the **Exodus** file via attributes (see Section 5.7.10.) Some limitations of element offsets are described in Section 5.7.9.1

5.7.9.1. Offset Shells and Lumped Mass

Offset elements can provide more accurate modeling of some structures, such as shells cladded on top of a volume. Offset elements necessarily involve coupling between the rotational and translational degrees of freedom. This results in off-diagonal coupling terms in the element stiffness and mass matrices.

Generally, the element stiffness matrix is fully populated and seldom is reduced. However, the mass matrix may be lumped (diagonalized) as described in Section 2.4.4.




Figure 5-42. – Offset examples for shell of thickness 0.1.

Lumping the mass matrix uncouples the translational and rotational degrees of freedom, which is inaccurate for offset shells. Specifically, while the total mass is conserved, the center of gravity and mass moments are not. The lumped mass looks as if it had not been offset. This is true even with mesh refinement. The models of the consistent and lumped mass are fundamentally different when element offsets are included. Mass lumping with offset shell elements is discouraged.

5.7.10. Spatially Dependent Properties via Exodus Attributes

Certain 2D element properties can be defined either in the input deck or via attributes on the input **Exodus** mesh. When a property is defined in the input deck it has a constant value for all elements of the block. The advantage of **Exodus** attributes is that a different value can be used in each element. This can be used to model properties such as variations in thickness in tapered shells.

The supported attributes are shown in Table 5-117.

Attribute Index	Keyword	Description
1	thickness	Thickness of the shell
2	fiber_orientation	Rotation of material with respect to local coordinate system
3	offset	Offset of the shell midplane with respect to the meshed surface: see Section 5.7.9

Table 5-117. – Shell parameters that can be set via attributes.

If an input deck value is given for a shell property, such as thickness, it will override any value given in the **Exodus** attributes. Attributes cannot be used with multi-layer shells. For Multi-layer shells all the element properties must be defined in input deck.

5.8. *HexShell*

The 8 noded HexShell is a hybrid solid/shell element. It is meshed as a standard hex element, but the formulation of the element is similar to that of a shell. Unlike a shell element, the thickness is determined by the mesh. But, the element is designed to operate with many of the same features as shell elements even when it becomes thin. Details of the element formulation are available in a separate report.²⁸ An introduction to HexShells is readily available in,⁵² and the verification manual⁵⁴ discusses the results of the verification problems from²⁸ for **Sierra/SD**.

The HexShell has a preferential thickness direction that must be set correctly. There are three ways to specify the thickness direction.

1. Using the **tcoord** , it may be specified by a coordinate frame.
2. An **Exodus** side set may be attached to one face of all the elements in a block using the keyword **sideset** . The thickness direction will be defined to be the normal to the sideset's surface. For example, if the sideset is placed on a side of the structure that lies on the x-y plane, then the thickness direction of the HexShell will be defined as the z direction, since that is the normal to the x-y plane.
3. **Sierra/SD** may attempt to determine the thickness direction from the topology. This is the default option (because it is the easiest for the user), but it is also the least robust.

Sierra/SD attempts to identify the element orientation first using **tcoord**. The **tcoord** keyword abbreviates thickness coordinate, and is only defined for HexShells. If **tcoord** is not specified, then **Sierra/SD** attempts to identify the element orientation second from the corresponding sideset. These methods do not depend on the decomposition, but the third method does depend on the decomposition. Lastly if no sideset is specified, **Sierra/SD** attempts to determine the thickness direction from the topology.

The element orientation may be identified in the output using the **eorient** keyword. See Section 7.1.43.

Thickness Determination by Topology

When the element thickness must be determined by the topology, the mesh must follow these requirements. The elements in the block must form a sheet. More than one disconnected portion of the sheet is possible, but all portions must adhere to these requirements.

- Every element in the sheet must have at least two neighbors, e.g. the sheet can't be a single element. NOTE... at this time, this is true for the parallel decomposed mesh too. The portions of the sheets found in each subdomain can not be a single element. We must be able to eliminate the thickness direction of each element by its neighbor connectivity.
- The elements in the sheet may vary in thickness, but the sheet must be exactly one element thick.
- The elements must be connected as a single sheet. Thus, if the sheet turns a corner, it must do so gently. The algorithm will fail if any element in the sheet is connected on the top or bottom to another element in the sheet.

Determining element thickness from the topology has known limitations, and is not planned for an update. This topology method is the oldest and depends on the body being laid out in a layer one element thick. Unfortunately, it is not well parallelized, as we do not have ghost elements. The other two methods do not depend on the decomposition.

HexShell Parameters

The **HexShell** requires a material specification. Optional parameters include the sideset or the coordinate frame and coordinate direction used to determine the thickness direction. The sideset keyword must be associated with a defined sideset in the model. The **tcoord** keyword requires a string and integer argument. The first is the Name of the coordinate system referenced. The second is the direction (1, 2 or 3) associated with the coordinate system.

#	Keyword	Arguments	Description
1	sideset	ID/name	sideset to specify thickness direction
2	tcoord	ID/name and direction	coordinate frame and coordinate direction
3	autolayer	# of layers and material	creates specified number of uniform layers of specified material

An example specification for a multi-layer HexShell is shown in input [5.2](#).

```
block 88
  HexShell
  sideset 88
  layer 1
    material 1
    coordinate 1
    thickness .4
  layer 2
    material 2
    coordinate 2
    thickness 0.6
```

```

end

block 89
  HexShell
  tcoord thickness 3 // azimuth
  material 1
end

begin cylindrical coordinate system thickness
  origin = 0 0 0
  Z point = 0 0 1
  XZ point = 1 0 0
end

block 100
  HexShell
  sideset 1 // the normal to sideset 1 will be the thickness
             // direction for block 100
  material 1
end

```

Input 5.2. Example HexShell Input. Block 88 is multilayer input with thickness direction determined using a sideset. Block 89 defines the thickness direction using a coordinate frame and the `tcoord` keyword.

HexShell Multilayers

The formulation of the HexShell supports multiple layers of orthotropic materials. Each layer has an associated material, normalized thickness and coordinate. The coordinate is provided to permit specification of the material coordinate. The thickness specifies the relative thickness of each layer. The total thickness is determined from the element topology, but relative thicknesses for each layer must be specified. If only one layer is specified, then the layer keyword is not required, and the relative thickness is irrelevant (and not required).⁴

There are two methods to specify multiple layers in a HexShell. The first, illustrated in input 5.2, provides complete flexibility over the material specification, orientation and thickness of each layer. An `autolayer` capability provides are much more limited specification that is useful for models of a single material with temperature dependence across the thickness. This `autolayer` feature creates the specified number of layers, of uniform thickness, of a single specified material.

⁴Layers for HexShells must specify the *relative* thickness of the layer. This is in contrast to layered shells which specify the absolute thickness (Section 5.7.8).

Materials for all HexShell specifications can be defined as a function of temperature, with the temperatures defined through the exodus file as element variables. The temperature can vary over both the elements and layers in the block.

```
block 1
  HexShell
  sideset 1
  autolayer 4
  material 1
end

Material 1
  name "steel"
  E function=1
  nu .3
  density 0.288
end

FUNCTION 1
  type Linear
  data 0 30e6
  data 1e6 20e6
END
```

Input 5.3. HexShell Autolayer Example. Here, exodus element variables define the temperature for each element on the block. Exodus layers must be of uniform thickness, and must be labeled “layer_temp1”, “layer_temp2”, etc.

When using temperature dependent materials, the temperature is obtained from the exodus file. The modulus is calculated as a function of temperature, and used in the element stiffness formulation. The temperature can vary both with layers and with elements. Any of the material parameters in either an isotropic or orthotropic material can be set to be temperature dependent. In the case of an isotropic material, any pair of two of the properties G , K , E , or ν can be temperature dependent.

The number of layers in the input file does not need to match the number of layers in the exodus file. The temperatures in the exodus file will be interpolated piecewise linearly to the center of the layer in the input file.

Temperature dependent orthotropic materials are supported for HexShells only. Temperature dependent densities are also supported.

Feature	Analytic Reference	Verification Section	Tested	Parallel Test	User Test
general	yes	54	Y	Y	some
multiple layers	no [†]	54	Y		

[†]Felippa’s report contains some verification. It has not been carried into **Sierra/SD**.

Table 5-118. – HexShell Verification Summary.

The mass properties of a layered HEXSHELL are computed approximately as follows.

1. The volume fraction, f_i , and density, ρ_i , of each layer is determined.
2. The contribution of the mass of the element is added to the nodes as if an element of density $\bar{\rho} = \sum_i \rho_i f_i$ filled the entire element.

The net affect of this is that the mass is computed as if an average density were applied. This could introduce minor errors if the element is thick and is much denser on one side than another.

For a HexShell if using `tcoord`, it is important to remember that the material definition may also use a non-default coordinate frame. In the next example, the thickness coordinate, `tcoord`, and the material definition use the same coordinate system.

```
begin cylindrical coordinate system 1000
  origin 0.0 0.0 0.0
  z point 1.0 0.0 0.0
  xz point 0.0 1.0 0.0
end
block 13
  HexShell
  tcoord 1000 1
  material 8
  coordinate 1000
end
```

5.9. *Beam2*

The **Beam2** element formulation is described in Section 3.14 of [52](#). This element is similar to the standard NASTRAN CBAR element, but it does not include a definition for a product of inertia or area shear factors. A product of inertia and area shear factors are included in the CBAR element in NASTRAN and are supported by the **Nbeam** element described in Section [5.10](#).

The use of a **Beam2** element requires a **block** definition with a **beam2**. The **block** definition must also have a **material** keyword referencing an isotropic material. Finally, the **Beam2** element must have a defined set of geometric parameters. Attributes may be defined in the input deck as follows.

```
block block_id
  Beam2
  material = material_id
  Area = area
  I1 = inertia_about_1
  I2 = inertia_about_2
  J = polar_moment_inertia
  orientation = x_orient y_orient z_orient
  offset = x_offset y_offset z_offset
end
```

The keywords are defined in following sections.

We define the local coordinate system before defining the attributes. The local coordinate system is defined by three axes – $\mathbf{x}_{elem}$, $\mathbf{y}_{elem}$, and $\mathbf{z}_{elem}$. The $\mathbf{x}_{elem}$ -axis is the axial direction. It lies along the length of the beam. The other two axes, the $\mathbf{y}_{elem}$ -axis and $\mathbf{z}_{elem}$ -axis, are determined by the orientation vector, $\mathbf{V}$. The local coordinate system and the orientation vector are shown in Figure 5-43. The orientation vector $\mathbf{V}$ lies in the plane defined by the $\mathbf{x}_{elem}$ -axis and the $\mathbf{y}_{elem}$ -axis – plane 1 in Figure 5-43. The $\mathbf{z}_{elem}$ -axis has the same orientation as $\mathbf{x}_{elem} \times \mathbf{V}$. Finally the $\mathbf{y}_{elem}$ -axis has the same orientation as $\mathbf{z}_{elem} \times \mathbf{x}_{elem}$.

The $\mathbf{x}_{elem}$ -axis and $\mathbf{z}_{elem}$ -axis define plane 2 in Figure 5-43. The $\mathbf{y}_{elem}$ -axis, which lies in the 1-plane, corresponds to a local 1-axis defined in a cross-sectional plane, a plane normal to the $\mathbf{x}_{elem}$ -axis. The $\mathbf{z}_{elem}$ -axis, which lies in the 2-plane, corresponds to a local 2-axis defined in the cross-sectional plane.

The **Beam2** element attributes are complicated. Four attributes are always required. A cross-sectional area, **area**, must be defined. The cross-sectional area can be defined with an **AREA** keyword. Two bending moments of inertia are also required. A bending moment of inertia for the 1-plane (bending about the $\mathbf{z}_{elem}$ -axis) is defined by the **I1** keyword. Bending moments in the 2-plane (or bending about the $\mathbf{y}_{elem}$ -axis) is defined using the **I2** keyword. I_1 and I_2 are the bending moments in their corresponding planes. I_1 and I_2 are not bending about their axes. This convention is consistent with many commercial codes including NASTRAN. A polar moment of inertia, **polar_moment_inertia**, for torsion about the $\mathbf{x}_{elem}$ -axis is required. The polar moment of inertia can be defined with the **J** keyword.

If the cross-section has the symmetry, $I_1 = I_2$, then the orientation vector $\mathbf{V}$ may be an optional attribute. Otherwise, $\mathbf{V}$ is a required attribute. It is necessary for the bending properties to have the correct global orientation. The components of the orientation vector




Figure 5-43. – Beam Orientation and Local Coordinate System.

can be specified with the values `x_orient`, `y_orient`, and `z_orient` using an `ORIENT` key word.

For attributes are provided in the **Exodus** mesh file, if an offset vector will be set through the attributes, then the orientation vector is always required.

The origin of the 1,2 coordinate system at the beam endpoints is the corresponding grid points by default. A user specified offset vector $\mathbf{V}_{off}$ translates the geometric location of the coordinate system origin. This offset vector is shown in Figure 5-44. For the **Beam2** element, one offset vector translates both ends of the beam. The `OFFSET` keyword is optional. The offset vectors move the beam neutral axis (the $\mathbf{x}_{elem}$ -axis) off the line that passes between the two grid points defining the connectivity of the beam. An offset is defined by a vector with values `x_offset`, `y_offset`, and `z_offset`. These values are associated with an `OFFSET` keyword.

When the offset option is used, the offset stiffness properties are equivalent to the stiffness generated by translating the beam by the offset direction and constraining the resulting offset nodes back to the untranslated nodes using rigid links. In addition, the offset mass properties are equivalent to the mass generated by translating the beam by the offset direction and constraining the resulting offset nodes back to the untranslated nodes using rigid links. For the **Beam2** element, *only* the component of the offset vector orthogonal to the element is used to compute the offset behavior for both the stiffness and mass.

For curved surfaces it is possible for the offset element to be inaccurate. The reason is that the radius changes, but the curvature does not change. Refer to section 5.7.9.1 for




Figure 5-44. – Beam Offset and Local Coordinate System.

limitations of element offsets.

The block parameters `area`, `inertia_about_1`, `inertia_about_2`, `polar_moment_inertia`, `x_orient`, `y_orient`, `z_orient`, `x_offset`, `y_offset`, and `z_offset`, may also be defined as attributes in the mesh file. Use the mesh file to specify attributes per element and the input block to specify attributes per element block. Input deck attributes override mesh file attributes. Attributes in the mesh file must be in the order specified in Table 5-119.

The **Beam2** element supports isotropic materials only.

The following section illustrates the use of the **Beam2** keyword in an element block definition. The element block has an integer block identifier of 3. This element block must consist of two node elements.

Table 5-119. – Attributes for Beam2.

	Keyword	Description
1	Area	Cross-section Area
2	I1	First bending moment
3	I2	Second bending moment
4	J	Torsion moment
5,6,7	Orientation	orientation vector
8,9,10	offset	beam offset vector

```

block 3
  Beam2
  Material 7
  Area 0.71
  I1 .05
  I2 5e-2
  J 0.994
  Orientation 1.0 -1.0 0.9
  Offset -3.14e-2 0.11 0.99
end

```

5.10. **Nbeam**

The **Nbeam** element was developed from the COSMIC/NASTRAN open source CBAR element. Unlike the **Beam2** element discussed in the previous section, the **Nbeam** element includes a definition for a product of inertia and definitions for area shear factors. The **Nbeam** element, currently, only has linear behavior implemented. If using a nonlinear solution method, the **Nbeam** element will not calculate a true internal force, but a linear force.

The use of a **Nbeam** element requires a **block** definition with a **Nbeam** keyword. The **block** definition must also have a **material** keyword referencing an isotropic material. Finally, the **Nbeam** element must have a defined set of geometric parameters. Most parameters for the **Nbeam** element may be entered either as attributes in the mesh file or through keywords in the **block** definition. Some parameters can be reset from default values only by use of the keyword definitions in the **block** definition. The general form of the **block** definition is as follows:

```

block block_id
  Nbeam
  material = material_id
  Area = area
  I1 = inertia_about_1
  I2 = inertia_about_2
  J = polar_moment_inertia
  I12 = product_inertia_12
  Shear_factor_1 = sfactor1
  Shear_factor_2 = sfactor2
  orientation = x_orient y_orient z_orient
  offset = x_offset y_offset z_offset
end

```

The various keywords in the above **block** definition are described in following paragraphs.

Local coordinate frame: Before describing the parameters for the **Nbeam** element, it is necessary to define the local coordinate system that is set up for beam elements in general. The local coordinate system is defined by three axes – $\mathbf{x}_{elem}$, $\mathbf{y}_{elem}$, and $\mathbf{z}_{elem}$. The $\mathbf{x}_{elem}$ -axis lies along the length of the offset beam. The other two axes, the $\mathbf{y}_{elem}$ -axis and $\mathbf{z}_{elem}$ -axis, are determined by an orientation vector, $\mathbf{V}$. The local coordinate system and the orientation vector are shown in Figure 5-45. The orientation vector $\mathbf{V}$ lies in the plane defined by the $\mathbf{x}_{elem}$ -axis and the $\mathbf{y}_{elem}$ -axis – plane 1 in Figure 5-43. The $\mathbf{z}_{elem}$ -axis is derived from the orientation vector $\mathbf{V}$ and the $\mathbf{x}_{elem}$ -axis by taking the cross-product $\mathbf{x}_{elem} \times \mathbf{V}$. Once the $\mathbf{z}_{elem}$ -axis is calculated, the cross-product $\mathbf{z}_{elem} \times \mathbf{x}_{elem}$ gives the $\mathbf{y}_{elem}$ -axis. As the **Nbeam** supports arbitrary vector offsets at each end, the orientation of the *offset beam* may differ from the orientation of the geometry (see “offset” below) that is not offset.




Figure 5-45. – Nbeam Orientation, Offset and Local Coordinate System. The coordinate system is in the plane of the *offset beam*. The plane is defined by the offset beam and the orientation vector, $\vec{V}$.

The $\mathbf{x}_{elem}$ -axis and $\mathbf{z}_{elem}$ -axis define plane 2 in Figure 5-43. The $\mathbf{y}_{elem}$ -axis, which lies in the 1-plane, corresponds to a local 1-axis defined in a cross-sectional plane, a plane normal to the $\mathbf{x}_{elem}$ -axis. The $\mathbf{z}_{elem}$ -axis, which lies in the 2-plane, corresponds to a local 2-axis defined in the cross-sectional plane.

Area: The cross-sectional area, **area**, must be defined either as exodus attributes or in the “block” section. The cross-sectional area can defined with an **AREA** keyword.

Bending Moments: The bending moments of inertia about orientation axes must be defined either in the exodus file, or the “block” section. A bending moment of inertia about the 1-axis (the local cross-sectional axis corresponding to the $\mathbf{y}_{elem}$ -axis), **inertia_about_1**, can be defined with the **I1** keyword. A bending moment of inertia about the 2-axis (the local cross-sectional axis corresponding to the $\mathbf{z}_{elem}$ -axis),

`inertia_about_2`, can be defined with the `I2` keyword. Finally, a polar moment of inertia, `polar_moment_inertia`, for torsion about the $\mathbf{x}_{elem}$ -axis is required. The polar moment of inertia can be defined with the `J` keyword.

The **Nbeam** element supports a product of inertia specification. The product of inertia about the 1,2-axes, `product_inertia_12`, is specified with the keyword `I12`. If the `I12` keyword does not appear, the value for `product_inertia_12` defaults to zero.

Shear Factor: The **Nbeam** element also has two area shear factor specifications. An area shear factor is a constant by which an average shearing strain on a beam cross-section must be multiplied in order to obtain the same transverse shear displacement as the transverse shear displacement that will be obtained from the actual shear strain distribution for the cross-section. Typically, the shearing strain varies over a cross-section. See Oden (Ref.⁴⁷) for a discussion of shear factors. An area shear factor for shear in the 1-direction, `sfactor1`, is specified with a `Shear_factor_1` keyword. If no `Shear_factor_1` keyword appears, the value for `sfactor1` defaults to 1.0. An area shear factor for shear in the 2-direction, `sfactor2`, is specified with a `Shear_factor_2` keyword. If no `Shear_factor_2` keyword appears, the value for `sfactor2` defaults to 1.0.

Orientation: The orientation vector $\mathbf{V}$ must be specified to assure that the bending properties of the beam have the correct global orientation relative to the rest of the structure. The components of the orientation vector can be specified with the values `x_orient`, `y_orient`, and `z_orient` using an `Orientation` keyword.

Offset: A user may specify offsets exactly as described Section 5.9. and shown in Figure 5-44. For the **Nbeam** element, the same offset vector is applied to both ends of the beam. The `OFFSET` keyword is optional. The offset vectors move the beam neutral axis (the $\mathbf{x}_{elem}$ -axis) off the line that passes between the two grid points defining the connectivity of the beam. An offset is defined by a vector with values `x_offset`, `y_offset`, and `z_offset`. These values are associated with an `OFFSET` keyword.

When the offset option is used, the offset stiffness properties are equivalent to the stiffness generated by translating the beam by the offset direction and constraining the resulting offset nodes back to the untranslated nodes using rigid links. For the **Nbeam** element, the full offset vector is used to compute the offset behavior, and different offsets may be applied at each end. (This behavior is different from the **Beam2** element, in which only the component of the offset vector orthogonal to the element is used to compute the offset behavior).

Note that for curved surfaces there may be modeling issues with offset elements, since there is no change in curvature with the change in radius. Refer to Section 5.7.9.1 for limitations of element offsets.

Many of the parameters described can also be defined as attributes in the mesh file. Attributes in the mesh file must be in the order specified in Table 5-120. If an attribute is

Table 5-120. – Attributes and Parameters for Nbeam.

#	Keyword	Description
1	Area	Area of beam
2	I1	First bending moment
3	I2	Second bending moment
4	J	Torsion moment
5,6,7	Orientation	orientation vector
8,9,10	offset	beam offset vector
11,12,13	–	offset of second node
–	I12	product of inertia
–	Shear_factor_1	shear factor 1-direction
–	Shear_factor_2	shear factor 2-direction

entered in both the mesh file and the input file, the value in the input file will supersede the value in the mesh file.

The **Nbeam** element is restricted to isotropic materials. No stress or strain output is available for **Nbeam** elements.

The following section illustrates the use of the **Nbeam** keyword in an element block definition. The element block has an integer block identifier of 3. This element block must consist of two node elements.

```
block 3
  Nbeam
  Material 7
  Area 1.92
  I1 2.57375
  I2 4.81277
  J 0.025816
  I12 -1.45983
  Shear_factor_1 0.44021
  Shear_factor_2 0.33313
  Orientation 1.0 0.0 0.0
  Offset 0.5 0.5 0.5
end
```

5.11. *TiBeam*

This is a Timoshenko beam with consistent mass. For beams with an aspect ration of 100 there is little difference between the beam2 and the tibeam. Beams with an aspect ratio of 10 differ significantly. The tibeam uses the parameters of a the **Beam2** . In addition the

effective area in shear should be provided (the default is 2/3). For a beam with square cross section,

```
block 3
  tbeam
  Material 1
  // Shear_factor = 10(1+nu)/(12+11 nu)
  areay = .84967320261438
  areaz = .84967320261438
end
```

In this example, the area, I_1 , I_2 , J and the orientation are all set in the Exodus file. Recall that for a circular cross section, the approximate shear factor is given in terms of Poisson's ratio, ν , by $6(1+\nu)/(7+6\nu)$. At this time, although the parser warns that **areay** and **areaz** are ignored, they are parsed and applied.

5.12. **Truss**

This is the definition for a **Truss** element based on Cook (Ref.²²). Trusses have stiffness in extension only. The **Truss** has 1 attribute as shown in the table. A linear elastic, isotropic material is required.

#	Keyword	Description
1	Area	Area of truss

No stress or strain output is available for trusses.

5.13. **Ftruss**

The **Ftruss** is a truss element whose stiffness is a function of the truss length.

Trusses have stiffness in only the axial direction. While they exist in a 3 dimensional world, forces orthogonal to the axial direction result in no resistance, i.e. they are singular. The axial force for an **Ftruss** element is defined as,

$$\vec{F}(\vec{L}_n, t) = -K(|\vec{L}_n|, t_n) \vec{L}_n \quad (5.1)$$

where $\vec{L}_n$ is the vector from the first point to the second at time t_n . Note that $|\vec{L}_n|$ is the instantaneous length of the truss. The force is *always* in the direction of the instantaneous element.

Denote by K_o the stiffness of a standard truss, and by L_o the nominal truss length. $F = -K_o dx$ implies that $K = \frac{K_o dx}{L_o + dx}$. The definition in equation 5.1 was chosen so that a force may be applied when dx vanishes.

Table 5-121. – Ftruss Attributes and Parameters.

#	Name	Type	Default	Comment
1	Area	<i>Real</i>	0	required if a material is specified.
2	Scale	<i>Real</i>	1	multiplier for the function
-	Function	<i>string</i>	<i>required</i>	function identifier (see section 2.8)
-	Material	<i>string</i>	<i>optional</i>	If the material specification is provided, it must point to a valid material (sec. 4), and an area must also be provided.

If a standard (non user) function is used, the stiffness is a function of truss extension only. It may not be a function of both extension and time.

Input to the **Ftruss** element is similar to that for the truss element. The attributes and parameters are listed in table 5-121, and a demonstration example is provided below. ⁵

```

block 88
  Ftruss
  function 88
  scale 1.0
  material 17      //optional material
  area 0.01        //area required iff material defined
end

```

If the material keyword is not found, no mass matrix is generated for the element. If a material is found, then *area* must also be defined. Like a standard truss, *area* is the first **Exodus** attribute. The area and material properties are used only to compute the mass properties of the element, and may be omitted.

The *scale* term can be defined using the input file, or alternatively, it may be defined using the second **Exodus** attribute.

5.14. ConMass

Concentrated masses are used to apply a known amount of mass at a point location. The **Exodus** file element topology is a sphere. Support for concentrated masses as two noded elements in exodus has been deprecated.

Parameters for the **ConMass** are listed below. Because of difficulties in translation or generation of the model, the parameters found in the **Exodus** file are not normally used for a **ConMass** . This avoids the confusion generated when mass constant defaults may

⁵Recall that attributes are ordered data that may be specified in the **Exodus** file, providing a variable which changes with each element. Parameters may be specified in the input file, and are applied uniformly to all elements in the block.

have been taken from beams for example. As a result, all parameters must be specified in the input or the analysis will fail.

This behavior can be tedious however, if many concentrated masses are found in the model, and if the analyst is confident that the attributes are appropriate for these elements. In this case, use the **ConMassA** element. It is identical to the ConMass, but uses the default attributes from the **Exodus** file. Typically seven attributes would be specified there.

A concentrated mass must have a mass. If no inertia tensor is specified, then the concentrated mass has 3 degrees of freedom, displacements. On the other hand a concentrated mass with an inertia tensor has both displacement and rotational degrees of freedom.

#	keyword	Description
1	Mass	concentrated mass
2	Ixx	<i>xx</i> moment of inertia
3	Iyy	<i>yy</i> moment of inertia
4	Izz	<i>zz</i> moment of inertia
5	Ixy	<i>xy</i> moment of inertia
6	Ixz	<i>xz</i> moment of inertia
7	Iyz	<i>yz</i> moment of inertia
8,9,10	offset	offset from node to CG

As an example element block,

```

block 5
  ConMass
  Mass 1000.0
  Ixx 1.0
  Iyy 2.0
  IZZ 1.5
  offset 30.0 40.0 50.0
end

```

The **ConMass** moments of inertia are defined at the location of the **ConMass** . The offset can be used to specify inertial terms about a different point.

A **ConMass** element will activate either 3 or 6 degrees of freedom on the node the mass is located. Every **ConMass** element will activate "DispX", "DispY", and "DispZ". A **ConMass** element with non-zero inertial terms or an offset will activate "RotX", "RotY", and "RotZ". In a case such as a spring-mass system where only one translational degree of freedom is desired, the mass should be constrained in the other directions. If **ConMass** elements are attached to solid elements, through shared nodes or a 2D element, either the inertial terms should be set to zero or the rotational degrees of freedom should be constrained. Failing to properly constrain the **ConMass** may result in a solver out of bounds error or incorrect results.

5.15. Spring

The **Spring** element provides a simple spring connection between two nodes in a model. Note that the direction of application of the spring should be parallel to a vector connecting the nodes of the spring. It is usually preferable to have the nodes of the spring be coincident. Springs are defined in the **Exodus** database using **Beam** or **Bar** elements.

The **Spring** element has three required parameters (translational spring stiffness). Rotational parameters are supported using the **RSpring** element described in section 5.16. Currently, there is no way to attach off-diagonal elements, i.e. there is no K_{xy} spring element. If that is required, a combination of a spring and a multi-point constraint must be used.

Springs can be defined in user defined coordinate systems.

#	Keyword	Description
1	Kx	translational spring constant in X
2	Ky	translational spring constant in Y
3	Kz	translational spring constant in Z

As an example element block,

```
block 51
  spring
  coordinate 7
  Kx 1e6
  Ky 1.11E7
  Kz 1000
end
```

5.15.1. Spring Parameter Values

It is strongly recommended that all three values of the spring constants be nonzero. This is especially important in parallel analysis performed using domain decomposition. Many domain decomposition tools may partition the model such that zero spring constants lead to singular domain stiffness matrices. This is true even if other elements may eliminate the singularity.

While setting nonzero spring stiffness helps to avoid solver problems, underlying domain decomposition problems may still exist for parallel calculations. Domain decomposition tools employ heuristics for connection of springs to solids; the models are not compatible. Finite length springs often result in constraints on rigid body modes. ⁶ Springs fill an

⁶This is not specific to parallel solutions. Most often, finite length springs introduce strain for a model rotation.

important analysis need, but analysts may find that in many cases it may be better to replace the spring elements by solid element meshes which more accurately represent the physical connection. While there are more degrees of freedom in the calculation, the accuracy is enhanced, and domain decomposition problems are mitigated.

5.16. *RSpring*

The **RSpring** element provides a simple rotational spring connection between two nodes in a model. It is usually preferable to have the nodes of the spring be coincident. RSprings are defined in the **Exodus** database using **Beam** or **Bar** elements.

The **RSpring** element has three required parameters (rotational spring stiffness). It is strongly recommended that all three components have some stiffness. This is particularly important when doing parallel analysis (see the discussion in section 5.15.1). Translational stiffness require the use of the **Spring** element described in Section 5.15. Currently, there is no way to attach off diagonal elements, i.e. there is no K_{xy} spring element. If that is required, a combination of an **RSpring** and a multi-point constraint must be used.

RSprings can be defined in user defined coordinate systems. The relevant parameters are listed in the table.

#	Keyword	Description
1	Krx	rotational spring constant in X
2	Kry	rotational spring constant in Y
3	Krz	rotational spring constant in Z

As an example element block,

```
block 52
  Rspring
  coordinate 7
  Krx=1e6
  Kry = 1.11E7
  Krz 0.1
end
```

5.17. *Spring3 - nonlinear cubic spring*

The **Spring3** element provides a nonlinear spring connection between nodes in a model. Note that the direction of application of the spring should be parallel to a vector connecting the nodes of the spring. It is usually preferable to have the nodes of the spring be coincident. Springs are defined in the **Exodus** database using **Beam** or **Bar** elements.

The **Spring3** element has nine required parameters (translational spring stiffness). There is no way to attach off diagonal elements, i.e. there are no K_{xy} spring elements. If that is required, a combination of a spring and a multi-point constraint must be used.

The force applied by the **Spring3** is defined as a cubic polynomial in each of the coordinate directions. Thus,

$$F_x = Kx1 \cdot u_x + Kx2 \cdot u_x^2 + Kx3 \cdot u_x^3 \quad (5.2)$$

For linear analyses, only the first term is used.

Cubic springs may be defined in user defined coordinate system.

#	Keyword	Description
1	Kx1	translational linear spring constant in X
2	Ky1	translational linear spring constant in Y
3	Kz1	translational linear spring constant in Z
4	Kx2	translational quadratic spring constant in X
5	Ky2	translational quadratic spring constant in Y
6	Kz2	translational quadratic spring constant in Z
7	Kx3	translational cubic spring constant in X
8	Ky3	translational cubic spring constant in Y
9	Kz3	translational cubic spring constant in Z

As an example element block,

```

block 51
  spring3
  coordinate 7
  Kx1 1e6
  Ky1 1.11e7
  Kz1 0
  Kx2 0
  Ky2 0
  Kz2 0
  Kx3 1e4
  Ky3 1.11e5
  Kz3 0
end

```

5.18. *Dashpot*

A dashpot represents a damping term proportional to velocity. Dashpot elements combine a viscous friction damper with a simple linear spring. The spring is included to avoid singular stiffness matrices when dashpots are connected without springs. Dashpots are

currently only used in transient dynamic, direct frf and complex eigen analyses. For other analyses only the spring term will be used.

The damping factor is the damping matrix entry. It has units of *force·time/length*. For a single degree of freedom system with a mass= M , the following equation is satisfied.

$$K \cdot u + c \cdot \dot{u} + M \cdot \ddot{u} = f(t) \quad (5.3)$$

Currently, dashpots are defined in the basic coordinate system only. Because they are single degree of freedom elements, the direction must also be defined (i.e. cid=1, 2 or 3). There are three parameters. All are required.

#	Keyword	Description
1	K	translational linear spring constant
2	c	damping factor
3	cid	coordinate direction (1, 2 or 3)

As an example element block,

```
block 51
  dashpot
  cid=1 // dashpot is in the X direction
  K=1e6
  c=1e5
end
```

Dashpots may be represented in the **Exodus** file with any linear element. The Truss element most closely mimics the dashpot's single degree of freedom behavior, and may be the best definition for domain decomposition tools.

Caution should be exercised when using dashpots (or any single degree of freedom element). The remaining degrees of freedom must be properly accounted for, or the system matrices will be singular. Care should also be exercised to ensure that if the nodes of the dashpot are not coincident, that the constraint force lies along the axis of the element - failure to do this can result in models that have nonzero rotational modes. There may also be important domain decomposition issues with dashpots. See Section 5.15 for a discussion.

5.19. SpringDashpot

The **SpringDashpot** element provides a general, fully coupled spring and dashpot connected to a pair of nodes. It is a linear element only, and is not corotational. It supports stiffness and damping in the translational and/or rotational degrees of freedom. The relevant parameters are described in Table 5-122.

#	Name	Description
1	Kxx	Translation Stiffness, K_{xx}
2	Kyy	Translation Stiffness, K_{yy}
3	Kzz	Translation Stiffness, K_{zz}
4	Kxy	Translation Stiffness, K_{xy}
5	Kxz	Translation Stiffness, K_{xz}
6	Kyz	Translation Stiffness, K_{yz}
7	Krxx	Rotation Stiffness, Kr_{xx}
8	Kryy	Rotation Stiffness, Kr_{yy}
9	Krzz	Rotation Stiffness, Kr_{zz}
10	Krxy	Rotation Stiffness, Kr_{xy}
11	Krxz	Rotation Stiffness, Kr_{xz}
12	Kryz	Rotation Stiffness, Kr_{yz}
13	Bxx	Translation Damping
14	Byy	Translation Damping
15	Bzz	Translation Damping
16	Bxy	Translation Damping
17	Bxz	Translation Damping
18	Byz	Translation Damping
19	Brxx	Rotation Damping
20	Bryy	Rotation Damping
21	Brzz	Rotation Damping
22	Brxy	Rotation Damping
23	Brxz	Rotation Damping
24	Bryz	Rotation Damping
25	coordinate	coordinate frame

Table 5-122. – SpringDashpot Parameters.

As shown in the table, all the elements of the matrices may be entered for this element. An example follows.

```
block 100
  SpringDashpot
    Kxx = 1e4
    Kyy = 1e4
    Kzz = 1e4
    Kxy = -1e4
    Kyz = -1e4
    Byz = 3.2
end
```

5.20. *Hys*

The **Hys** element provides a simple, one dimensional approximation of a joint going through microslip. Many simple joints can be represented by their *hysteresis* loop, a curve in the displacement vs. force plane. The relevant parameters of this element are indicated in the table, and illustrated in Figure 5-46.

#	Keyword	Description
1	Kmax	maximum slope of f vs u curve
2	Kmin	minimum slope of f vs u curve
3	fmax	maximum possible force
4	dmax	maximum possible displacement

The **fmax**, **dmax** pair define the limits of applicability of the element. The element will fail if the internal force exceeds **fmax** or the displacement exceeds **dmax**. The slope of the curve at the origin is **kmax**. It represents the small amplitude response of the system. The slope at the extremum, i.e. at (**dmax**,**kmax**) is **kmin**.

A **Hys** element uses a Beam or truss element in the **Exodus** file. At the current time, the element may only be defined in the *X* direction. An example of the **Sierra/SD** input is shown below.

```
block 2
  Hys
    Kmax 4.5e+7
    Kmin 3.0e6
    fmax 5.92
    dmax 0.9833e-6
end
```




Figure 5-46. – Hys element parameters.

5.21. Joint2G

The **Joint2G** element was added with the Iwan models. It models the adherence between surfaces.

Each **Joint2G** element connects a pair of nodes (or *grids*, hence the “G” in Joint2G). It permits users to specify independently the constitutive behavior of each of the degrees of freedom connecting its node pair.

The constitutive behavior is implemented through a constitutive class that provides generalized scalar forces in response to corresponding generalized displacements. The constitutive response was decoupled from the element to support general constitutive models.

Joint deformation under cycling loading is complex. The Iwan element can accurately model bolted joints, but depends on parameters that must be calibrated (using experimental data).

5.21.1. Specification

The meshed objects that map into the **Joint2G** element are defined in the **Exodus** database using **Beam** or **Bar** elements. The **Joint2G** ignores the **Exodus** file attributes. Properties are specified in the block and PROPERTY cards. In the example below, properties are assigned to element block “2”.

```
block 2
  coordinate 5
  shear_axis 2
  Joint2G
  kx=iwan    1
  ky=elastic 1.0e6
  kz=elastic 1.0e6
  krx=null
  kry=null
  krz=null
end
```

The above statement declares block 2 to be of type **Joint2G**. It also declares the constitutive response in the x to be the 4 parameter Iwan model.⁵⁶ The model parameters are specified in “Property 1” defined below. In this case, the four parameters chosen are χ , $\phi_{\max}$, R , and S . The Iwan properties can be specified alternatively by the parameter set χ , $\phi_{\max}$, F_S , and β .

```
property 1
  chi      = -0.82139
  phi_max  = 1.0325e-04
```

```

R      = 7.608594e+06
S      = 5.616950e+06
end

```

The constitutive behavior in the y and z directions is elastic with stiffness specified by the third argument - 1.0×10^6 in this case.

In this example, there is no specification for constitutive behavior in the three rotational directions. The *Null* specification merely means that those degrees of freedom in the relevant nodes are not activated (“touched”) by this element. Because of artifacts associated with parallelization, it is recommended that if any of the rotational degrees of freedom are active (not *Null*), they all should be active.

The directions (“x”, “y”, and “z”) employed above are those associated with the coordinate system declared for the block. In the example shown, there is an explicit reference to coordinate system 5. If there is no such explicit reference to a coordinate system, then the “x”, “y”, and “z” directions are those of the global coordinate system.

In the case when the joint2G element is used in conjunction with a Tied Joint, then the **shear_axis** can be used to specify the “x” direction for the constitutive response of the joint2G. Note that the **shear_axis** parameter is only meaningful when the joint2G is used in conjunction with a Tied Joint.

The **shear_axis** parameter allows the user to specify the “x” direction for the constitutive behavior. Since **shear_axis** is set to 2 in the above example, the “x” direction will be derived from the second component of coordinate 5. For more information on the **shear_axis** parameter, we refer to Figure 5-61 and Section 5.35.

5.21.2. Constitutive Behavior

Elastic Undamped, linear elastic behavior is defined by the **elastic** followed by the value of the parameter. No **property** section is required.

Damper Linear, damped behavior is obtained using a **damper** in the **Joint2G** definition, and using a property definition to specify the stiffness and damping terms. Typically each direction will require a different property definition.

```

block 3
  Joint2g
  kx=damper 1
  ky=damper 2
  kz=damper 3
  krx=null
  kry=null
  krz=null
end

```

```

property 1
  DAMPER
  K=1e6
  MU=.2
end

```

4-Parameter Iwan Model (iwan) The Iwan element is a collection of spring slider elements designed to provide a predicted model of joint behavior (including energy loss). Joint modeling with Iwan elements is described elsewhere.^{57,61} Descriptions of the relationship of the Iwan element to other joint elements are also available.⁵⁸

The schematic of the Iwan model is shown in figure 5-47. Parameters for the behavior may be specified using either an older definition (Table 5-123), or a newer set (Table 5-124). The newer parameters are summarized below.

#	Keyword	Description
1	chi	Exponent, χ , describing slope of force-dissipation curve at small amplitudes
2	R	Constant coefficient in distribution
3	phi_max	Maximum break free pseudo-force
4	S	Strength of singularity in break free force distribution
	alpha	Geometric factor specifying nonuniform spacing of dphi (optional, default = 1.2)
	sliders	Number of slider elements (optional, default = 50)

Table 5-123. – Older Iwan 4-parameter model.

#	Keyword	Description
1	chi	Exponent, χ , describing slope of force-dissipation curve at small amplitudes
2	beta	shape parameter of force/dissipation curve
3	K_T	Tangent stiffness at low loads
4	FS	Maximum break free pseudo-force
	alpha	Geometric factor specifying nonuniform spacing of dphi (optional, default = 1.2)
	sliders	Number of slider elements (optional, default = 50)

Table 5-124. – Revised Iwan 4-parameter model.

chi: determines the slope of the dissipation-force curve. Typically $0 < \chi < -1$. A value of zero corresponds to a coulomb type loss in Mindlin solutions. A value of $\chi = -1$ corresponds to a viscous like (but amplitude dependent) loss with dissipation proportional to the square of the amplitude. Dissipation follows the relation,

$$\text{Dissipation} \approx (\text{Amplitude})^{\chi+3}$$




Figure 5-47. – Iwan Constitutive Model.




Figure 5-48. – Dimensionless hysteresis curves for the four-parameter Iwan model with $\chi = -1/2$ and two values of β .

beta: determines the shape of the dissipation-force curve. Beta affects both the shape of the hysteresis curve within microslip (Figure 5-48), and the abruptness of the transition from microslip to macroslip as shown in Figure 5-49. $0 \leq \beta < \infty$.

KT: determines the slope of the force-displacement curve at low amplitudes. This is equivalent to a spring constant, and is used as such in analyses for which the element is treated linearly.

FS: determines the force at which the last slider gives out, and element goes into macroslip. The Iwan element is a statistical distribution of spring/slider elements. This is a point on that distribution.

The Reduced Iwan improves on the Iwan by requiring less calibration. Mechanisms to capture nonlinearity and dissipation are provided. Accuracy however depends on the fit with the experimental data under varying loading conditions.

Reduced Iwan Plus Pinning (RIwan) Reduced Iwan model, described later in the section, can be used as a constitutive model for a **Joint2G** definition using the **riwan**. This can be used for any desired direction, using the following format that requires associated property block:

```
block 3
  Joint2G
  kx=damper 1
  ky=riwan 2
  kz=damper 3
```




Figure 5-49. – Dimensionless static loading curves for the four-parameter Iwan model with $\chi = -1/2$ and three values of β , as the model goes into macroslip.

```

krx=null
kry=null
krz=null
end

property 2
FS=4e3
Kt=1.5e7
chi=-0.5
beta=0.005
Kp=2e7
dp =2e-3
end

```

Description: The RIwan element is a modification made by Brake (Ref.¹⁴), of the revised 4 parameter Iwan model described in the previous section. The modifications include (a) simplification of the Iwan force by assuming that, upon load reversal, the distribution of friction elements resembles scaled version of the original distribution, and (b) incorporation of pinning forces through approximation of Hertzian contact. The monotonic load-displacement curve for this element is shown in figure 5-50, where the initial nonlinear curve corresponds to microslip, the plateau corresponds to macroslip, and the steep linear curve represents pinning. The reader is referred to Ref.¹⁴ for further details, but the model can be summarized as follows:

$$F_{RIwan} = F_{Pin} + F_{Sliding}. \quad (5.4)$$




Figure 5-50. – Reduced Iwan Load Displacement Curve.

where the pinning force is given by the linear relationship,

$$F_{Pin} = \begin{cases} K_p(u - \delta_p), & u > \delta_p \\ 0, & \delta_p \leq u \leq -\delta_p \\ K_p(u + \delta_p), & u < -\delta_p \end{cases} \quad (5.5)$$

where δ_p is the pinning displacement shown in 5-50. The sliding force is given by

$$F_{Sliding} = \begin{cases} F_0 + \frac{F_S - F_0}{F_S} F_{Iwan} \left(u \frac{F_S}{F_S - F_0} \right) & \text{loading} \\ F_0 - \frac{-F_S - F_0}{-F_S} F_{Iwan} \left(-u \frac{-F_S}{-F_S - F_0} \right) & \text{reverse loading} \end{cases} \quad (5.6)$$

where F_0 is the shifted central force with $-F_S < F_0 < F_S$. Table 5-124 and figure 5-50 define $F_S, K_T, \chi, \beta, K_P, \delta_P$. We use the 4 parameter model Iwan force. $F_{Iwan} =$

$$\begin{aligned} & \frac{F_S(\chi+1)}{\phi_{MAX}^{\chi+2} \left(\beta + \frac{\chi+1}{\chi+2} \right)} \left(\left(\frac{1}{\chi+2} - \frac{1}{\chi+1} \right) u^{\chi+2} + \frac{\phi_{MAX}^{\chi+1}}{\chi+1} u \right) \\ & + \frac{F_S}{\phi_{MAX}} \frac{\beta}{\beta + \frac{\chi+1}{\chi+2}} \min(u, \phi_{MAX}) \end{aligned} \quad (5.7)$$

One Dimensional Gap Model (gap) The **Gap element** model attempts to represent the behavior of a gap closure with a bilinear elastic element. For proper numerical behavior, the stiffness of the open gap should not be more than a few orders of magnitude less than




Figure 5-51. – Eplas Model.

the stiffness when the gap is closed. The **Joint2G** implementation of the **Gap** model is identical to the axial behavior of NASTRANs cgap/pgap element and the axial behavior of the stand alone version of the **Gap** element implemented in **Sierra/SD** (Section 5.23).

#	Keyword	Description
1	Ku	Unloaded Stiffness
2	Kl	Loaded Stiffness
3	U0	Initial Gap Opening
4	F0	Preload (force at U0)

```
property 1
  ku  = 1e5
  kl  = 1e6
  U0  = 0.01
  F0  = 200
end
```

Elastic Plastic Hardening Model (eplas) **eplas** element is an elastic-plastic 1-dimensional element with linear isotropic hardening. Both the plastic strain and the hardening variable are initialized to zero. The parameters are illustrated in Figure 5-51.

#	Keyword	Description
1	k	Linear Stiffness
2	kp	Hardening Stiffness
3	fyield	Force at Yield

```

property 1
  k      = 1e6
  kp     = 1e5
  fyield = 1e4
end

```

One Dimensional Spring-Dashpot Model (damper) A **damper** represents a damping term proportional to velocity. Damper elements combine a viscous friction damper with a linear or cubic spring. The spring is included to avoid singular stiffness matrices when dampers are connected without springs. Dampers are currently only used in transient dynamic, direct frf and complex eigen analyses. For other analyses only the spring terms will be used. The behavior of this element is similar to **dashpot** , but also includes cubic terms.

The damping factor is the damping matrix entry. It has units of *force·time/length*. For a single degree of freedom system with a mass= M , the following equation is satisfied.

$$K \cdot u + \mu \cdot \dot{u} + M \cdot \ddot{u} + K_3 \cdot u^3 + \mu_3 \cdot \dot{u}^3 = f(t) \quad (5.8)$$

#	Keyword	Description
1	K	stiffness
2	Mu	viscous damper coefficient
3	K3	<i>Optional</i> cubic stiffness coefficient
4	Mu3	<i>Optional</i> cubic damping coefficient

```

property 1
  K      = 1e6
  Mu     = 1e2
  K3     = 1e4    // optional, default=0
  Mu3    = 0.1   // optional, default=0
end

```

Additional Constitutive Behavior The philosophy employed in the implementation of the **Joint2G** element of decoupling the constitutive behavior from the element machinery should facilitate the implementation of other constitutive models. Among those whose implementation is foreseen are the following:

- Bouc-Wen hysteresis model
- Preisach hysteresis model

5.22. Line Weld

Line welds offer a way to join two shells by a collection of virtual Joint2G elements (section 5.21). The position of the line weld is defined by a collection of beams which are meshed contiguously with one of the shells. Using a contiguously-meshed beam instead of the shell directly to define the line weld position enables more general weld definitions (e.g. skip welds).




Figure 5-52. – Line weld definitions for attaching the purple and cyan shells. The line weld follows the red beam, which is contiguously meshed with the purple shell.

In figure 5-52, the line weld will be used to join block 3 (in red) to surface 1 on the cyan shell. This is accomplished in three steps, each illustrated by figure 5-53:

- Create a new virtual node (in green) coincident with every existing block node
- Tie each virtual node to the nodes of the shell on which it lies
- Create a zero-length, virtual Joint2G between every block node and its coincident virtual node

Currently, line welds in **Sierra/SD** are only available for the elastic Joint2G axial constitutive model. These can be thought of as a general translational/rotational spring, with axes as shown in figure 5-52. Input 5.4 gives an example input for the line weld shown in figure 5-52. Note that rotational stiffness is only supported along the direction of the beam (the “r-axis”). Any rotational components along the other two axes will be ignored.

In keeping with Sierra/SM syntax, line welds may be specified using a force vs. deflection function as shown in Input 5.5. The current line weld capability is purely linear, and the line weld stiffness k is determined by estimating the derivative with a forward finite-difference stencil, i.e.,

$$k \approx \lim_{h \rightarrow 0^+} \frac{f(h) - f(0)}{h}.$$

A warning is issued if the corresponding backward stencil indicates that the derivative is not well-defined, i.e. if

$$\lim_{h \rightarrow 0^+} \frac{f(h) - f(0)}{h} \neq \lim_{h \rightarrow 0^-} \frac{f(h) - f(0)}{h}.$$




Figure 5-53. – Line weld Joint2G connections. The Joint2G elements are initially zero-length, and the green and red nodes coincident, but they are depicted with a finite deformation to delineate the individual components.

A fixed step size is used, so although the code will detect the case where the limits are not equal, it will not know whether the one-sided limits are well-defined. To aid the user in assessing the validity of these operations, the estimated stiffness values and other critical information is reported in the .rslt file.

Gap removal is enabled by default, and can be disabled using the **gap removal** option. The gap removal solution method 3.38 can make debug easier. Disabling gap removal can cause artificial grounding of rigid body modes, but can avoid collapsing elements when the search tolerance is large or when element quality is low. As with contact, the output mesh represents the model after gap removal.

```
begin line weld
  surface = surface_1
  block = block_3
  search tolerance = 1e-4
  r displacement elastic = 1.0e6
  s displacement elastic = 1.0e6
  t displacement elastic = 1.0e6
  r rotation elastic = 1.0e6
  gap removal = on
end
```

Input 5.4. Line weld input corresponding to figure 5-52

```
begin line weld
  surface = surface_1
  block = block_3
  search tolerance = 1e-4
```

```

r displacement function = R_force_function
r displacement scale factor = 1.0
s displacement function = S_force_function
s displacement scale factor = 1.0
t displacement function = T_force_function
t displacement scale factor = 1.0
r rotation function = Rrot_force_function
r rotation scale factor = 1.0
gap removal = on
end

```

Input 5.5. Force function syntax for line welds

Adding `line_weld` to an **outputs** or history block selects line weld output. To get the line weld output in the history file, the beam block from which the line welds were created must be included in the history section using **block** . The keyword `line_weld` adds the additional outputs shown in table 5-125 to the beam elements from which the line weld was created. The new virtual line weld elements created in the mesh have the usual output associated with a Joint2G element, such as **eforce** .

Line welds output force per unit length in the weld local coordinate system as `line_weld_force_rst` and `line_weld_moment_rst` .

Name	type	Description
<code>line_weld_weld_active</code>	int	A value of one marks beams within search tolerance to the surfaces that have active welds
<code>line_weld_initial_weld_length</code>	real	Length of the line weld beam element
<code>line_weld_force</code>	vector	Force per unit length produced by the line weld in the global XYZ system
<code>line_weld_moment</code>	vector	Moment per unit length produced by the line weld in the global XYZ system
<code>line_weld_force_rst</code>	vector	Force per unit length produced by the line weld in the element local RST system
<code>line_weld_moment_rst</code>	vector	Moment per unit length produced by the line weld in the element local RST system

Table 5-125. – Line weld output: note that these are intended to exactly match the equivalent outputs in **Sierra/SM**.

5.23. Gap element

Gap elements are modeled after the non-adaptive NASTRAN CGAP/PGAP elements. They are intended to provide a simple, penalty type element suitable for modeling simple connections. Note that these elements (like all beam-like elements) when embedded in solid meshes can result in difficult domain decompositions, and lead to load imbalance.

The **Gap element** is inherently nonlinear. In linear analysis, the element behaves approximately like a spring with the stiffness determined by KU and KL and a transverse stiffness, KT. The parameters of the element are listed in the table below and shown graphically in Figure 5-54.

#	Keyword	Description
1	KU	Unloaded stiffness
2	KL	Loaded stiffness
3	KT	Transverse stiffness (closed)
4	U0	Initial gap opening
5	F0	Preload, i.e. force at U0
6	coordinate	Coordinate frame.

The unloaded stiffness, KU, represents the stiffness of the element when the gap is open. It must be greater than zero. The loaded stiffness, KL, represents the stiffness when the gap is closed (as shown in the figure). The stiffness is KL when $U_A - U_B$ is greater than U0.

The initial gap opening and preload define the corner point in the force/deflection curve as shown in Figure 5-54. Typically, these will be zero.

A Gap element provides for transverse stiffness and friction. When the gap is closed, the transverse stiffness is KT. If the gap is open, the transverse stiffness is reduced to $KT' = KT \times KU/KL$.

The coordinate frame is an optional attribute of the gap element. The gap open and closes along the X axis of the frame. Note that the direction of the coordinate frame is important. The element determines a quantity $U_A - U_B$ along this coordinate axis. *This axis may not align with the coordinate alignment of the elements, which can lead to confusion.* If the coordinate frame is not provided, each Gap element will have a coordinate frame generated such that the gap opens and closes along the line between the two points. If the points are coincident, then a coordinate frame *must* be provided.

The Gap element is a simple penalty type element that somewhat mimics the effect of a physical gap. Choice of the value of KL is important to success of the element. Good values are somewhat in the range of the neighboring element stiffness. Too large a value can lead to matrix condition problems. Too small a value results in excessive softness and penetration in the gap.

Because the element is nonlinear, it has a significant impact on solutions. As described in Section 3.25 (and the **update_tangent** keyword), the default behavior for the nonlinear

solver is a partial Newton iteration. This means that the tangent stiffness matrix is not updated between iterations. Thus, if K_L and K_U are different, the solver will be using the wrong slope in the newton loop. Many, many iterations may be required for convergence. You may want to turn on the ‘nlresidual’ option in the echo section (see 7.7) which will put convergence information into the results file.

An example is shown below.

```
block 2
  GAP
  KL 4.5e+7
  KU 3.0e6
  KT=1e6
  f0 5.92
  u0=0.9833e-6
  coordinate 5
end
```




Figure 5-54. – Gap element Force-Deflection Curve.

Gap Issues The Gap element is definitely more complex than most elastic elements. Here is a partial list of “gotchas” that we have observed.

- Gaps should normally be zero length elements. Like springs, a gap that has a physical length will not be invariant to rigid body rotation. See Section 1.8.7. One approach to this would be to use a combination of beam and Gap elements. Note however, that if K_T is zero, and the gap opens and closes along the line between the beam endpoints, the element is invariant to rotation.
- The Gap element may use a coordinate frame to define its direction. *In this case the direction is not set by the nodal coordinates.*
- The direction of the Gap element must correlate to the displacement difference from $UB - UA$. It is easy to get this direction reversed.




Figure 5-55. – Mass bouncing off a Gap. With this large time step the model is not conserving energy. Reducing the time step is required to correct the problem.

- If you set **U0**, you must also set **F0**. This element does not constrain the force/displacement curve to go through zero. The input must do this. The Gap element may thus be used to enforce an initial displacement or force. That may not be what you want. It can cause slow convergence on the initial time step.
- Significant numerical damping may be required for convergence. Closing the gap can cause energy to be moved into higher frequencies. Without numerical damping, this energy can multiply until the solution becomes unstable. Numerical damping is best introduced by setting “**rho**” in the time integrator. Values of “**rho**=0.2” to “**rho**=0.7” have worked well. It is problem dependent.

Physically closing a gap would cause some energy loss, either by microslip, or by a small amount of local plastic deformation. Numerical damping can dissipate this energy that is removed from the physical system by means that are not included in the finite element model.

- This Gap element may not conserve energy. This is demonstrated in Figure 5-55, where a mass is dropped onto a gap. A completely elastic rebound would take the mass back to zero. Instead, it rebounds significantly above zero. This issue comes about because of time discretization. The mass “penetrates” the gap region too far, which stores too much energy in the element. It is then expelled with too much velocity. The only solution with this element is to reduce the integration step.

- Setting either KU or KL to zero is a recipe for disaster in parallel. Use a small positive value even if physically the unloaded stiffness may be zero.

5.24. **Gap2D**

The **Gap element** of the previous section provides a useful construct for planar type interactions. A common modeling issue is a bolt hole that is too large. To model this interaction an ellipsoidal Gap element (or **Gap2D**) may be required.

The **Gap2D** element operates like the Gap element except that the gap could open in 2 dimensions. The gap is open provided that the element displacement is within an ellipse defined by the major and minor axes.

$$\left(\frac{u_x}{U0X}\right)^2 + \left(\frac{u_y}{U0Y}\right)^2 < 1 \quad (5.9)$$

The major and minor axes of the ellipse are defined in the x and y direction of the coordinate frame.

Parameters of the **Gap2D** element are listed below.

#	Keyword	Description
1	KU	Unloaded stiffness
2	KL	Loaded stiffness
3	KT	Transverse stiffness (z direction)
4	U0X	Initial gap opening, major direction
5	U0Y	Initial gap opening, minor direction
6	coordinate	Coordinate frame.

While the gap geometry is defined as an ellipse, stiffness is not. In the open section of the element, the in-plane stiffness is KU, and is independent of direction. Likewise, in the closed gap region, the in-plane stiffness is independent of direction, and is defined by KL. The out of plane stiffness for this element is *always* KT. Note that the transverse stiffness behavior is significantly different from that of the standard **Gap element**.

The definitions above define the gradient of the force only, and for this nonlinear force, the value of the force depends on the path chosen for integration. For this element, we define the force as the integral along the shortest line from the origin.

In Figure 5-56, two possible integration paths are shown for arriving at the point (x_1, y_1) . In the first path, we integrate to $(x_1, 0)$ and then up to (x_1, y_1) . The y component of force is $f_y^{(1)} = KL \cdot y_1$. In path 2, we follow the straight line through (x_b, y_b) . The associated force is $f_y^{(2)} = KU \cdot y_b + KL(y_1 - y_b)$. For this element, we always choose the shortest line path (path 2). This ensures that the force is not history dependent.




Figure 5-56. – Gap2D force diagram.

5.25. **GasDmp**



GasDmp is currently BETA release.
Enable with the “`--beta`” command-line option.

The **GasDmp** element is a nonlinear, beam-like element that simulates the damping forces on MEMS devices due to gas pressure as MEMS beams vibrate. The element has no stiffness, but has damping roughly proportional to velocity/ L^3 , where L is the distance from the beam to the substrate. The element is experimental. Contact Troy Skousen of *Sandia National Labs* or Professor Burak Ozdoganlar at Carnegie Mellon for details.

Inputs to the **GasDmp** element are as follows.

#	Keyword	Description
1	W	Beam width (length units)
2	dL	Considered length of beam (length)
3	mm	Molecular mass of gas (mass)
4	p0	Ambient pressure of gas (pressure)
5	T	Ambient temperature of gas (temperature)
6	muRef	Reference viscosity (pressure * time)
7	TRef	Reference temperature (temperature)
8	ww	Viscous temperature exponent

Currently, the parameters are implemented through the input file and not through the Exodus_II file.

5.26. Nmount

The **Nmount** element is a Navy-specific mount element that provides an external force at user-specified points in the model. These forces are formed from a constitutive equation that is supplied by the user in the form of a subroutine. The Nmount capability provides an interface capability that allows the user to input their own subroutine that evaluates the constitutive equation.

An example of the user interface is shown in input 5.6. Mount orthogonal directions must be provided either as attributes in the **Exodus** file, or using the “Orientation” keyword in the “block” section. The relation of the orientation vector to internal element coordinates is shown in Figure 5-57. Remaining information is provided in the “block” section. Each mount type requires a separate block entry. Mount parameters are provided as text input in the block section. A list of built in Nmount types is listed in Table 5-126.

Each mount type may require a different number of mount parameters. If more parameters are provided than required for this mount, the additional parameters are ignored *without warning*. If fewer parameters are provided than are anticipated for the mount, the last parameters are set to zero, a warning is printed, and the analysis continues.

Mount type “7” is a special user subroutine mount. This mount type relies on Fortran or C functions compiled into the code at runtime. For this option **user subroutine file** must be defined in the file section. The user subroutines are incorporated into a custom build of the **Sierra/SD** executable with a command line like “sierra –make salinas -i myInput.inp”. See the Sierra/SM user manual for more details on the mount user subroutine formats and requirements.

```
block 41378
  Nmount
  mount type = 99
    parameters = 1.414 3.141 2.713
    orientation = 0 0.7 1
end
```

Input 5.6. Sierra/SD Mount Interface

Stability The Nmount element applies a force to the joining nodes in much the same way as an externally applied force. It provides no contribution to the stiffness matrix, and as such resembles an explicit element. Thus, stability issues can arise with this formulation. For certain models, damping has been shown to stabilize the formulation. The user may need to experiment with time step and damping levels to determine appropriate parameters for a stable solution.

Name	Index	Comment/Parameters
fail_truss	1	1 = Axial stiffness, Kr 2 = torsional stiffness 3 = critical value of tensile strain 4 = critical value of rotational strain
SpringDashpot	2	1 = Axial stiffness, Kr 2 = Stiffness in S frame, Ks 3 = Stiffness in T frame, Kt 4 = Axial damping, Cr 5 = Damping coefficient in S frame, Cs 6 = Damping coefficient in T frame, Ct
Snubber	3	1 = Model Direction (1=axial, 2=radial) 2 = Model Type (1=10k, 2=20k) 3 = Snubber gap 4 = offset 5 = input weight on compression 6 = input angle of radial action
TorsionalSpring	4	1 = Toronto spring constant, K_θ $M = K_\theta(\theta_2 - \theta_1)$
Test-spring	5	1 = K_x 2 = K_y 3 = K_z 4 = K_{θ_x} 5 = K_{θ_y} 6 = K_{θ_z} $F_j = K_j x_j$ and $M_j = K_{\theta_j} \theta_j$
NLSpring	6	1 = K_x 2 = K_y 3 = K_z 4 = W_x 5 = W_y 6 = W_z $\vec{F} = K x + W x^2$
subroutine	7	Mount behavior defined by mount force subroutine = <string> mount init subroutine = <string> mount size subroutine = <string>

Table 5-126. – Built in Nmount Models.




Figure 5-57. – Nmount Orientation

X_{elem} Normalized vector from node 1 to node 2. May change as the structure deforms.

$\vec{V}$ User provided orientation vector.

$$Z_{elem} = \frac{\hat{X}_{elem} \times \vec{V}}{|\hat{X}_{elem} \times \vec{V}|}$$

$Y_{elem} = \hat{Z}_{elem} \times \hat{X}_{elem}$. A normalized vector in the $X_{elem}\vec{V}$ plane, and orthogonal to X_{elem} .

5.27. **Rrod**

An **Rrod** is a *pseudoelement* which is infinitely stiff in the extension direction. The constraints for an **Rrod** may be conveniently stated that the dot product of the translation and the beam axial direction for a **Rrod** is zero. There is one constraint equation per **Rrod**.

The **Rrod** is specified using beams or trusses in the **Exodus** database, with a corresponding **block** section in the **Sierra/SD** text input file. No material is required. A block may contain Any number of connected or disconnected **Rrod** elements. The following is an example of the input file specification for an **Rrod** if the **Exodus** database contains beams in block id=99.

```
block 99
  Rrod
end
```

Attribute	default	description
RB_ID	-	translation identifier
CID_FLAG_INDEP	123456	independent coordinate flag
CID_FLAG_DEPEND	123456	dependent coordinate flag

Table 5-127. – Rbar Exodus Attributes.

5.28. *RBar*

An **RBar** is a *pseudoelement* which is infinitely stiff in extension, bending and torsion. The constraints for an **RBar** may be summarized as follows.

1. the rotations at either end of the **RBar** are identical,
2. there is no extension of the bar, and
3. translations at one end of the bar are consistent with rotations.

The **RBar** is specified using beams or trusses in the **Exodus** database, with a corresponding block section in the input file. No material is required and any number of connected or disconnected **RBar** s may be placed in a block. The following is an example of the input file specification for **RBar** s if the **Exodus** database contains beams in block id=99.

```
block 99
  RBar
end
```

RBar s can be reordered so that the number of **RBar** s connected to a single node is minimized. Having a large number connected to the same node results in a populated matrix and a slow computation. Therefore, reducing the number of connections can shorten run time. (see the *reorder_rbar* parameter in the **parameters** Section 2.3).

The **RBar** attributes are listed in Table 5-127, and are described below.

RB_ID Sometimes a collection of RBars is a description of a rigid body. This occurs for example when translating a NASTRAN model containing RBE2 elements. During translation these bars are grouped into rigid bodies based on their connectivity. The RB_ID is an index to that grouping.

CID_FLAG_INDEP By default, all degrees of freedom are active on both nodes of the RBar. Independent dofs are activated on the first node. The “CID_FLAG_INDEP” allows control over which degrees of freedom are activated. The flag is specified as an integer which is sum of components. ²

²It is an unusual descriptor, but it was designed to somewhat mimic the NASTRAN cid flag.

100000	X degree of freedom
20000	Y degree of freedom
3000	Z degree of freedom
400	R_x degree of freedom
50	R_y degree of freedom
6	R_z degree of freedom

Thus, ‘123456’ activates all dofs, and ‘123000’ activates only translations.

CID_FLAG_DEPEND By default, six dofs are eliminated from the bar. By setting this attribute to a non-default value, constraint equations may be skipped. The values are the same as the “CID_FLAG_INDEP” described above.

5.28.1. Interaction of RBars

If two RBars or sets of RBars share a node then they are effectively merged into a single rigid set. E.g. if an RBar connects nodes 1 and 2, and another RBar connects nodes 2 and 3 then the effect is the set of nodes 1, 2, and 3 act as a single rigid set of nodes. The same rule applies to intersection of rigid sets 5.33 or the intersection of RBars and rigid sets as the rigid set capability is effectively a means to automatically generate RBars.

5.29. RBE2

Sierra/SD has no support for the NASTRAN **RBE2** element. However, in most cases the **RBE2** element is not that different from a collection of **RBar** s.

5.30. RBE3

The **RBE3** pseudo-element’s behavior is taken from the eponymous NASTRAN element. More detail is found in the Theory Manual section **Sierra/SD** Elements subsection Rigid Elements subsubsection **RBE3**

The element is used to apply distributed forces to many nodes while not stiffening the structure as an **RBar** would. The **RBE3** uses the concept of a reference node.

Because all the nodes in an **RBE3** are not equivalent, each **RBE3** requires its own block ID. In the **Exodus** file, all links connecting to a single **RBE3** are defined in a single element block. The input file then specifies that this is an **RBE3** element block, as shown in the example below. If the model requires many **RBE3** elements, a separate block must be specified for each.

Usage The optional parameters for the **RBE3** pseudo-element are shown in the table below. These parameters must be specified in the input file, not as attributes of the **Exodus** file.

Keyword	value	Description
refc	<i>string</i>	reference node coordinates
WT	6 <i>reals</i>	relative weight of coordinates

refc The **REFC** parameter sets the degrees of freedom to activate on the reference node. For instance **REFC**='12' activates equations that constrain degrees of freedom associated with *X* and *Y* translations. No other degrees of freedom are affected. If the **REFC** keyword is not provided, it defaults to **REFC**='123456', i.e. constraint relations will be provided for the 6 structural degrees of freedom on the reference node.

WT. The contributions of each of the coordinates of the independent nodes may be scaled by **WT**. Most typically this would be used to determine the relative weight of rotational degrees of freedom on the independent nodes to the computation of the reference node rotations. The default value is **WT**= 111000 which means that the rotations do not contribute to the **RBE3**.

Generally we recommend there be no contribution from the rotations. The rotation of the element may then be determined solely from the translational degrees of freedom on the independent nodes.

The formulation of the **RBE3** is based directly on the published method from MSC/NASTRAN. Details of the method are described in the Theory Manual section **Sierra/SD Elements subsection Rigid Elements subsection RBE3**.

Cautions in using RBE3 Albeit convenient, the **RBE3** is not a true element. It can introduce complexity in the solution.

- A **RBE3** may connect a large portion of the model. This degrades linear solvers efficiency. As a consequence convergence may be slow.
- A **RBE3** connected to many nodes requires a lot of memory. This memory is stored on a single processor.
- No two MPCs should be linked together. Linear solvers may fail in this case.
- Accelerations (see Section 6.1.3) cannot be prescribed on an **RBE3** or any other MPC.
- The element has no logic to determine which degrees of freedom of the independent nodes are active. Thus, if you specify $WT = 111111$ the element will try to determine its rotation based on a combination of the translational and rotational degrees of freedom on the independent nodes. If the rotational degrees of freedom are inactive, they are treated as zero. This is rarely what is wanted.

- Care must be taken to ensure that only one node of the **RBE3** has multiple connections to its links. Further, every link in the **RBE3** must be connected to the reference node.
- A **Joint2G** with side averaging uses two **RBE3** elements.
- Many user issues are caused by **RBE3** elements.

Example RBE3 The following is an example of the input file specification for an **RBE3** if the **Exodus** database contains beams in block id=99.

```
block 99
  RBE3
  refc=123456
  wt=1 1 1 0 0 0
end
```

5.31. Superelement

Keyword	value	default	Description
file	<i>string</i>		input file containing matrices
format	<i>string</i>	netcdf	format of superelement input file: netcdf or DMIG
savememory	<i>yes/no</i>	no	controls storage of matrices in memory
diagnostic	<i>integer</i>	1	0 - run no diagnostics, 1 - compute $K_r * R_{BM}$, 2 - compute <code>eig(Kr,Mr)</code>
map	<i>integer</i>		table of node/cid pairs
map	<i>string</i>		“ascending_id” or “sorted” or “locations”
skip_output	<i>yes/no</i>		option to disable netcdf output
sensitivity_param	<i>integer/real</i>		parameter index and value of sensitivity
	parameter may be	used	multiple times to specify different parameters

A superelement is an abstract concept with different realizations in commercial codes. **Sierra/SD** does not support a full automatic superelement capability. The **Sierra/SD** Craig-Bampton reduction 3.5 reduces an entire model to a reduced-order model. Importing a reduced model (or superelement) into **Sierra/SD** is also supported, typically using **mksuper**. **Sierra/SD** supports superelements that can import of a mass and stiffness matrix into a full system model. This linearized approach may be used in any type of analysis.

Evaluating the suitability of a superelement approximation is left to the user. Few if any users have used superelements correctly based on only the information in *User's Manual*. Tutorials particularly the ones in⁵³ are recommended.

Limitations

- The superelement must be small enough (have a sufficiently small number of degrees of freedom) to fit (in the virtual memory) a single MPI rank. No consideration for superelements which span processors is made.
- Nodes on the superelement interface may be shared across processors. Interior degrees of freedom are local to a single processor.
- Output of the interface node degrees of freedom will be made in the base model in the usual way. Output of internal superelement quantities will be made in the superelement database file. The superelement modal degrees of freedom will be stored in the **Exodus** and MATLAB output on the X-degree of freedom of newly created virtual nodes.
- No automatic data recovery is available.
- Only a single level of superelement is supported.
- The mass properties report is computed by lumping mass to the interface dofs. For a superelement formed from a free-free system this will preserve the total translational mass of the superelement. However, the rotational inertia and center of gravity will not be exactly preserved in the superelement mass properties. If the superelement has an internal constraint or if for other reasons the superelement cannot exactly reproduce rigid body modes then the translational mass properties will not necessarily be preserved by the superelement.
- No geometric stiffness effects are currently accounted for in superelements. The default at this time is for these elements to return zero geometric stiffness.

User Input

The following input is provided by the user. If **format=DMIG**, the connectivity information is read from the **DMIG** file, and cannot be specified in any other way.

connectivity (Exodus): Note that codes such as NASTRAN input superelements by connecting to the nodes directly. Like any element with **Sierra/SD** the superelement must be mapped to a single processor. The superelement must be in the finite element database used to partition the elements. To provide the geometric connectivity to the model, the connectivity must be added to the **Exodus** file in one way or another.

If the superelement has the same number of nodes as a standard element, the analyst may choose to use such an element to provide the connectivity. This can facilitate

visualization of the model. **Sierra/SD** does support an element with *more* nodes than required for the connectivity map. Thus, a Hex-8 could be used to define the connectivity for a superelement with 7 nodes on the interface. The connectivity map cannot have more nodes than the element.

The **mksuper** utility function will add a superelement to an **Exodus** database. See.⁵³

connectivity map (Exodus): The equations for the system matrices must be associated with the nodes and degrees of freedom in the model. The following example creates a map for an eight degree of freedom reduced order matrix. The first column of the map is associated with the node index in the element. The second degree of freedom defines the coordinate direction (typically 1 to 6 for x , y , etc).

```
// node cid
map 0      0
    0      0
    1      1
    1      2
    1      3
    2      1
    2      2
    2      3
```

In this example, the first two rows of the system matrices are associated with internal degrees of freedom (DOFs) such as fixed interface modes. These interior dofs are indicated by a zero for both the node index, and the coordinate direction. Row 3 of the matrix is associated with the first node in the element connectivity, and with the x coordinate direction. Row 8 is associated with the second node, and the z coordinate direction.

There must be exactly as many rows in the connectivity map as there are rows in the system mass and stiffness matrices.

If the node index is negative, the row of the matrix associated with that degree of freedom will not be mapped to the system matrix. This can be used to “clamp” a generalized degree of freedom.

*The node index is not the node number in the **Exodus** file. It is the index into the element connectivity. Thus, for a four node element, the index must never exceed 4. This permits the use of **gjoin** and other tools without the need to reorder these terms in the input file.*

Alternate formats may be used to provide the map between rows of the system matrices and degrees of freedom of the residual structure. For these alternate formats to be used, the netcdf file containing the superelement data must include the **cbmap**

data, which provides an *internal* mapping between internal rows and columns and the internal nodes. These methods include the following.

map ascending_id or sorted If the user specifies the node number connectivity of the superelement in an ascending node order, then we can automatically generate the map. ³ Note, either *ascending_id* or *sorted* may be used here, they refer to identical algorithms.

map locations If the nodal coordinates of the superelement are stored in the netcdf reduced order model file, then the best match among coordinates of the residual and the superelement can be used to determine the map. This method works best if the superelement and residual have the same coordinate locations and if there are no collocated nodes in the interface. Each superelement interface node will be mapped to the closest finite element model node. No search tolerance is needed as the closest node is always found, even if it is far away. However, when using this option care should be taken that superelement interface nodes and finite element nodes match in space. If a superelement interface degree of freedom is mapped onto a finite element node with significantly different coordinates the superelement behavior may be substantially degraded and the superelement may no longer be able to represent rigid body modes. A warning is emitted by **Sierra/SD** if a node match cannot be found within a distance of one one-millionth of the characteristic model size.

system matrices: The system matrices may be provided in a **netcdf** or **DMIG** file. These matrices are available as output of the CBR reduction process (Section 3.5) and may also be generated with other tools such as Nasgen. The file must contain the following.

Kr. The reduced stiffness matrix. This is required for all analysis.

Mr. Most analyses require a reduced mass matrix as well. Its dimension must match that of the stiffness matrix.

Cr. A optional reduced damping matrix may be used. It must be of the same dimension as **Kr**.

maps that connect the degrees of freedom of the superelement to the degrees of freedom of the residual structure.

An accurate reduced **Kr** for 3D analysis should have exactly 6 zero energy modes. It must be symmetric (**Sierra/SD** will try to symmetrize it). Typically, **Mr** would be non-singular. Failure to meet these requirements can confuse the entire solution procedure, and lead to erroneous solutions.

³With the **mksuper** application, it is easy for the user to set up an element with ascending order, but most tools do not know how to visualize the element. Visualization may be easier using standard elements, but the the restriction that the connectivity have ascending node ids is confusing.

transfer matrices (Exodus): Output of results on interior points in the superelement are facilitated using optional output transfer matrices (OTM). These are described in the section on Craig-Bampton reduction (3.5). These matrices are written to the output **Exodus** file *only* if superelement output is requested in the **Craig-Bampton reduction** output specification. The following matrices apply.

OTM Nodal output transfer matrix.

OTME Element output transfer matrix.

OutMap An optional node map for the OTM.

OutElemMap An optional element number map for OTME.

skip_output: Optionally provides a means of disabling all output to the **netcdf** results files. This is particularly useful if the analyst wishes to use the same **netcdf** data for multiple superelements in the model. Without this keyword, each of these superelement blocks would be writing to the same file location, resulting in corrupted data.

output specifications: In the text input file superelement output is selected from either the **outputs** or **echo** sections by adding the word **superelement**.

If requested in the **outputs** section, then a new **Exodus** file will be generated from the information and name of the **netcdf** file. The number of nodes in the new file is the sum of the number of nodes on the interface and the number of nodes in the OTM. The number of elements is the number of elements in the OTME. All elements will be placed in a single element block.

Because we don't know the connectivity of the elements in the OTME, all such elements will be defined as sphere elements, and will be collocated on a single node in the model. This impedes visualization, but the element data is preserved for other types of post processing.

Likewise, no coordinate information is available for the interior nodes of the model. These elements will be located at the origin of the system.

sensitivity_param: If the **Craig-Bampton reduction** that generated the superelement included a sensitivity analysis, then the **netcdf** file containing the superelement matrices also contains derivatives of the reduced matrices with respect to the parameters. This information can then be used in the superelement block to set the superelement parameters as needed. This uses the linear Taylor series expansion of the sensitivity information of the Craig-Bampton model to compute the updated reduced matrices, and thus by-passes the need re-generate the Craig-Bampton model when the parameters are perturbed. The **sensitivity_param** allows the user to input specific values of the parameters for the superelement.

The above parameters are entered in the **block** section of the input file. For example,

```

block 10
  superelement
  file='example.netcdf'
  // node cid
  map 0    0
      0    0
      1    1
      1    2
      1    3
      2    1
      2    2
      2    3
  diagnostic=1
  sensitivity_param 1 0.01 // thickness in CBR shell model
  sensitivity_param 2 30e6 // modulus in CBR shell model
end

```

In this case, there are two sensitivity parameters, one for the thickness of a shell block in the Craig-Bampton model, and the other for the Young's modulus in that same block. Note that the format is assumed to be **netcdf** because the keyword **format** is not specified.

DMIG Input Files

Matrix	Acceptable Names
Stiffness (Kr)	K2GG,KAAX
Mass (Mr)	M2GG,MAAX
Damping (Cr)	C2GG,BAAX

Table 5-128. – Acceptable names of matrices within **DMIG** input files.

- **Sierra/SD** does not check the file extension of **DMIG** input files. Files generated by **Sierra/SD** use the **.dmig** extension, whereas direct output from NASTRAN uses **.pch**. Both are acceptable as input.
- Allowable matrix names are listed in table 5-128.
- Information required by **Sierra/SD** in parsing **DMIG** files comes from the matrices themselves: all comments are ignored, including e.g., those providing information about the dof map, number of interface modes, etc.
- The stiffness damping matrix **K4AXX** is ignored.

Finally, we show an example input deck for **DMIG** format:

```

block 32
  superelement
  format = DMIG
  file=rom4.dmig
end

```

5.32. *Dead*

A **dead** element has no mass and no stiffness. It may be of any dimensionality, solid, planar, line or point. Interior nodes to a block of **dead** elements will not be included in the computation of the model. There are no parameters for **dead** elements.

Note that for sideset-based load and boundary conditions, the boundary condition sideset is defined by sides of elements. If an element block is marked dead, then the boundary condition on the sides of those elements will generally also be deactivated. Special attention needs to be given to any sideset that is on the internal boundary between a dead block and an active block. Deactivation of the dead block may deactivate the boundary condition on that internal sideset if the sides are associated with the dead elements. In visualization tools the sideset normal will point out of the element it is attached to. Visualization of the sideset normal can be helpful to determine which block the sideset is attached in ambiguous internal sideset cases. As a general rule it is recommended that use of internal sidesets be avoided due to the model ambiguity they can cause.

5.33. *Rigidset*

Rigid Sets are intended as a usability tool to permit the analyst to treat a set of nodes as completely rigid. The input is straightforward.

```

Rigidset set1
  sideset 1
  sideset 2:5
  nodeset 88-90,wing
END

```

The above definition would establish a single set that is tied together. For purposes of error reporting only, the name “set1” is associated with this example set. If multiple *independent* sets are required, then multiple rigid set definitions may be made.

Table 5-129 shows the parameters. Any number of **rigidset** sections may be introduced. Each acts independently. Exodus sideset or nodeset information may be included in the definition. The rules for defining multiple nodesets, sidesets, or blocks at once are the same as the history output section, 7.3.

Table 5-129. – Rigidset Parameters.

Parameter	type	description
sideset	int/name/list	sideset id
nodeset (<i>not recommended</i>)	int/name/list	nodeset id
CenterNode tiedto node	integer	see below

Tied Node tied node or center node A rigid set can be attached to a Tied Joint 5.35 or Joint2g 5.21. In this case, a “reference” node may be generated and tied to another block or element. This is accomplished with the keywords below.

CenterNode tiedto node XX block YY

The **CenterNode** command will create a bar element with two nodes, and associate it with block YY. One end of the bar element is node XX. Node XX must exist in the input mesh. The other end node of the bar element will be created on the fly at the centroid of the rigid set. Note this capability only works for blocks with a single element. There are examples in.⁵³ Figure 5-58 illustrates the concept.




Figure 5-58. – Rigidset/TiedJoint Centernode Connection. The model illustrates the connection of a physical rigid set to a physical reference node via a virtual center node and virtual connection block.

5.33.0.1. Voltage Rigid Sets For many models using piezoelectric materials, a free surface of the piezoelectric tile may be plated with a purely conductive material such as copper. The conductive layer results in an equipotential surface. To simplify modeling an equipotential surface, **Sierra/SD** enables voltage rigid sets, which enforce a spatially constant voltage on the nodes associated with specified nodeset or sideset. Specifying voltage rigid sets are done in the rigid set block as shown in the following.

```
RIGIDSET unique_identifier
  sideset 1
  voltage
END
```

Input 5.7. Voltage Rigid Set

5.33.0.2. Limitations Rigidsets meet an important need to tie many nodes together. Generally they are much more robust than generating collections of **RBar** s or other rigid elements. However, it is easy to generate redundant constraints through this input. Redundant constraints cause most linear solvers to fail, and **Sierra/SD** may not always provide diagnostics. Generally,

1. Rigidsets should be completely disjoint, i.e. should share no common nodes. If two rigid sets do share any node they are effectively merged as a single larger rigid set.
2. If a Rbar is connected to any node of a rigid set then effectively the combination of Rigid set and connected Rbars behave as a single rigid system. Similarly, connecting two rigid sets via a Rbar effectively merges the two sets into a single larger rigid set.
3. None of the nodes in the rigid set should be constrained (as through a boundary condition).
4. While nodesets can be used to define rigid sets, this is not recommended because parallel decompositions may put only one or two nodes on a processor. So few nodes may introduce local singularities in rotation that impact the linear solver. If possible, use a sideset to define the rigidset.

5.34. *Rrodset*

Like the **rigidset** of Section 5.33, the **Rrodset** provides a convenient means of tying together a surface. All the limitations of the rigid set apply here. Unlike the rigid set, the **Rrodset** constrains only the distance between nodes on the faces, and no rotational degrees of freedom are constrained. The **Rrodset** acts much like a Kevlar skin; it resists stretching, but does not impede bending.

For a quadrilateral face, the **Rrodset** is equivalent to applying a rigid rod to each of the edges of the face. A constraint is also placed across one of the diagonals of the face as shown in Figure 5-59. An example is shown below.

```
Rrodset
  sideset 5
END
```

Like the **rigidset**, the **Rrodset** may be used to connect a “reference” node to a block.




Figure 5-59. – Rodset Constraints. The black lines indicate the edge of the element. Red lines are corresponding linear constraints.

5.35. *Tied Joint*

The *Tied Joint* models a joint structure. At the heart of the *Tied Joint* is an *Iwan* model. The *Tied Joint* supports flexibly mixing many models. An *Iwan* element may be used to represent the shear response, and multipoint constraints may be used to represent the normal response. Energy loss of the joint can be approximated by the *Iwan* element, and normal surfaces alignment can be preserved by the constraints.

5.35.1. Input Specification

Refer to Figure 5-60 for reference to the model definition. An example input is shown in input 5.8. There are several sections to the model definitions. Parameters of the input are summarized in Table 5-130. Details are below.




Figure 5-60. – Tied Joint Geometry. The two side set surfaces are shown separated for clarity. A virtual element is created which connects only the shear components of the joint. Normal components are interconnected using Tied MPCs.

```
Tied Joint
  Normal Definition = slip
  surface 3,5
  connect to Block 11 // Joint2g block
End

// definitions for the referenced Joint2g
Block 11
```

```

Joint2g
Kx=Iwan 1
Ky=Elastic 1e6
coordinate=5 // for anisotropic shear parameters
End

```

Input 5.8. Tied Joint Example

surface	2 sidesets connected by the joints, in slip constraints the 1st sideset is the face-surface and the 2nd the node-surface as in tied data
tied nodes	a sideset connected by the joint, the face-surface in slip constraints as in tied data
tied faces	a sideset connected by the joint, the node-surface in slip constraints as in tied data
normal slip none	if slip , normal-only node-face constraints
search_tolerance	if slip , defines node-face MPCs
edge_tolerance	if slip , defines node-face MPCs
shear connect to block side	reference block for whole joint average , rigid or Rrod

Table 5-130. – Tied Joint Parameters.

Name: Optional name of this joint. Useful primarily in diagnosing error messages.

Surface: Exactly two sidesets should be provided, these two sidesets will be connected via the joint. For node-face MPCs involved in normal direction constraints the nodes of the second surface (node-surface) are constrained to the faces of the first surface (face-surface)

Normal Definition: In the *Tied Joint*, the joint behavior in the normal direction can be governed by the joint element or by distributed node-face MPCs on the sidesets.

Slip implies the surfaces will remain in tied contact, and shear effects are managed by the “shear definition”. Tied node-face constraints will be created similar to the Tied Data [8.1](#) command. These constraints only tie the normal deformation of the sides together. The shear behavior of the joint will be managed by the single joint element. In this case, the joint element will be located at the nearest node to the centroid of the node-surface.

None implies that no specific node-face normal constraints will be generated. Surfaces may separate or interfere and the joint normal behavior will be controlled only by the whole-joint element. In this case, the joint will be located at the centroid of the node-surface.

Connect to block: A reference to a block containing parameters for the whole-joint element, a single element that connects the two sides of the joint. Usually a Joint2G element is used. If the normal definition is **none**, then the whole-joint element must specify behavior in all six dofs. Otherwise, if the normal definition is **slip**, then tied constraints are used to constrain the normal motion of the joint and only the three dofs associated with plane motion in the whole-joint element are used.

For non-isotropic shear behavior, the block may include a coordinate command. The frame may be curvilinear (e.g. cylindrical), in which case whole joint quantities are evaluated at the centroid of the surfaces (see coordinates, 2.7). To reference the basic (or default) frame, use coordinate frame “0”. The coordinate frame is specified in the connected element frame.

For curvilinear coordinate frames, it may be difficult to exactly specify the orientation of the centroid of the surface. Any user defined coordinate frame will be projected to the plane of the surface at the centroid, and a new coordinate frame is generated for specification of the orthogonal, in-plane coordinates. The $\tilde{X}$ and $\tilde{Y}$ axes of the user-defined coordinate system are projected to the plane, with the new third axis ($\tilde{Z}'$) in the normal direction.

In the case when the whole-joint element is a **Joint2G** element, the **shear_axis** can be used in the block definition to define the coordinate direction used for the first in plane constitutive component. We refer to Figure 5-61 for a description of the local coordinate system used to specify the constitutive behavior of the **Joint2G** element. The surface normal, n , is obtained as the normal on the node that is closest to the centroid of the sidesets that define the Tied Joint. This normal direction defines the $\tilde{Z}'$ axis of the local coordinate system. The **shear_axis** definition specifies which of the 3 axis of the user-specified coordinate system (in this example coordinate 5) is intended to be the first shear direction for the constitutive response. Thus, if **shear_axis** is set to 1, then $\tilde{X}'$ is defined as the part of the $\tilde{X}$ axis from the user-defined coordinate system that is orthogonal to $\tilde{Z}' = n$. If the $\tilde{Z}$ axis of the user specified coordinate system lines up exactly with the normal at the node, then the shear direction will be in exactly the same direction as the $\tilde{X}$ axis in the user-defined coordinate system. Generally, they will not line up perfectly, and this is the main reason why the **shear_axis** is needed. Once $\tilde{Z}' = n$ and $\tilde{X}'$ are defined, the remaining component of the coordinate system $\tilde{Y}'$ can be obtained by a cross product.

Side: The **side** defines additional constraints on the surfaces and how the tied joint Joint2G element spreads load to the surfaces.

A tied joint is between two surfaces A and B . To create the tied joint first a point in space is selected near the centroid of the joint. Next two new collocated nodes A_c




Figure 5-61. – The surface normal, n , is defined by the normal of the surface at the centroid point. The shear axis direction (in this example $\tilde{X}$) is projected onto the surface as $\tilde{X}'$. Together, $\tilde{X}'$ and surface normal provide the basis for the generated coordinate frame.

and B_c are created at this point. Node A_c will track the average displacement of side A . Node B_c will track the average displacement of side B . The whole joint element, a Joint2G, is connected between nodes A_c and B_c and controls the stiffness and damping of the joint.

average See Figure 5-62. A_c and B_c are connected to side A and B via RBE3 constraints. One RBE constrains node A_c to have the average displacement of side A and another RBEs constrains node B_c to have the average displacement of side B . The RBE3 constraints are added as bar elements to the output mesh file. Side average is the default for “normal definition” of “slip”.

rigid provides a means of constraining all the nodes on each surface to move together as a rigid set (see users 5.33). See Figure 5-63. Side A is turned into a rigid set and side B is turned into a rigid set. Node A_c is added to rigid set A and node B_c is added to rigid set B . An RBE3 is not required in this case as the rigid set already guarantees that the displacement at A_c and B_c track the average displacement of the sides. Side rigid is default for “normal definition” of “none”.

Rrod provides a means of constraining all the nodes on the surface to move together as a Rrod set (see users 5.34). See Figure 5-64. Side A is turned into a rrod set and side B is turned into a rrod set. As with side=average RBE3 constraints are needed to constrain nodes A_c and B_c to the average displacement of the two sides. The RBE3 constraints are added as bar elements in the output mesh file.

The side **rigid** option will create the stiffest overall joint behavior, **average** the softest, and **rrod** something in between. Using option **rigid** to avoiding a complex RBE3 constraint can also avoid numerical issues that sometimes come with the **average** or **rrod** options.

Not all Tied Joint specifications are fully consistent. In particular, the specification of the “normal definition” and the “side” descriptions are not fully independent. Table 5-131 summarizes some dependencies between these two parameters.




Figure 5-62. – Construction of tied joint with side=average.




Figure 5-63. – Construction of tied joint with side=rigid.




Figure 5-64. – Construction of tied joint with side=rrod.

Normal Definition	Side	Status
none	average	RBE3s constrain joint end nodes to have the average displacement of each surface.
	Rrod	RBE3s constrain joint end nodes to have the average displacement of each surface. Additionally each surface is turned into an rrod set.
	rigid	Each surface is turned into an rigid set which also contain the joint end nodes. An RBE3 is not needed in this case as the rigid set already constraints the joint end nodes to have the average displacement of the surfaces.
slip	average	Node-face normal only constraints are added to the surfaces. RBE3s constrain joint end nodes to have the average displacement of each surface. The shear deformation of the whole-joint model is based on the relative lateral motion of each side.
	Rrod	Node-face normal only constraints are added to the surfaces. RBE3s constrain joint end nodes to have the average displacement of each surface. Rrod constraints stiffen the surface. The shear deformation of the whole-joint model is based on the relative lateral motion of each side.
	rigid	Invalid. Overly constrained joint. Fatal error.

Table 5-131. – Tied Joint, “Normal” and “Side” dependencies.

Output Specifications Because the *Tied Joint* is not fully represented in the **Exodus** database (except as a collection of surfaces), standard element output capabilities are insufficient to represent the data. The data is divided into two categories: configuration and results.

5.35.1.1. Configuration Output The configuration output is only available in the text output of **Sierra/SD**, i.e. in the .rslt file. It is requested with the keyword “input_summary” in the “ECHO” section (see 7.7). This includes the following.

1. The type of the normal enforcement.
2. Surface information.
3. Centroid of the surface pairs (if applicable).
4. Owning processor for the shear elements (if applicable).
5. Shear models.

Results Output The only results output that is currently available for a Tied Joint consists of the forces in the **Joint2G** element that connect the two surfaces of the Tied Joint together. Currently, these forces can be only obtained in the history (or frequency) file for a transient or nonlinear transient analysis. They cannot be written to the global **Exodus** output file. If we consider the same example that is given in input 5.8, we could obtain the element forces as follows for a transient analysis

```
History
  block 11
  EForce
end
```

or, for a frequency domain analysis,

```
Frequency
  block 11
  EForce
enD
```

where in this example block 11 is the **Joint2G** block that connect the two surfaces of the Tied Joint together.

6. BOUNDARY CONDITIONS AND INITIAL CONDITIONS

6.1. *Boundary conditions*

Boundary conditions are specified within the **boundary** section. Node sets, side sets, blocks, or node lists may be used to specify boundary conditions. The rules for defining multiple nodesets, sidesets, or blocks at once are the same as the history output section, 7.3, and also follows the rules for integer lists detailed in Section 2.1. Specialized coordinate systems that can be used are described in section 2.7. The example in input 6.1 illustrates the method.

```
boundary
nodeset end_nodes      // for nodes in nodeset "end_nodes"
  x = 0                // constrain x=0
  coordinate 1          // for x/y/z in coord system 1

nodeset 1              // for nodes in nodeset 1
  x = 0.1              // constrain x=0.1
  y = 0                // and y=0
  RotZ = 0             // and the rotational dof about Z

nodeset 13:15, widget  // for nodesets 13-15 and "widget"
  accelx = 0.3          // constrain the x-acceleration
  function=1            // varying in time with function 1
  disp0 = 0.1           // and initial condition
  vel0 = 0.2            // disp0 = 0.1*0.3 and vel0 = 0.2*0.3

sideset 5              // for nodes in sideset 5
  absorbing            // use an absorbing bc

sideset 2              // for the acoustic sideset 2
  p = 0               // fixed acoustic pressure
                      // (pressure release condition)

sideset acousticSide   // for sideset "acousticSide"
  pdot = 1.0          // constrain the time derivative
                      // of acoustic pressure for
                      // enforced accelerations
  function = 2         // varying in time with function 2
  p0=1.0              // and initial condition p0 = 1.0

sideset 6              // for sideset 6
  impedance_pressure=0.5 // use a pressure impedance bc
  impedance_shear = 0.5 // and a shear impedance bc
```

```

sideset 7          // for sideset 7
  slosh = 0.6      // use a slosh bc

sideset 8          // for sideset 8
  infinite_element // use infinite elements
  use block ie_block // where the infinite element
                    // parameters are in the
                    // user-defined block "ie_block"

sideset piezoelectricSide // for "piezoelectricSide"
  V = 0             // fixed voltage
                    // (electrical ground)

block fixed_block   // for nodes in block "fixed_block"
  fixed             // constrain all dofs

node_list_file='nlf.txt' // for nodes in the file "nlf.txt"
  fixed             // constrain all dofs
end

```

Input 6.1. Example boundary section

The descriptors for the displacement boundary conditions are, X, Y, Z, RotX, RotY, RotZ, P, and **fixed**. Their application and meaning are listed in Table 6-132. An optional equals sign separates each descriptor from the prescribed value. The value **fixed** implies a prescribed value of zero for all degrees of freedom.

Note however that the syntax checking in the boundary block does not check for duplicate boundary conditions, and silent failures are possible. For example in the example shown in input 6.2, part of the input block is silently ignored.

```

Boundary
  nodeset 10 // This
    rotx = 0 // is
    roty = 0 // parsed
    rotz = 0 // as
  nodeset 10 // expected
    accelx=370.
    function=1
    accely=380. // Parsing
    function=2 // silently
    accelz=390. // ignores
    function=3 // these!

```

Keyword	Description
<i>prescribed displacement keywords</i>	
X	X Component of displacement
Y	Y Component of displacement
Z	Z Component of displacement
RotX	Component of Rotation about X axis
RotY	Component of Rotation about Y axis
RotZ	Component of Rotation about Z axis
fixed	Constrain all components of rotation and translation
P	Acoustic pressure
V	Voltage
<i>prescribed acceleration keywords</i>	
AccelX	scaling factor on X component of motion
AccelY	scaling factor on Y component of motion
AccelZ	scaling factor on Z component of motion
RotAccelX	scaling of rotational motion about X axis
RotAccelY	scaling of rotational motion about Y axis
RotAccelZ	scaling of rotational motion about Z axis
AccelV	second derivative of voltage
disp0	initial displacement
vel0	initial velocity
Pdot	derivative of acoustic pressure
P0	initial acoustic pressure
<i>prescribed displacement keywords (directfrf-only)</i>	
DispX	scaling factor on X component of motion
DispY	scaling factor on Y component of motion
DispZ	scaling factor on Z component of motion
RotDispX	scaling of rotational motion about X axis
RotDispY	scaling of rotational motion about Y axis
RotDispZ	scaling of rotational motion about Z axis
FreqV	scaling factor on voltage
FreqP	scaling factor on acoustic pressure

Table 6-132. – Dirichlet Boundary Enforcement Keywords.

```
enD
```

Input 6.2. Example Silent Failure of Parsing

In order to apply functions 2 and 3, it is necessary to provide a nodeset for each function as in input 6.3.

```
Boundary
  nodeset 10
    rotx = 0
    roty = 0
    rotz = 0
  nodeset 10
    accelx=370.
    function=1
  nodeset 10
    accely=380.
    function=2
  nodeset 10
    accelz=390.
    function=3
enD
```

Input 6.3. Corrected syntax

The way that the parser works for the boundary block is that the text is divided up into consecutive chunks by the keywords `nodeset`, `sideset`, `block`, `node_list_file`, and `end`. This is the pattern followed in input 6.1. Within each chunk, any number of Dirichlet boundary enforcement keywords (see Table 6-132) may be provided. Surprisingly, only the first function in a chunk is parsed; any others are ignored. There is no warning for ignored text in the boundary block at this time.

6.1.1. Prescribed Displacements and Pressures

In linear statics, one may prescribe a nonzero displacement by entering a value following the coordinate direction, as shown in input 6.4.

```
Boundary
  nodeset 1
    x = 3
    y = 0
```

```

z = 0
rotx = 0
roty = 0
rotz = 0
end

```

Input 6.4. Prescribed Displacement for Statics

For acoustics, pressures may be fixed by specifying $p = 0$, as in Table 6-132 on sideset 2. This corresponds to a pressure release condition.

For linear statics, there must be no **function** entry following the entry. Prescribed displacements have the same limitations as prescribed accelerations described in the next section. The load in this case is introduced by the prescribed displacement. However, the **loads** section must exist (for error checking purposes) even if it is empty.

6.1.2. Prescribed Voltage

For electro-mechanical coupled physics problems, constant voltage boundary conditions may be specified on nodesets or sidesets using the keyword **V** in the Boundary block. In the following example, electrical grounds were set at sideset 1 and nodeset 1, and a non-zero constant voltage boundary condition set at sideset 2.

```

Boundary
  sideset 1
    V = 0
  nodeset 1
    V = 0
  sideset 2
    V = 2
end

```

Input 6.5. Prescribed Voltage

6.1.3. Prescribed Accelerations

, ,

In transient dynamics, the acceleration on a portion of the model may be prescribed as a function of time. As shown in Table 6-132, acceleration are specified using **accelX**, **accelY**, **accelZ**, **RotaccelX**, **RotaccelY**, **RotaccelZ**, **disp0**, **vel0**, **Pdot** and **accelV**

A function must be used to apply the time-dependent boundary accelerations. Optional initial displacement and velocity can also be specified; if not, they default to 0. In the example above, the x acceleration of nodesets 13-15 and “widget” will be prescribed as $0.3 \times f(t)$, where $f(t)$ is defined in function 1. The *accelx* factor also scales the initial displacement and velocity. Thus initial displacement is given as $0.1 * 0.3$ and the initial velocity is $0.2 * 0.3$.

Prescribed accelerations are ultimately enforced in the code by integrating to produce a prescribed displacement as shown in Equation 6.1. The *start time of the function*, *tstart*, is not the start time of the analysis. Accounting for this is important in hand-off analyses. A function is required. Not listing a function will generate an error message.

$$u(t) = \text{scale factor} * \left[\int_{tstart}^t \left(\int_{tstart}^t f(t) dt \right) dt + (t - tstart) * v_0 + u_0 \right] \quad (6.1)$$

In the case of an acoustic sideset or nodeset, the prescribed value is the first time derivative of acoustic pressure, denoted above as *Pdot*. This is because, internally, **Sierra/SD** solves for the velocity potential, and the first time derivative of the velocity potential is the acoustic pressure. Thus, by specifying the first time derivative of pressure, one is prescribing the acceleration of the velocity potential.

In the case of a prescribed voltage time history, the prescribed value is the second time derivative of the voltage, i.e. voltage acceleration (*accelV*). However, since the first and second time derivatives of voltage do not contribute to the equations of motion (see theory reference for details), **Sierra/SD** also permits direct specification of the voltage time histories. A voltage time history may be directly defined by the function following the keyword *timeV*.

An additional point to consider when applying prescribed accelerations is that the initial velocity and displacement (denoted as *disp0* and *vel0*), are also necessary to completely define the boundary condition. These values account for the constants of integration obtained when integrating the prescribed acceleration to obtain the corresponding velocity and displacement on the sideset or nodeset. In the case of acoustics, only one initial condition is needed (*p0* which specifies the initial acoustic pressure), since only the first time derivative of acoustic pressure is specified. In the case of prescribed voltage acceleration, the descriptors *disp0* and *vel0* are used to define, respectively, the initial voltage and initial time derivative of voltage. By default *disp0*, *vel0*, and *p0* vanish.

There are some limitations with the prescribed acceleration capability, which are listed in the following, and in Table 6-133. First, prescribed accelerations are not currently set up to work with multicase solutions. Also, they only work in the standard (Cartesian) coordinate system. Prescribed accelerations can be used in meshes that have nonlinear or viscoelastic elements, as long as the prescribed accelerations are not applied directly to the nonlinear or viscoelastic elements. Note that the nodes involved in prescribed accelerations cannot coincide with nodes that are involved with MPCs.

Finally, note that when prescribed accelerations are used, they induce a load on the structure. Thus, in many cases the **loads** section serves no purpose, unless an additional external load is applied. In these cases, however, an empty **loads** block is still needed in the input file. An error message will be generated if the input file has no **loads** section.

- | |
|--|
| <ol style="list-style-type: none"> 1. No support for multibase. 2. Only in basic coordinate directions. 3. Cannot be used on nodes attached to viscoelastic elements. 4. Cannot be used on nodes attached to nonlinear elements. 5. Cannot be used on nodes connected to rigid elements or MPCs. 6. Load section is required, even if empty. |
|--|

Table 6-133. – Limitations for Prescribed Acceleration Boundary Conditions.

6.1.4. Prescribed Frequency-Varying Displacements

For the direct frequency response solution method, a portion of the model may be prescribed displacements as a function of frequency. As shown in Table 6-132 the descriptors for prescribed frequency dependent displacements are **DispX**, **DispY**, **DispZ**, **RotDispX**, **RotDispY**, **RotDispZ**, **FreqP**, **FreqV**. A function must be used to apply the frequency dependent boundary condition.

6.1.5. Node_List_File

To make it easier to apply boundary conditions, a *node_list_file* option is provided. In this option, the user provides an additional text file that contains a list of global node ids separated by white space. No comments, or other characters are allowed in the file, as shown in input 6.1. The remainder of the boundary condition specifications are unchanged.

There are several limitations place on collections of nodes specified in this manner.

1. This is an inefficient method of supplying the nodes. It is recommended that nodesets or sidesets be employed when practical.
2. No node distribution factors may be provided.
3. The output **Exodus** file will have no record of this list.

4. The global node numbers are the mapped **Exodus** global numbers. This is the arbitrary node numbering provided by the analysts, and may not correspond to the 1 to N ordering used by some other programs. ⁴
5. There is NO requirement that the nodes be sorted in the list, but repeating a node in the list can have undefined results, i.e. don't do it.

6.1.6. Nonreflecting Boundaries

Nonreflecting boundary conditions for acoustics and for elasticity may be specified using the “absorbing” keyword.

This section allows the user to specify an exterior boundary for acoustic, elastic, or coupled structural acoustic simulations. Once specified, first-order non-reflecting boundary conditions are applied on this surface. The boundary is specified with a sideset. The sideset can be placed either on acoustic or elastic elements. The code automatically determines whether the sideset is placed on acoustic or elastic elements, and then applies the appropriate boundary conditions.

Only pressure waves need to be absorbed for acoustic elements, and the absorbing boundary could represent an infinite fluid surrounding a structure. For elastic waves both pressure and shear waves need to be absorbed, and the absorbing boundary could represent an infinite elastic medium, such as in a seismic problem.

An example of this syntax is given below.

```
Boundary
  sideset 5
    absorbing
    radius = 1.0
end
```

The parameter “radius” specifies the radius of the sphere that defines the absorbing boundary. For a planar absorbing surface, one can either specify no radius, or a large radius (the radius is equal to infinity for a planar surface). In those cases, the absorbing boundary condition reduces to a plane-wave absorbing condition. We also note that the radius parameter refers to the distance from points on the spherical surface to the center of curvature, not to the origin of the coordinate system. Thus, it is independent of the coordinate system that is specified. For example, one could shift the coordinates of the nodes of the acoustic mesh by any constant, but the radius parameter would remain the same.

⁴Earlier versions of **Sierra/SD** used the 1 to N ordering, where N is the maximum number of nodes in the model. Current versions always use the mapped ordering for node references.

6.1.7. Impedance Boundary Conditions

Impedance boundary conditions are partially reflecting and partially absorbing. Thus, they are somewhere in-between a rigid wall and an absorbing boundary condition. They reduce to these special cases for certain choices of the impedance parameters.

An example syntax for an absorbing boundary condition is given below

```
Boundary
  sideset 6 // sideset on acoustic material
    impedance = 0.5
  sideset 7 // sideset on elastic material
    impedance_pressure = 0.5
    impedance_shear = 0.5
end
```

In this case, sideset 6 is attached to acoustic elements, and sideset 7 is attached to elasticity elements. For acoustic elements, only one impedance parameter is needed, and it corresponds to an impedance condition for pressure waves only (acoustic elements support no shear waves). For elasticity elements, the **impedance_pressure** and **impedance_shear** correspond to impedance for pressure and shear waves, respectively. This example specifies that sideset 6 is to have an impedance of $Z = 0.5\rho c$, where ρ is the density and c is the speed of sound. Thus, the “impedance” parameter that is parsed in is the multiplier on the characteristic impedance ρc . Similarly, for the elasticity element the pressure and shear impedance would be $Z_P = 0.5\rho c_P$ and $Z_S = 0.5\rho c_S$, where c_P and c_S are the speeds of sound for the pressure and shear waves, respectively.

Currently, impedance boundaries are only set up to work with the standard characteristic impedance ρc . Thus, specifying the “radius” parameter with an impedance boundary condition will have no effect.

We note that if the impedance parameters are all set to 1.0, the problem reduces to the absorbing boundary described in the previous section. If set to 0, the impedance condition becomes a pressure-release boundary for acoustics and a free boundary for an elasticity element. If set to a large number, the impedance boundary condition reduces to a rigid-wall condition for acoustics, and a fixed condition for elasticity elements.

6.1.8. Slosh

Slosh boundary conditions are applied at free surfaces that are effected by gravity. This type of free surface is typically only important on the surface of a liquid such as water. It contributes to the mass matrix, resulting in “surface” wave modes.

An example syntax for an absorbing boundary condition is given below

```

Boundary
  sideset 7
    slosh = 0.102 // 1.0/9.8 (m/s^2)
end

```

This specifies that sideset 7 is to have a slosh boundary condition. In this case, the slosh coefficient needs to be set to $\frac{1}{g}$, where g is the gravity constant. Thus, for SI units, the slosh coefficient is 0.102. Currently, slosh boundary conditions are only valid for acoustic elements. Applying them to elastic elements will generate an error.

6.1.9. Infinite Elements

In this section, we describe how to use infinite elements for acoustics. These elements serve as both high-order absorbing boundary conditions, and far-field calculators that allow the analyst to compute the solution at far-field points outside of the acoustic mesh. This latter step is a post processing step.

The infinite element specification begins with a sideset on the **Exodus** file of interest. Currently, that sideset has to be an ellipsoidal surface or part of an ellipsoidal surface. Thus, a full spherical surface, hemispherical surface, or a quarter of a sphere would all be acceptable. Infinite element accuracy will degrade if the element surfaces on the boundary do not adequately represent the ellipsoidal surface. The finite element surfaces will be faceted, but enough elements on the boundary are needed to represent the ellipsoidal curvature.

Once a sideset is identified for the infinite element surface, the **boundary** section in the input deck would be modified as follows.

```

Boundary
  sideset 1
    infinite_element
    use block 57
end

BLOCK 57
  infinite_element
  radial_poly = Legendre
  order = 5
  source_origin = 0 0 0
  ellipsoid_dimensions 15 15 30
  neglect_mass = yes
END

```

where block 57 contains the infinite element parameters. The number 57 is arbitrary; the user can pick any number (or name) that is not assigned to a block in the input mesh

Parameter	Description	Options	default
radial_poly	the type of polynomial for radial expansion	Lagrange, Legendre, Jacobi	Legendre
order	the order of the radial basis	0-19	0
source_origin	the origin of the ellipsoid	3 real numbers	0 0 0
ellipsoid_dimensions	radial dimensions of ellipsoid axes	3 real numbers	0 0 0
neglect_mass	indicates whether to neglect infinite element mass	yes or no	yes
correct_mass	whether to correct negative mass terms.	yes or no	yes
useplanelineintersectmethod			

Table 6-134. – Available parameters for the infinite element section.

(Exodus) file. The parameters are summarized in Table 6-134. Currently, only Legendre polynomials are available for the radial basis. The order of the polynomial can vary from 0 to 19. Order 0 corresponds to a simple absorbing boundary condition. Higher orders will be more accurate, but also more computationally expensive. The source point is the location of the center of the ellipsoid that the infinite elements emanate from.

The `ellipsoid_dimensions` parameters indicate the axial dimensions of the ellipsoid in the global coordinate system. They are specified as ellipsoid radii instead of ellipsoid diameters. In the case of a sphere, all 3 parameters are equal and the radius of the sphere. These parameters are currently required, and an error will be generated if they are not specified.

The **`neglect_mass`** keyword indicates whether to neglect the mass matrix contributions from the infinite elements. By default, **`neglect_mass`** is yes. Note that for a spherical surface, the mass matrix contributions from an infinite element are identically zero. However, when numerically generated, small entries will be present in the mass matrix, and thus an option is provided to include these terms in the analysis. Neglecting the mass, **yes**, is recommended in most cases.

Infinite elements only require a specification of a sideset on the surface of interest. No elements need be set up explicitly on this interface. Internally, **Sierra/SD** constructs virtual elements and virtual nodes that define the actual infinite elements, but the analyst need not build a layer of elements on the boundary of the sideset.

The infinite element formulation in **Sierra/SD** uses a Petrov-Galerkin formulation, instead of a standard Galerkin formulation. As a result, nonsymmetric system matrices are encountered with infinite elements. This restricts the solver options to the GDSW solver for time and frequency domains (i.e. **`directfrf`**). Infinite elements can be used either with purely acoustic problems, or with coupled structural acoustics. The formulation is the same, and the GDSW solver is required for the solutions since nonsymmetric matrices are encountered.

6.1.9.1. Far-Field Postprocessing The infinite element formulation allows the analyst to compute the response outside of the acoustic mesh as a post processing step. The response can be computed at any point outside the mesh, and for any time interval. Currently, the **linesample** capability is used to write out the far-field data (see Section 7.5). This data may be written in a readable MATLAB format, which can easily be read in to create plots of the data.

The output will be written to a MATLAB m-file with the name “linedata.m” or “linedata.exo”, depending on which option is selected for output. One file is written per analysis (results are joined analogous to history file output). For example, reading this file in will create vectors **FieldTime** and **displacement** . The acoustic pressure is found in **displacement1** .

We note that the infinite element output in the far-field is always given with respect to some time shift. Details of this are given in the theory notes on infinite elements. The shifted times are included in the **linesample** output for the analyst to use. These allow for plotting the time histories against the appropriate time vectors.

The shifted time output is available in the **linesample** output in a nodal array called **FieldTime**. The dimension of the **FieldTime** array is the same dimension as the acoustic pressure output, since each node in the **linesample** output has its own **FieldTime** array. One **FieldTime** array is available for each sample point in the **linesample** output.

The following command in MATLAB will plot the pressure for the first sample point.

```
FieldTime = nvar09;  
pressure = nvar01;  
plot(FieldTime(1,:), pressure(1,:))
```

The **linesample** points defined in the **linesample** file can contain points that are both inside and outside of the acoustic mesh. For points that are inside of the mesh, the **FieldTime** array for each node will be identically equal to the time array. For points outside of the acoustic mesh (i.e. inside of the infinite element mesh), the **FieldTime** values will be larger than the corresponding time values in the **Time** array, since the acoustic waves will take additional time to reach these far-field points.

6.1.10. Perfectly Matched Layers

Perfectly Matched Layer (PML) elements enforce acoustic baffle boundary conditions. A detailed explanation of theory and implementation of our PML formulation is available.¹⁷ These elements serve as an absorbing boundary condition for outgoing acoustic waves, much like infinite elements. Unlike infinite elements, they are linear elements, and do not exhibit the large matrix condition numbers and convergence issues that can accompany infinite elements. Table 6-135 summarizes the parameters.

While PML are a separate block of elements in the finite element boundary, in an input file they are treated like a boundary condition. The **boundary** section is set up as follows:

```
Boundary
  sideset 1
  pml_element
  use block 217
end
```

```
Block 217
  pml_element
  pml_thickness 1
  stack_depth 1
  source_origin = 0 0 0
  ellipsoid_dimensions 15 15 30
  loss_function = polynomial
  loss_params 0 960 960 0
end
```

where block 217 contains the PML parameters. The number 217 is arbitrary; the user can pick any number (or name) that is not assigned to a block in the input mesh (Exodus) file. The output mesh will contain PML elements in block 217.

Table 6-135. – PML Element Parameters.

Parameter	Description	Options
pml_thickness	length of PML extrusion from boundary	Real
stack_depth	number of elements through PML thickness	Integer
source_origin	the origin of the ellipsoid	3 real numbers
ellipsoid_dimensions	radial dimensions of ellipsoid	3 real numbers
loss_function	type of function describing PML decay	singular or polynomial
loss_params	constants in function describing PML decay	4 real numbers

For PML, a loss function is a definition of the rate of decay of the outgoing wave. While the choice of the loss function $\sigma(d)$ is discussed in the literature, [11,12,44](#) papers in the literature use a range of formulations and implementations, and it is still unclear what the best choice is for any given problem. Typically, the loss function starts at a low value (often zero) to minimize numerical reflections, and increases at an increasing rate to maximize the loss terms near the outer boundary. One option is the polynomial loss function, that includes the constant, linear, quadratic, and cubic terms, with four parameters that define the loss function

$$\sigma(\xi) = c_1 + c_2 \frac{\xi}{t} + c_3 \frac{\xi^2}{t^2} + c_4 \frac{\xi^3}{t^3} \quad (6.2)$$

where c_1 , c_2 , c_3 , and c_4 are specified in the input file, ξ is the distance along the normal from the Gauss point to the inner ellipsoid boundary, and t is the total thickness of the PML layer. The loss function is normalized such that changing the thickness of the PML layer does not change the maximum value of σ . Another option is the singular loss function.¹² (equation (6.3)), which is unbounded at the outer boundary.

$$\sigma(\xi) = \frac{c_1}{t - \xi} \quad (6.3)$$

PML elements can be extruded from either Tet4 or Hex8 meshes. Note that Tet4 is the default; for PML elements extruded from Hex8 meshes, it is necessary to specify **Hex** by modifying the **boundary** section as follows:

```
Boundary
  sideset 1
  pml_element
  use block 217
  hex
end
```

6.1.10.1. Limitations PML is only supported for certain formulations and element types. The implemented PML formulation only applies in the frequency domain, and will throw a fatal error for eigen or time domain solutions. Additionally, PML is only supported for three and four noded boundary faces, limiting the exterior of the acoustic domains to Tet4, Wedge6, and Hex8 elements.

6.1.11. Periodic Boundary Conditions

Periodic boundary conditions can be applied in Sierra-SD using **begin-periodic** block. Built on the tied data algorithm, periodic boundary conditions are currently supported for structural/solid surfaces, with the displacement on the side B surface specified relative to the displacement of the side A surface. The A and B surfaces within a begin-periodic block may not share nodes, but surfaces across various begin-periodic blocks can share nodes (e.g. standard representative volume element (RVE) model would have three begin-periodic blocks, with the surfaces in each block sharing nodes with surfaces in other blocks, at the RVE edges). The surfaces can be curved, but the user must ensure that they have the same geometry, translated in space. Meshes on the two surfaces need not match. Matching meshes are however recommended whenever feasible, to ensure accuracy of the computed stresses on the surfaces. See discussion in Tied Surfaces section above regarding inaccuracies in stresses at the tied surfaces with mismatched meshes. The underlying algorithm results in non-homogeneous MPCs from a node-to-face constraint algorithm. Naturally, for matching meshes, these MPCs reduce to node-to-node MPCs.

Like a **tied data** block, each begin-periodic block represents a *single* pair of opposing sides connected by periodic boundary conditions, as shown in the example below.

```

begin-periodic
  side A = 12
  side B = 18
  name "PBCx-12-18"
  geometric offset = 10.0 2.0 3.0
  search tolerance = 1e-7
  Ux = 0.5
  Uz = 0.25
end

```

```

begin-periodic
  side A = 1
  side B = 6
  name "PBCy-1-6"
  search tolerance = 1e-7
  Ux = 0.5
  Uy = -0.5
end

```

In defining surfaces, care must be exercised to *ensure that the normal vectors of the two surfaces point towards each other*. The keyword “geometric offset” represents the distance from the A surface to the B surface in x, y and z directions. Similar to the TIED DATA block, the keyword “search tolerance” represents the normal distance from a node on the B surface, shifted by the geometric offset, to the face on the A surface. See Figure 8-72 for further details. The keywords “Ux,” “Uy” and “Uz” represent the displacement of the B surface relative the A surface in x, y and z directions respectively. The description of all the parameters are shown in Table 6-136.

6.1.12. Usage Guidelines

The user is referred to the verification manual for examples of periodic boundary conditions(PBC), one with unidirectional PBC and the other with multi-directional PBC, and their application to simulating periodic volume elements (PVE). Here, we provide some guidelines and cautionary remarks associated with PBC and PVE.

1. *Matching Meshes* The two sidesets that are connected by PBC do not need to have matching meshes, but are recommended whenever feasible. Such matching meshes can be built using various tricks at the mesh-generation stage. For example, for symmetric PVEs (with deterministic microstructure), one could split the PVE using plane(s) of symmetry, meshing one part and mirroring the meshed volume to get the complementary volume, eventually building the mesh for the overall volume with matching meshes on the opposite side.

2. *Non-matching Meshes:* When the microstructure is random or lacks symmetry, PBC may need to be applied to connect non-matching meshes. Such situations may encounter

local oscillations in stresses on the surface, associated with the underlying node-to-face contact strategy. These oscillations are expected to decay quickly going into the volume, making the volumetric average more accurate. Since micromechanics modeling often involves average stresses, the error in homogenized global material properties is expected not to be significant provided that the representative volume element (RVE) is large enough (which will also be a requirement from the standpoint of representing random microstructure).

3. *Geometric Offset*: Theoretically, geometric offset should not be needed in the periodic-boundary block. However, whenever the opposite faces have non-matching meshes, there can be errors in automatic computation of offset, which involves centroid computation of the sideset that may involve discretization errors. Given this, whenever feasible, it is recommended that the geometric offset be provided in the input. Note that this comment applies only to non-matching meshes; automatic computation of geometric offset would not have any errors for matching meshes.

4. *Imposition of Homogenized Strain*: The user is referred to the verification manual (problem on PVE) for details of imposing global strain tensor through multi-directional PVE.

Table 6-136. – Parameters for Periodic Boundary Conditions.

Parameter	type	description
Name	<i>String</i>	name of periodic boundary condition block defaults to <i>periodicBC</i>
side A	<i>Integer</i>	sideset A id
side B	<i>Integer</i>	sideset B id
Geometric Offset	<i>3 Reals</i>	x,y,z components of B sideset location relative to A sideset automatically computed if not explicitly specified (explicit specification is recommended for non-matching meshes)
Search Tolerance	<i>Real</i>	search tolerance normal to the face defaults to 1e-8
Ux	<i>Real</i>	x-displacement of B sideset relative to A sideset default value is 0.0
Uy	<i>Real</i>	y-displacement of B sideset relative to A sideset default value is 0.0
Uz	<i>Real</i>	z-displacement of B sideset relative to A sideset default value is 0.0

6.2. *Exodus Mesh Boundary Condition Input*

Several boundary conditions may be set to values specified in the mesh geometry file or input sources. This is used to **hand-off** loads an earlier finite element analysis, even a previous **Sierra/SD** simulation. The pressure determined in another code, and output to a

sideset and can put input on the sideset, determining the right-hand side in a **Sierra/SD** analysis.

The **Exodus** file input data is defined at one or more time slices that do not necessarily correspond to **Sierra/SD** time steps. Figure 6-65 summarizes the rules for interpolating the input time slides to **Sierra/SD** time steps.

The value of a boundary condition at a time ...

- ... before any input **Exodus** time steps is the initial input value.
- ... between two input **Exodus** time steps is determined by linear interpolation.
- ... after the last input **Exodus** time step is the last input **Exodus** time step.

An example of interpolating Exodus data do **Sierra/SD** time steps is shown in Figure 6-65.




Figure 6-65. – Example of Interpolation of Exodus Data to Analysis Steps.

6.2.1. SpatialBC Functions

This appears to be a capability that was added for a specific purpose that was mistakenly exposed to users.

The Spatial boundary condition function is used to set a boundary conditions from the input **Exodus** mesh geometry file. It resembles the **randomlib** function 2.8.8. The difference is that the **randomlib** function uses a nodeset associated with the specified sideset. With Spatial boundary condition, the nodeset is specified directly.

The variable input from the **Exodus** file is specified with the parameter `exo_var`, followed by either `scalar` or `vector`, sets the variable input from the **Exodus**.

In the example the input scalar is acceleration in the Z direction:

```
boundary
  nodeset NS_top
  function zFromMesh
  accelz = 1
end
function xFromMesh
  type = spatialBC
  nodeset NS_top
  exo_var scalar AccZ
end
```

Here `AccZ` is the exact name of a field on the Exodus nodeset. Exodus field names are case insensitive.

Sierra/SD has only a couple regression tests of `SpatialBC`, and each test uses exactly the same boundary condition. All of them result in a confusing and misleading warning about missing sidesets. The vector input feature is completely untested.

6.2.2. Input an Acoustic Point Source from a Volume

The 3 tests of `ReadNodal` exercise the capability to input a scalar acceleration from the whole mesh (body) or an element block. In acoustic point source analysis volume velocities may be input from an **Exodus** file using a `ReadNodal` function. In theory, the **Exodus** file contains the first or second derivatives of the volume velocity at the corresponding times. The name of the field input from the **Exodus** file is specified with the `exo_var` keyword. Table 6-137 lists the run time parameters for `ReadNodal` functions. An example is provided in input 6.6. The keyword `exo_var` must be followed by two keywords, specifying first whether the data is a scalar or a vector, and second specifying the name of the variable on the **Exodus** database. However the vector input feature is untested. Thus, Table 6-137 specifies that a scalar variable with the name **volume_acceleration** should be available on the **Exodus** database. The `interp` parameter is the same as was described for the `randomlib` functions in section 2.8.8. It specifies the type of temporal interpolation.

`ReadNodal` determines what, if anything, to read by accepting the first match found. One after another, it searches for 3 things, and errors out if the last search come up empty. The first search is for a nodal variable with the specified name (i.e., **volume_acceleration** in the example above). Next it looks for an element variable with the same name. Finally it searches for a face variable.

Table 6-137. – ReadNodal function parameters.

Keyword	Values	Description
type interp	ReadNodal	required to specify function temporal interpolation scheme none=nearest linear=linear interpolation
exo_var	scalar volume_acceleration	either scalar or vector, followed by variable name

```
function 55
  type=ReadNodal
  interp=none
  exo_var scalar volume_acceleration
end
```

Input 6.6. Example ReadNodal Function Specification

6.2.3. Input an Acoustic Point Source from a Node Set

Input volume velocities for acoustic point source analysis from the input **Exodus** mesh geometry file with the **ReadNodalSet** function. Velocities are input at each time specified in the **Exodus** file. Spatially dependent velocities are input at each nodeset node.

Table 6-137 lists run time parameters for **ReadNodalSet** functions. The only difference from the **ReadNodal** parameters is the required **nodeset** parameter. The force is applied over a single nodeset of the model. This nodeset must match the definition in the load section.

```
function 55
  type=ReadNodalSet
  interp=none
  nodeset=NS_top
  exo_var scalar volume_acceleration
end
```

Input 6.7. Example ReadNodalSet Function Specification

6.2.4. ReadSurface functions

A **ReadSurface** function reads in data from either the entire **Exodus** file, or a nodeset or sideset that covers a surface of interest. If a set is specified in the **function** block, then data corresponding to that set is read in from the **Exodus** file. Otherwise, the variable is read from the entire mesh as a nodal variable (rather than a nodeset variable).

Once the data is read, **Sierra/SD** integrates the data over the surface to create a time and spatially-varying forcing function. One difference between this function and the **ReadNodal** function is that the data from **ReadSurface** is used in a surface integration to generate the load, whereas the data from **ReadNodal** does not need to be integrated, and thus can be inserted directly into the force vector.

This function is used to read in surface velocities or accelerations which are used as a boundary condition for acoustic analysis. It can also be used for applying time and spatially-dependent pressure or traction loads on a structure. For this case, the load output variable currently only outputs element data for values read from a sideset. Nodal values, like those used in **randomlib** functions, are output as 0. There is 1 and only 1 test of read surface with a traction load.

As currently implemented, the **ReadSurface** function operates only by reading data from an external exodus data file. The name of the variable to read from the **Exodus** file must be specified in the input deck using the **exo_var** keyword. Also, the variable must be specified to be a scalar or a vector, using the syntax given in input 6.8. Pressures require a scalar variable, and tractions require a vector variable. An example for **ReadSurface** functions is given in input 6.8. A **sideset** matching the corresponding Load sideset is required. The **interp** selected the temporal interpolation algorithm as was described for the **randomlib** function. The default option, linear interpolation, is the only option available.

In input 6.8, the keyword **exo_var** specifies the type of data, such as **vector** or **scalar**, and the name of that variable in the **Exodus** file. In the case of a vector, the name of the variable as given in the input deck should be the base name of the variable, without the suffix of 'x', 'y', or 'z'. For example, for the data given in input 6.8, a vector nodal variable with name 'vel' should be available in the **Exodus** file. Thus, the data in the **Exodus** file would have names **velX**, **velY**, and **velZ**. In the case of scalar data, the base name given (i.e. **vel** in input 6.8), should match exactly the name of the nodal variable in the **Exodus** file.

```
loads
  sideset 1
    traction 1 1 1
    function 55
end
function 55
  type=ReadSurface
  interp=linear
  exo_var vector tracLoad
```

```
end
```

Input 6.8. Example ReadSurface Function Specification

ReadSurface functions ignore distribution factors when used to apply loads. See section [6.3.14.1](#).

6.2.5. ExodusRead functions

A `ExodusRead` function reads in data from elements. Exodus read is currently only available with the acoustic boundary conditions `point__volume__acceleration` and `point__volume__velocity`.

ExodusRead function operates only by reading data from an external exodus data file. The name of the variable to read from the **Exodus** file must be specified in the input deck using the `exo_var` keyword. The acoustic load conditions this function type supports are scalars.

```
loads
  block 1
    point_volume_accel = -1.0
    function = 77
  end
function 77
  type=ExodusRead
  exo_var scalar ddvol
end
```

Input 6.9. Example ExodusRead Function Specification

6.2.6. In Core Transfer Functions

A limited capability exists to perform in core transfer of variables between physics codes via MPI messages rather than **Exodus** files. In this use case both codes are run simultaneously with the results from the load producing code constantly being fed to the load using code. The main use case for this capability is to **hand-off** boundary conditions from one code to another without the need for huge intermediate **Exodus** files with a large numbers of steps.

6.2.6.1. Transfer from Fuego

One field that may be transferred in core is the divergence of Lighthill's tensor from Fuego to **Sierra/SD**. This is a nodal value that can be used as an acoustics load term. An example of the relevant input for this case is:

```
PARAMETERS
  mpmd_transfer_type = fuego
END
LOADS
  nodeset = 1000
  lighthill = 1.0 #Sets load type and scale factor
  function = from_fuego
END
FUNCTION from_fuego
  type transfer
END
```

An example of the job execution syntax for this case is:

```
$ mpirun -n 1 fuego -i Fuego.i : -n 1 salinas -i Salinas.inp
```

6.2.6.2. Transfer from SPARC

Another field that may be transferred in core is the traction load from SPARC to **Sierra/SD**. This is a face-based term that can be used for structural loads. An example of the relevant input for this case is:

```
PARAMETERS
  mpmd_transfer_type = SPARC
  mpmd_transfer_sidesets = 5
END
LOADS
  sideset 5
  traction = 1.0 1.0 1.0 #Sets load type and scale factors
  function = from_sparc
END
FUNCTION from_sparc
  type transfer
END
```

An example of the job execution syntax for this case is:

```
$ mpirun -n 1 sparc -i sparc.i -c 0 : -n 1 salinas -i Salinas.inp
```

6.3. *Loads*

Loading conditions are specified within the **loads** section. The following example illustrates the method.

```
Loads
  nodeset 3
    force = 1.0 0. 0.
    scale = 1000.
    function = myLoadFunction
  nodeset nose
    coordinate 11
    force = 0. -1 0
  nodeset 7
    point_volume_vel = 1
    scale = 1.0
    function = 1      // time history of dV/dt,
                      // where V is the volume of the source
  body
    gravity
    0.0 1.0 0
    scale -32.2
  block wing
    thermal_load
    function = 1
  block 12
    thermal_load
    function = 3
  sideset 7
    pressure 15.0
  sideset tail
    traction = 100.0 20.0 0.0
    coordinate 0
  sideset 13
    acoustic_vel 1.0
    function = 1
  sideset 14
    pressure = 1
    follower=yes
  node_list_file='force.nodes'
    force=1.0 0 0.
    scale = 100.
    function=2
END
```


Loads may be applied to node sets, side sets, blocks, node lists (see Section 6.1.5), or the entire body (in the case of inertial loads). Pressure loads may be applied using side sets. The pressure is always normal to the surface. All loads applications are additive. Forces should not be applied to sidesets.

The components of each load specification are listed in Table 6-138. The syntax followed is to first define the region over which the load is to be applied (either **nodeset** , **sideset** , **block** , **node_list_file**, or **body**). Each such region defines a *load set*. For each such definition, one (and only one) load type may be specified. However, any region definition (except **node_list_file**) may be repeated so that forces and moments may be applied using the same region.

Following the definition of the load type, a vector (or scalar in the case of pressure loads) must be specified, except in the case of a thermal load, where no vector or scalar multiplier is needed. The vector is the load applied in the basic coordinate frame unless a coordinate frame is also specified (see Section 2.7).

6.3.1. Load

Loading conditions for individual cases in a multcase solution are specified within the **load** section. See Section 3.2 for information on **load** specifications for multcase solutions. The **load** section is *identical* to the **loads** section described in the previous section (6.3), except the section begins with the keyword **load** , and a load step identifier is required. The following example illustrates the required input. ⁵

```
LOAD=57
  nodeset 3
    force = 1.0 0. 0.
    scale = 1000.
    function = 2
  nodeset 5
    force = 0. -1 0
END
```

Unlike the **loads** section, there may be multiple **load** sections in the file. Each **load** can then be applied to different cases in the multcase solution by calling out the **load** identifier in the case block of the multcase solution, as described in Section 3.2.

6.3.2. Scale Factors for the Load

The total load *on each degree of freedom* is the product of the load vector, the scale factor, and the nodeset distribution factor found in the **Exodus** file. For pressures and tractions, the load is also multiplied by the area of the face. Note that in some cases the nodeset

⁵The identifier is usually an integer, but any unique string is acceptable.

distribution factor may be zero. ⁶ In that case, the total applied force will also be zero. The Exodus distribution factors can be ignored by setting the parameter **AllowExodusDistFacts** to false.

6.3.3. Sideset Loading

The **pressure** , **acoustic_vel** , and **acoustic_accel** boundary conditions may only be applied to side sets. The total pressure is the product of the scale factor, pressure (scalar) and sideset distribution factors. By default, pressure loads are not follower loads, i.e. pressures are applied in the direction of the undeformed element normal for the entire simulation. The **follower** keyword may be applied to user defined functions if a follower load is required. See section 6.3.6 for follower stiffness specification.

If the pressure loading is not normal to the sideset, the **traction** capability should be used. NOTE: Pressure will act on a surface in a compressive sense, while a traction can be specified as any vector which will act on the **sideset** specified in the direction given by the triple values specified after **traction** . Also, traction loads are applied on the faces of the shell elements in a piecewise manner, i.e., the traction load acting on a face of the element is assumed constant. If the distribution factors on the nodes of the element vary, the average of the load (element per element) is assumed.

Traction loads may be specified in either the *global* coordinate frame (default), or in a coordinate frame projected onto the surface.

If the analyst provides a coordinate frame with the traction definition, then that frame is projected onto the surface of each element. Figure 6-66 illustrates that projection. Note that when a coordinate system is used, there can be a mesh dependence on which direction the forces are applied. Note that the third coordinate of the traction will always be applied along the surface normal, and that the third component of the vector will always correspond to the surface normal (and hence will be applied as a pressure).

⁶Because the nodeset distribution factors are part of the **Exodus** file, and may be difficult to check, errors in the distribution factors are common. Analysts are urged to carefully examine the distribution factors.

Section	Keyword	Parameters
Region (<i>defines application area</i>)	body nodeset sideset block node_list_file	- <i>id/name</i> <i>id/name</i> <i>id/name</i> <i>filename</i>
Load Type (<i>defines application method</i>)	force moment gravity pressure point_volume_vel point_volume_accel acoustic_vel acoustic_accel Lighthill surface_charge traction thermal_load energy_load	<i>val1 val2 val3</i> <i>val1 val2 val3</i> <i>val1 val2 val3</i> <i>value</i> <i>value</i> <i>value</i> <i>value</i> <i>value</i> <i>value</i> <i>val1 val2 val3</i> - -
<i>optional specifications</i>		
Coordinate Frame (<i>for vector loads only</i>)	coordinate	<i>id/name</i>
Scale Factor Multiplier	scale	<i>val1</i>
Function (<i>Required for transient analysis</i>)	function	<i>id</i>
follower	follower	<i>yes/no</i>

Table 6-138. – Load Specification Keywords.




Figure 6-66. – Coordinate Frame Projection for Traction

6.3.4. Spatial Variation

Variation of the load over space is accomplished using node set or side set distribution factors or a function. If these are provided in the **Exodus** file, the load set is spatially multiplied by these factors. The total loading is the sum of the loads for each load set summed over all the load set regions.

6.3.5. Required Section

When prescribed accelerations are applied in the **boundary** section (6.1), they induce a load on the structure. In these cases the **loads** section may serve no purpose, unless an additional external load is applied. In these cases, an empty **loads** block is still needed in the input file. An error is generated if the input file has no **loads** section.

6.3.6. Follower Stiffness

The follower stiffness that corresponding to an applied pressure load may be included. A follower pressure load applied to a structures will “follow” the structure during deformation, always remaining normal to the surface where they are applied. As such, the applied force due to a pressure load depends on the deformed state, and this induces a follower stiffness matrix that contributes to the overall stiffness matrix of the structure.

The boundary where the pressure is applied is specified with a sideset. Also, the magnitude of the applied pressure field must be specified, as shown in the example below. The follower stiffness matrix scales linearly with the magnitude of the applied pressure.

```
Loads
  sideset=1
  pressure = 10.0
  follower=yes
END
```

In the above example, sideset 1 is used to denote the surface where the pressure is applied. The parameter "pressure" specifies the magnitude of the applied pressure field.

6.3.7. Acoustic Loads

The **acoustic_vel** , **acoustic_accel** , **point_volume_vel** , **point_volume_accel** , **Lighthill** loading conditions are specifically designed for acoustic elements, and thus may only be applied to acoustic elements. In all cases, a time function is required that defines either the time or frequency dependence of the loads.

The **acoustic_vel** and **acoustic_accel** keywords specify the fluid velocity and fluid acceleration in the normal direction of the element faces in the sideset, respectively. Note that these are the counterparts to the **pressure** load for structures in the sense that they are Neumann boundary conditions.

We note that the **acoustic_vel** and **acoustic_accel** approaches should yield the same acoustic response, provided that the **acoustic_vel** time function is precisely the time integral of the **acoustic_accel** function. This time integration must include the constant of integration. If the two time functions for **acoustic_vel** and **acoustic_accel** are complementary in this way, the acoustic pressure output from these approaches will be the same up to first order. They are not exactly the same since the time derivative of velocity potential is needed to generate the acoustic pressure for output, and that time derivative is only first-order accurate.

An example of the **acoustic_vel** keyword is given below.

```
Loads
  sideset 1
    acoustic_vel = 1.0
    function = velFunc
END
```

In this case, sideset 1 is given a prescribed normal velocity of amplitude 1, with a time dependence given by function *velFunc*.

Currently, a given load case cannot contain both an **acoustic_vel** and an **acoustic_accel** input. Only one or the other can be specified in a given load case,

though for a multicasel solution the **acoustic_vel** and **acoustic_accel** inputs could be present in separate load cases. We also note that for coupled structural acoustics, only the **acoustic_vel** keyword is applicable. For analysis involving only acoustic elements, either keyword can be used.

The **point_volume_vel** and **point_volume_accel** keywords prescribe an acoustic point source on a nodeset. This force is the product of the fluid density with the first and second derivatives, respectively, of volume of the source. The **function** for the point source contains the time history of the first (for **point_volume_vel**) and second (for **point_volume_accel**) time derivative of volume.

Since the code scales by density in the internal calculations, there is no need to multiply the time history of volume by density to get the acoustic force. Thus, for point sources the **scale** parameter is typically set to 1.0, unless a direct scaling is desired. The units of the input time functions for **point_volume_vel** and **point_volume_accel** are volume per unit time and volume per unit time squared, respectively. The density need not be multiplied by these functions, since the code is already dividing by density internally (see the theory notes on structural acoustics for a more detailed discussion.)

Currently, the point acoustic source is only implemented for the time domain (transient) calculations.

The **Lighthill** keyword prescribes the divergence of the Lighthill tensor at nodes. The double divergence of the Lighthill tensor provides a source term for noise generation in the pressure formulation of acoustics. The divergence of the Lighthill tensor is a time varying vector quantity that can only be applied as a nodeset load using the readnodalset function. The Lighthill load is only implemented for the time domain (transient) calculations. The Lighthill load is only valid for the pressure formulation of acoustics and can only be used with **acoustic_accel** loading. For Lighthill loading, the **scale** parameter is typically set to 1.0, unless a direct scaling of the load being read in is desired. The **outputs** keyword **acousticlighthill** outputs the acoustic Lighthill source term.

```
Loads
  nodeset 1
  point_volume_accel = 1.0
  function = accelFunc
end

Function accelFunc
  type LINEAR
  name "volume_acceleration"
  include inc/volume_acceleration.inp
end

Outputs
  acousticlighthill
end
```

Input 6.10. Example Input for Point Acoustic Load. In this case, nodeset 1 would consist of a single node, and the file "volume_acceleration.inp" would contain the second time derivative of volume velocity of the source, with units of volume per time squared. Note that the amplitude of the point source is taken to be 1.0, and that it does not include the density multiplier.

The sign conventions of the **acoustic_vel** , **acoustic_accel** , **point_volume_vel** , and **point_volume_accel** keywords are important. For the **acoustic_vel** and **acoustic_accel** cases, the equations of motion are given by,

$$\frac{1}{c^2}\ddot{p} - \Delta p = - \int_{\Gamma} \rho q (a \cdot n) d\Gamma \quad (6.4)$$

or, in discrete form,

$$M\ddot{p} + Kp = f \quad (6.5)$$

where ρ is the density, q is the surface shape function, a is the acceleration vector on the surface, n is the normal to the surface, and Γ is the portion of the surface where the loading is defined. M , K , and f are the mass, stiffness, and discrete force vectors. We denote $a \cdot n = a_n$ as the normal component of acceleration. We also note that this force has a negative sign in front of the integral, which comes from the variational formulation. This implies an inverse relationship between surface acceleration and acoustic pressure. Thus if the acceleration is oriented in the same direction as the normal, then a_n will be positive, and thus the total force vector will be negative. Intuitively, this makes sense, since if the acceleration is in the same direction as the surface normal, mass will be ejected from the acoustic space, causing a decrease in pressure. Conversely, if the acceleration is oriented in the opposite direction as the surface normal, then a_n will be negative. This will cause the total force vector to be positive, resulting in a positive pressure. This makes sense, since in this case mass will be added to the acoustic space, causing an increase in pressure.

For the **point_volume_vel** and **point_volume_accel** boundary conditions, the equations of motion are given by

$$\frac{1}{c^2}\ddot{p} - \Delta p = \rho \frac{\partial^2 V}{\partial t^2} \delta(x - x_0) \quad (6.6)$$

or, in discrete form,

$$M\ddot{p} + Kp = f \quad (6.7)$$

where $\frac{\partial^2 V}{\partial t^2}$ is the second derivative of the volume change with respect to time, and $\delta(x - x_0)$ is the Dirac delta function that makes the term zero everywhere except where $x = x_0$. We note that V is the volume of fluid added to the surrounding acoustic space, *not* the volume of the point source per se. Thus, the sign of the acoustic pressure will be related to the sign of $\frac{\partial^2 V}{\partial t^2}$. A positive $\frac{\partial^2 V}{\partial t^2}$ would result in a positive acoustic pressure, implying that fluid mass is added to the surrounding acoustic space. Conversely, if $\frac{\partial^2 V}{\partial t^2}$ is negative, mass will be subtracted from the acoustic space, and thus a negative acoustic pressure will result.

The previous examples involved spatially constant functions of time. Acoustic boundary conditions with spatially-varying functions of time are supported through the **ReadNodal** and **ReadSurface** functions as described in Sections 6.2.2 and 6.2.4 respectively.

6.3.8. Thermal Loads

The **thermal_load** option is used in conjunction with a spatial temperature specification for the structure. The temperature distribution can either be specified via the input **Exodus** file, or on a block-by-block basis, as described below. Based on the temperature distribution, a thermal load is computed and then applied to the structure.

If the solution method is selected to be statics, the **thermal_load** option will provide the thermal load necessary to solve the thermal expansion problem. If the solution method is transient dynamics, the same thermal load will be applied as in the statics case, but modulated by the function that is specified below the **thermal_load** keyword. This corresponds to a thermal shock analysis. Thus, for a transient dynamics problem that includes damping, and with a function that is equal to 1.0 for all time, the transient analysis would eventually converge to the same solution as obtained in the statics analysis, which would be the solution from a classical thermal expansion analysis. On the other hand, for a transient dynamics problem with a **thermal_load** in which the associated time function is not equal to 1.0, the thermal load will be scaled according to that time function. For example, in the case of a mesh that has block-by-block values of temperature $T\_current$ specified in the input deck, and a thermal load function that ramps up from zero to one, the actual thermal load applied to the structure will be multiplied by that time function. In this case, the full thermal load will only be seen after the ramp in the time function is completed.

If it is desired to apply a thermal preload to a structure, we generally recommend using a **statics** analysis rather than a **transient** analysis, since in the latter case the preload that will be computed will be a dynamic preload that will oscillate about the static preload solution. If damping is used, this dynamic preload will converge to what would be obtained from using a **statics** analysis. However, in some cases such as when rigid body modes are present, a **transient** analysis may be the only option for applying the preload.

The temperature field can either be read from an **Exodus** file, which would typically be the result of a thermal analysis, or it can be specified on a block-by-block basis in the input deck. For temperature fields that change from element to element, the temperatures must be read in from the **Exodus** input file. For more uniform temperature distributions, it is more efficient to specify them block-by-block in the input deck. Note that when using thermal loads, the temperature data is expected to either be in the mesh (exodus) files, or specified using the input deck (i.e. block-by-block). If temperature is specified in the **Exodus** file and on a block-by-block basis in the input deck, then the input deck values take precedence.

Output stress and strain. Sometimes it is of interest to output the stress after a thermal load analysis. In this case, the stresses that are output to the **Exodus** file will be

the *mechanical* stress, rather than the combined thermal-mechanical stress. There is a known bug in the way that thermal stresses are computed, particularly when temperature comes from the **Exodus** file. If thermal stresses are needed extreme care should be taken. Mechanical stress is the same as elastic stress. **Strain**, **elastic_strain**, and **thermal_strain** output are all available.

Input deck syntax. If temperatures are specified using the input deck, then each block must be given its own temperature. In the example below, there are 2 blocks, and each is given a different temperature.

```
BLOCK 1
  material 1
  T_current 100
END
BLOCK 2
  material 2
  T_current 200
END
```

Note that if $T_{current}$ is specified for some blocks and not for others, the code will error out.

When temperatures are read in from the **Exodus** file, the material properties can be specified as temperature-dependent. This implies that each element will have different material properties. More details are given in the section on temperature-dependent material properties.

For thermal statics or thermal transient analysis, each material block must be given two additional parameters, the reference temperature, $T_{ref} = \mathbf{Tref}$, and the coefficient of thermal expansion, $\alpha_t = \mathbf{alphat}$. These parameters are defined via the thermal strain, which is given by

$$\epsilon_{thermal} = \mathbf{alphat} (T_{current} - T_{ref}) \quad (6.8)$$

An example is the following.

```
MATERIAL 1
  E 10e6
  nu 0.3
  Tref 300.0
  alphat .001
  density 0.1
END
```

The defaults for **Tref** and **alphat** are both 0.0. This implies that if they are not specified, then the material will not contribute to the thermal analysis (see equation 6.8).

Note: currently, isotropic thermal strain is supported for isotropic, isotropic viscoelastic, and anisotropic materials.

Shell and beam type elements are not supported in thermal strains. If a material with a coefficient of thermal expansion is used in a shell, beam, or other unsupported element **Sierra/SD** will generate an error.

The default **Exodus** file labels for the temperatures are shown in the table below. This is the default variable format for **Sierra/SD**. However, it is also possible to read in element variables and variables of different names. Using the keyword **thermal_exo_var** in the **parameters** section (2.3) allows you to specify the name of the temperature variable in the **Exodus** file. A fatal error is generated when **thermal_load** is used with **energy_exo_var**. A nodal variable of this name is expected. If there isn't one, an element variable is used. And if no element variable of the given name is found, an error will be generated.

Name	Definition
TEMP	the nodal temperature

The **thermal_load** load case can be used in a multcase solution method. In that case, the stresses and internal forces from the thermal analysis are used as initial conditions for the next case. For example, for a fixed-fixed cantilever beam that is subjected to a uniform temperature increase, the beam will undergo a stretch due to the thermal static analysis, and will have residual stresses. If this beam were then subjected to an eigen analysis in a subsequent case, the modes would be modified due to the geometric stress stiffening. Conversely, for a fixed-free beam, there would be no residual stresses and thus no effect on subsequent cases. Note that the displacements from thermal analysis are not carried over to subsequent cases. Thus, to get the total displacement from a thermal analysis followed by transient, one would need to add the displacement results from the two cases separately.

When temperature is read from the **Exodus** input mesh using either the default temperature file name or the **thermal_exo_var** keyword a constant temperature can be read from a single time step with the keyword **thermal_time_step** in the parameters block. Alternatively a time variant temperature can be updated periodically from the mesh file with use of the keyword **nUpdateTemperature** in the solution case input. The **nUpdateTemperature** keyword defines how often (every n steps) temperature should be updated. In a transient run updating of temperature is moderately expensive thus it may be advantageous to only update temperature at a finite step interval if temperature is changing slowly.

When a new temperature value is read that temperature is used immediately to update the applied thermal strain. Additionally, the updated temperature can be used with the **nUpdateDynamicMatrices** keyword to use apply the changing temperature to updated material properties. When reading transient temperature data from the input mesh the closest **Exodus** time step at or below the current time step is used. Temperature does not interpolate between **Exodus** steps.

The **thermal_time_step** keyword must be specified in the **parameters** block, to specify which time step of the previous thermal analysis should be used to extract temperature data. The following gives an example.

```
Parameters
  thermal_time_step 10
  thermal_exo_var "TEMP"
end
```

The **Exodus** files can contain multiple time steps of temperature data. The user can select which time step is to be used for defining temperature data in **Sierra/SD**, using the keyword **thermal_time_step** . In this example the tenth time step will be read in from the **Exodus** file. The default value for the **thermal_time_step** is 1.

The **nUpdateTemperature** keyword is placed in the solution case to specify how often to update the temperature that is read in.

```
SOLUTION
  transient
  nUpdateTemperature 5
END
```

Here a new temperature is read in every 5th time step. If the transient solution specifies the last time step from the thermal analysis then the final temperature will be used.

The next example presents input for thermal statics analysis.

```
SOLUTION
  statics
END

Parameters
  thermal_time_step 10
end

Loads
  body
    thermal_load
END
```

6.3.9. Energy Deposition Input and Loads

Input from energy deposition are similar to thermal loads (section 6.3.8). These loads are specified when energy is deposited directly in the structure as with an X-ray deposition. For consistency with other applications, the energy is defined as *specific energy*, i.e. the energy per unit mass. Such direct energy deposition is converted to a change in temperature after which thermal strains and loads are computed exactly as for the **thermal_load** approach.

Energy is converted to a change in temperature using the specific heat of the material (see Section 4.4.8).

$$\tilde{E} = C_v \Delta T$$

where $\tilde{E}$ is the specific energy of the body, C_v is the specific heat capacity for constant volume, and ΔT is the change in temperature.

The energy load is specified using the keyword **energy_load**. All other parameters are identical to *thermal_load*. Note that by the nature of these loads there is often an exponential decay in energy as a function of depth. For this reason, it is advantageous to specify the loads at Gauss points, particularly when using higher order elements. Energy loads can also be specified at nodes or centroids. Nodal energy loads should be avoided because it is not clear which materials specific heat to use when converting an energy load to a temperature change for nodes on the interface between two materials. For this case **Sierra/SD** will use the specific heat of the material that is processed last (which is not typically what is required by the analyst). **Sierra/SD** generates a fatal error if the specific heat is not specified for a material with an energy load.

Energy may also be used as an input for thermally dependent material properties. To ensure that the energies are converted to temperature before determining the material properties, identify the variable name from the **Exodus** file with the **energy_exo_var** and **energy_time_step** keywords, rather than the **thermal_exo_var** and **thermal_time_step** keyword. Using **thermal_exo_var** with **energy_load** generates a fatal error.

6.3.10. Consistent Loads

The loads for every 3-D and 2-D element is calculated in a consistent fashion when a pressure load is applied. For more details on the implementation, see the programmer's notes. It is important that consistent loading be used. This is especially true for shell elements where the consistent loading is required to properly apply rotations.

6.3.11. Pressure_Z

Depth dependent pressure loads may of course be applied using a user defined function. To simplify this loading condition, depth dependent pressure may also be applied using the

pressure_z keyword. An example is shown in input 6.11. This loading is applied only in the basic coordinate frame, and the analyst must specify that the pressure is either “below” or “above” an offset to the coordinate axis. The pressure is always proportional to the depth. In the example of input 6.11, the pressure is zero at $x = 5$, 10 at $x = 4$, 20 at $x = 3$. At depths above the “waterline”, the pressure is zero. Any of the basic coordinate directions (x , y , or z) may be used as a reference.

```
// depth dependent pressure for a waterline at x=5.
Loads
  sideset 2
    pressure_z 10.0 below x = 5
  sideset 20 // air
    pressure_z 1e-4 above x = 5
END
```

Input 6.11. Depth Dependent Pressure Load Example. This load section applies a pressure to sideset 2 which is proportional to the distance below $x = 5$.

6.3.12. Surface Charge

For electro-mechanical coupled materials such as dielectric or piezoelectric materials, surface charges may be applied to a specified sideset using the keyword **surface_charge** in the loads block. The surface charge is mechanically analogous to a pressure, where surface charge represents a charge per unit area and is applied only to the voltage degrees of freedom. Surface charges can be applied for direct frequency response and transient solution methods. An example is provided below.

```
Loads
  sideset 1
    surface_charge = 1 // scalar multiplying the specified function
    function 1 // surface_charge function
END
```

Input 6.12. Surface Charge Example

6.3.13. Static Loads

Static loads only require the definition of the load region and load keyword (e.g. force) with its accompanying parameters. A function may be used instead. In this case, the function will be evaluated at time $t = 0$.

6.3.14. Time Varying Loads

Additional options provide the capability of varying the load over time. The **load** options include,

- **scale** with one parameter provides a scale factor to be applied to the entire load set. Only one scale may be provided per load set.
- **function**. A time varying function may be applied by specifying a function ID. Only one function may be applied per load set. The function is defined in the **function** section (see section 2.8 on page 84). The loads applied at time t for a particular load set will be the sum of the force or moment vectors summed over the nodes of the region and multiplied by the scale value and the value of the time function at time t .

NOTE:

If no function is applied for a particular load, then the function is defined as 1.0 for all time. All loads will be applied to the transient solution, regardless of whether an explicit time function is defined.

6.3.14.1. Reading Loads from Exodus Data Loads may be read in from previous analyses when stored in the input exodus file. These are read using an appropriate function. See sections 6.2.2, 6.2.3, and 6.2.4 for functions which read data on nodal values, on a node set and on a surface respectively.

Note that exodus read functions do not trigger follower stiffness calculations when used as follower loads. A warning is issued when the follower stiffness calculation is skipped.

Also, note that exodus read functions ignore distribution factors on sidesets. A warning is issued in this case when the distribution factors are ignored.

6.3.15. Random Pressure Loads

Input for random loads can be complicated, though the loads are not uncommon and are important for many applications. ¹ This type of random pressure loading is developed for use of direct transient loading typical of a turbulence load on a hypersonic vehicle. Throughout the development, we maintain a concept of flow direction, and correlation distances that may be different in flow and transverse directions. By computing the random pressure fields as part of the time evolution, we avoid the need to compute these complex quantities before the run. The approach requires a linear solve at each solve to compute the loads.

¹A hypersonic vehicle is a prime example of a random loading. Turbulence provides a time varying loading which has a limited spatial and temporal correlation on the surface of the hypersonic vehicle.

The most general type of input is the correlation matrix, which is the inverse Fourier transform of the spectral density matrix. The **RandomPressure** load option provides a simplified means of specification of the loading. The material in this section is consistent with and builds on Section 2.8.7.

Subsection Random Pressure Loading section Loads and Materials of the Theory Manual⁵² describes the approximations involved in the implementation. These approximations are summarized in Figure 6-67.

The simplified correlation matrix is not general, but may be useful for a large class of problems. It has the following limitations.

1. The system must be time stationary.
2. The correlation function must be separable (a product of temporal and spatial correlations).
3. The same PSD shape must apply throughout the entire hypersonic vehicle body. The PSD may be scaled as a function of z , but there may be no change in the shape.
4. The PSD must have a cutoff. The time integration must occur above this cutoff frequency.
5. By default, the temporal function is represented by a *sinc* function. This may be replaced by a user defined temporal function.

Figure 6-67. – RandomPressure Loading Approximations.

The random loading is a component of the loads section. An example is shown here, and described in Table 6-139.

```
Loads
  sideset 22
    RandomPressure
      correlation_length_z = 2.0 // required
      correlation_length_r = 0.67 // required
      cutoff_freq = 16.8 // required
      correlation_function = 20 // defaults to sin(x)/x
      psd_scale_function = 10 // defaults to Sigma=1
      NTimes = 5 // defaults to 5
      coordinate 1 // defaults to basic frame
      MinimumNodalSpacing = 1.0e-5 // defaults to 1.0e-8
      NumberOfInitializationSteps = 100 // defaults to 5
END
```

Details for the parameters to the correlation matrix input are described below.

Parameter	Type	Default	Comment
correlation_length_z	real	required	spatial decay in flow direction
correlation_length_r	real	required	spatial decay orthogonal to the flow
cutoff_freq	real	required	cutoff frequency
correlation_function	string	$\frac{\sin(t\omega_c)}{t\omega_c}$	defaults to basic frame smallest allowable inter-node spacing random number seed iterations to improve initial spatial distribution
psd_scale_function	string	$\Sigma(z) = 1$	
NTimes	int	5	
coordinate	string	0	
MinimumNodalSpacing	real	$1.e - 8$	
Random_Seed	int	ignore	
NumberOfIniti... ...alizationSteps	int	5	

Table 6-139. – Random Pressure Inputs.

correlation_length_z Spatial decay in the flow direction, L_z . The flow direction is the Z axis of the coordinate frame. The correlation function $C(\Delta Z)$ is proportional to $\exp(-\Delta Z/L_z)$, where ΔZ is the distance between two points in the flow direction.

Correlation_Length_R Correlation_length_r is the spatial correlation distance in the radial or transverse direction. The correlation function is proportional to $\exp(-\sqrt{(\Delta x^2 + \Delta y^2)}/L_r)$.

Cutoff_freq The cutoff frequency, F_c is important to the operation of the RandomPressure algorithm. No energy may be found in the PSD above this frequency. The time integrator may not sample the system lower than this frequency, i.e. $dt < 1/F_c$.

NTimes The matrix is proportional to the number of time values assembled, and affects the interpolation as described in the Theory Manual section Loads and Materials subsection Alternative Derivation Based on Lagrange's Equations subsubsection Separation of spatial and temporal components

Typically, only a small number of terms are required. Note that there are $2 * NTimes + 1$ terms in the sum, and the dimension of the correlation matrix grows commensurately. The number may depend on the interpolation time step and on the shape of the PSD. Default=5 (which produces 11 terms in the sum).

CORRELATION_FUNCTION The temporal time function, whose argument is $(t_1 - t_2)$. By default this function is $\sin(x)/x$, with $x = \pi F_c(t_1 - t_2)$. It must be an even function of the argument.

PSD_SCALE_FUNCTION provides a means of scaling the power spectral density as a function of flow direction. This type of input requires that the PSD have the same shape at all locations, but the value may be scaled. Scaling the PSD effectively scales the standard deviation of the pressure. Default is no scaling. The function must be positive for all values of the coordinates.

coordinate is an optional coordinate frame that is used to define the flow direction. The $\tilde{Z}$ component of that frame is the direction of flow. By default, the basic frame is used.

MinimumNodalSpacing Some models can contain co-located nodes on the surface where the random pressures are to be applied. This can cause the correlation matrix to be singular, since the repeated nodes would result in two identical rows in the correlation matrix. The **MinimumNodalSpacing** keyword allows the analyst to specify the smallest inter-node spacing (absolute) that is allowed on the surface where the random pressure is being applied. Any nodes that are closer than that tolerance will be treated as identical in the correlation matrix manipulations. The **Exodus** file and corresponding nodal output will not be changed. This will avoid a singular correlation matrix, but does not alter the mesh database.

NumberOfInitializationSteps Initially the pressure spatial distribution may be too correlated. Mesh resolution exacerbates this issue. Increasing the **NumberOfInitializationSteps** mitigates the issue. Each initialization step requires about as much CPU time as an implicit time step. Default=5 (values below 5 are discouraged).

OMEGA_C Deprecated. Use Cutoff_freq.

ALPHA_Z Deprecated. $\alpha_z = 1/L_z$.

BETA_T Deprecated. $\beta_t = 1/L_R$.

The computation of the random pressure loads depends on matrix factorizations (subsection Random Pressure Loading section Loads and Materials of the Theory Manual⁵²).

However, the Cholesky matrix factorizations are defined only if the correlation matrix is (numerically) non-singular. At this time, the code stops with an error if this occurs. A common cause of this error is using too many time steps NTimes with too small a time step. For this reason, the condition number of the temporal correlation matrix is always evaluated, and, if it is singular, the cutoff frequency is decreased. In this case the warning message

```
Singular temporal correlation matrix
Increasing Delta_T to ...
```

will be printed in the result file for processor 0. Another source of ill conditioning is the use of large correlation lengths correlation_length_z or correlation_length_r, or a fine mesh.

For this reason inverse condition number estimates are printed in the result files. An inverse condition number is the relative distance to a singular matrix, and is denoted Rcond, for reverse condition number. In double precision, an Rcond below 10^{-12} indicates that the factorization may fail. The precise statements in the results files are

```
TemporalCorrelationMatrixRcond = ...
Estimated SpatialCorrelationMatrixRcond = ...
Estimated CorrelationMatrixRcond = ...
```

6.3.16. Frequency Dependent Loads

Frequency dependent loads may be applied for frequency response analysis. The real part of these loads is applied exactly as above with the understanding that the functions referenced apply to frequency not time. Frequency dependent loads may include an imaginary component. This is done by prefixing the load types listed above by the letter “*i*”. Thus, the imaginary part of the load uses these load types.

For Complex Analysis	
Option	Parameters
iforce	<i>val1 val2 val3</i>
imoment	<i>val1 val2 val3</i>
igravity	<i>val1 val2 val3</i>
ipressure	<i>val1</i>
itraction	<i>val1 val2 val3</i>

A function should be associated with each such load. An example follows.

```
Loads // example for FRF analysis
  nodeset 1
    force=1 0 0      // the real part of the load
    function=11
  nodeset 2
    iforce=1 0 0     // the imaginary part of the load
    scale .707
    function=12
END
```

6.3.17. Modal Force Loading

Modal force loading can be used to directly load specific modes in the **modaltransient** solution case. An example input follows.

```
ECHO
  modalvars
END
Loads
  body
  ModalForce
```

```

    function GravFunction
END
FUNCTION GravFunction
    type table
    tablename 28
END
TABLE 28
    dimension 2
    size 100 9
    delta 0.000005 1
    origin 0.000005 0
    datafile=Qforce.txt
END

```

The format of Qforce.txt is identical to the output from a traditionally loaded modal transient solution with **echo modalvars**. **echo modalvars** does not print out a load at time zero, so either the load at time zero should be added to the datafile, or **origin** should be used to specify the first time in the file. Loads can be linearly interpolated between time values in the datafile. With redundant modes, some modes can change between subsequent **Sierra/SD** runs. Care should be taken to ensure the load are applied correctly.

Full details of table input are discussed in Section 2.8.19. Only two-dimensional tables are supported for modal force loading, with rows representing equally spaced time steps and columns representing each mode. Modal force loading is the only load type that supports two-dimensional table inputs.

6.3.18. Rotational frames

Often when analyzing rotating structures, it is convenient to perform the analysis in the rotating frame where the structure is not undergoing large displacement. Analysis in that frame introduces “fictional” or “pseudo” forces with centrifugal,² Coriolis and Euler contributions. These are termed “forces”, but the contributions are introduced from operating in a non-inertial coordinate frame as described in the Theory Manual section Loads and Materials subsection Analysis of Rotating Structures.

The associated keywords are found in Table 6-140.

The Galerkin framework used for finite elements, introduces matrices associated with these pseudo forces. In addition to the standard mass and stiffness matrices that arise in linear structural dynamics, force-based matrices are also common. These include follower stiffness matrices from applied pressures, and Coriolis/centrifugal matrices in rotating structures.

²There is often confusion about the description of the “centrifugal” or “centripetal” term. The *centripetal* force is a real force applied in the inertial coordinate frame which causes an object to travel in a circular path. The *centrifugal* force is the pseudo-force that appears from inertial terms in a rotating coordinate frame.

Option	Parameters
angular_velocity	<i>vel1 vel2 vel3</i>
angular_acceleration	<i>accel1 accel2 accel3</i>
coordinate	coordinate name

Table 6-140. – Rotating Frame Parameters.

Input 6.13 provides the corresponding **Sierra/SD** input for a rotational load applied to a body. The centrifugal stiffness and Coriolis coupling matrices are both derived from the rotational velocity of the structure, which uses the keyword **angular_velocity**. The vector angular velocity components are specified after the **angular_velocity** keyword.

An angular acceleration, $\dot{\Omega}$, may also occur, as when an aircraft carrying a weapon makes a rapid course correction. This angular acceleration results in a pseudo-force, called the Euler force, that is tangent to the angular acceleration vector. Application of angular acceleration is restricted to linear and nonlinear statics analysis in **Sierra/SD**.

For static loads analysis angular acceleration and angular velocity are applied independently. A similar static loads analysis of a rocket provides envelope survivability information during launch.

```
LOADS
  body
    angular_velocity = 0.0 2.0 0.0
    coordinate = 1
  body
    angular_acceleration = 0.0 0.0 3.0
    coordinate = 3
END
```

Input 6.13. Application of centrifugal and Euler forces. The loads above apply an angular acceleration of 3 radians/ s^2 in the Z -direction of coordinate frame 3, and an angular velocity of 2 radians/s in the Y -direction of coordinate frame 1.

Angular acceleration is applicable only in statics.

Left Hand Side contribution

Angular velocity introduces both left-hand side matrices and right-hand side force vectors. The algebraic expression for dynamics can be written as follows.

$$(K_m + K_g + K_{cen})u + (C + C_{cor})\dot{u} + M\ddot{u} = f_{extern} + f_{cen} \quad (6.9)$$

where,

K_m is a material matrix,
 K_g is the geometric stiffness matrix correction,
 K_{cen} is the centrifugal softening term,
 C is the damping/coupling matrix,
 C_{cor} is the Coriolis coupling matrix,
 M is the mass matrix,
 $f_{external}$ is the external force vector,
 f_{cen} is the centrifugal force, and
 u is the displacement.

Except for K_g , every matrix term is constant, depending only on the geometry and the elements. The Coriolis and centrifugal terms are also independent of displacement, u , though they depend on Ω .

For linear analysis (both linear statics and linear transient dynamics), the geometric stiffness terms is zero. However, since this term depends on stress, which is proportional to displacement, the geometric stiffening is typically proportional to the square of the angular velocity. As the geometric stiffening is typically of the same magnitude as centrifugal softening (also proportional to Ω), confusion can arise.

In multicase analyses, the matrices are typically generated only once; exceptions occur for nonlinear solutions and for the **tangent** method. It is recommended that linear solution cases include an update to the tangent stiffness matrix as part of a multicase solution. An example is shown in input 6.14.

```

Solution
case s1
  statics
  load=1
case up
  tangent
case s2
  statics
  load=1
End
  
```

Input 6.14. Example using Tangent Update

Limitations

There are a number of limitations for the rotational frames implementation.

1. Static analysis appropriately applies the centrifugal and Euler forces. The left-hand side matrix for geometric stiffness is only properly updated if the **tangent** step is applied.
2. In a single case solution, quadratic eigen solutions will include Coriolis and centrifugal terms if angular velocity is specified in the **loads** section. This is the case even though there is no true load for Q EVP.
3. Currently, Q EVP solutions can be computed for rotating structures only if there are no rigid body modes in the structure. An example is shown in input 6.15. In this case, a static preload with a rotational load is computed, followed by a tangent update, and then followed by a Q EVP analysis. This type of analysis would be useful for examining the effect of rotational loading on the modes of a structure. However, this will only work correctly if there are no rigid body modes in the structure.
4. The user is limited to one rotational frame per analysis. In other words, the whole body must be rotate together. One could not model a helicopter in a fixed frame and the associated rotor in another.
5. Rotational loads applied in the rotating frame are linear loads, and do not require a follower keyword.
6. For transient dynamics, the time varying function must be 1.0.
7. Angular acceleration is only applicable to statics analyses.
8. Superelements do not retain full accuracy. It is recommended that interface dofs for superelements retain either 3 or 6 degrees of freedom.
9. The Boundary section applies to all cases in a multcase solution.

```
Solution
case 'statics'
  statics
    load=1
case 'up'
  tangent
case 'qevp'
  qevp
    method=projection_eigen
    nmodes=100
    load=1
End
```

Input 6.15. Example of using qevp for Tangent Update

NOTE:

A time varying function with magnitude 1.0 for the full time span should be used for time varying solution cases. Additional work would be required to apply general loading patterns.

Finally, for optimal accuracy, the user can update the tangent matrix at the expense of greater computational expense.

6.3.19. Rigid Body Filter for Input

For some analyses, it is advantageous to remove rigid body components of a solution. The input forces may be filtered so that only self-equilibrated forces are applied. This process of removing the rigid body component from the solution is sometimes referred to as ‘Inertia Relief’ or ‘Inertial Relief’.

The rigid body filter is applied using input in the **parameters** section (2.3) and is illustrated in the example of input 6.16. The filtered force values can be output by requesting the **force** output option (section 7.1.34).

While the filter can ensure equilibrated loads, additional parameters may be required to help the linear solver address the singularity generated by floating structures. Typical input is provided here, with details in the appropriate sections. The **constrain_rbms** solver option must be used in conjunction with **FilterRbmLoad** *only* when the rigid body modes are in the null space of the system matrix. For example, **constrain_rbms** should be used for statics where the system matrix is the stiffness matrix, but **constrain_rbms** should not be used for transient where the system matrix includes inertial terms. Currently, only **GDSW** supports selectively constraining rigid body modes. The **FilterRbmLoad** parameter is supported for transient ³ and static solution cases. For other solution cases this parameter will have no effect on the solution. ⁴ The similar capability for modal solutions is presented in Section 3.36.

```
Solution
  statics
  solver=gds
  solver_options=gdsOptions
End

Parameters
```

³The FilterRbm option is only compatible with Newmark-Beta time integration. The generalized-alpha time integrator may be used, but the coefficient $\alpha_m = (2\rho - 1)/(1 + \rho)$ will be set to zero, where ρ is the integration parameter. The resulting integration scheme is still 2nd-order accurate. The rigid body motion equations will be integrated using the same scheme as the flexible body equations. This is reviewed in subsection Damping of Flexible Modes Only section Solution Procedures of the Theory Manual.

⁴Modal solutions, such as **modaltransient**, do not use FilterRbmLoad. However, see⁵³ for means of accomplishing the same process by direct use of the geometry rigid body modes.

```

FilterRbmLoad=allStructural
RbmTolerance=1e-10
End

solver_options gdsOptions
  constrain_rbms "X Y Z RotX RotY RotZ p"
End

```

Input 6.16. Rigid Body Filter Example Input

The names of the rigid body modes to constrain are the same as those setting boundary conditions 6.1. The ‘p’ keyword refers to the constant pressure in a structural acoustics problems, and the other six are the typical rigid body modes of a structure in Cartesian coordinates. ⁵

If the rigid body filter is activated, **Sierra/SD** calculates the corresponding rigid body modes, checks the residuals, and report a fatal error if the residual norms,

$$\frac{\|K\Phi_r\|_2}{\|K_d\|_\infty\|\Phi_r\|_2}$$

are larger than **RbmTolerance**. The default **RbmTolerance** is 10^{-10} . Each module that needs the rigid body modes will recalculate them.

6.3.20. RanLoads

The **RanLoads** section is used to provide input information for spectral input to a random vibration analysis. In a random analysis, the output response relates to the input, as

a_i is the output quantity at degree of freedom, i . For example, a_i may be the acceleration power spectrum, measured in $(in/s^2)^2/Hz$.
 follows. H_{ij} is the transfer function from input i to dof j .
 S_{jk} is the input power spectrum. Typically, this is in units of $(force)^2/Hz$.
 It is dimensioned to the number of independent inputs.

The relationship is

$$a_i(\omega) = \sum_{j,k} H_{ji}^T(\omega) S_{jk}(\omega) H_{km}(\omega). \quad (6.10)$$

The **RanLoads** section provides a specification for $S_{jk}(\omega)$. Note that this input will contain both a spatial and spectral component. In **Sierra/SD**, we require that each matrix element in the input power spectrum be expressible as a product of spectral and

⁵The keywords are “x”, “y”, “z”, “RotX”, “RotY” “RotZ” and “p”. The set of these keywords must be enclosed in quotation marks.

spatial components. Y_i is a spatial loading term associated with the i^{th} row and column of S , and F_{ij} is a spectral only matrix function.

$$S_{ij}(\omega, x) = Y_i(x)Y_j(x)F_{ij}(\omega) \quad (6.11)$$

It typically has units of $1/Hz$.

The **RanLoads** section contains the following required keywords.

Parameter	Argument	Description
matrix	<i>Int/String</i>	matrix-function identifier
load	<i>Integer</i>	row/column identifier

The **matrix** keyword identifies the appropriate **matrix-function** (see Section 2.8.17). The matrix-function determines the dimensionality of the input (using the **dimension** keyword). It also determines the spectral characteristics of the load.

The spatial characteristics (which correspond to Y_i in equation 6.11) are determined in **load** sections within the **RanLoads** definition. There must be exactly as many **load** sections as the dimensionality of input. For example, if the S_{FF} matrix is 3×3 , then there should be 3 separate load sections. Each load section within the **RanLoads** block must be followed by an integer indicating to which row/column it corresponds. The details of each **load** section are identical to the over all **loads** section (see 6.3) except that no time/frequency function is allowed. Note that only one load is required per row of the S_{FF} matrix, but each *entry* of the matrix may have a spectral definition (identified by a real and/or imaginary **function**).

The following example illustrates the definition of a single input specification. The loading is scaled so that a 1000 lb (or 454 kg) mass located on the input point (in nodeset 12 here) is scaled to produce a unit g^2/Hz loading.

```
RanLoads
  matrix=1
  load=1
    nodeset 12
      force=0 1 0
      scale 1.00e3 // needed to convert to g
      // loads input in lbs. The PSD is in g^2/Hz.
      // F = accel * mass
      //   = accel * (scale_factor)
      //   = accel * ((1000*.00259)*384.6)
END
```

Scaling the input force for a random vibration analysis can be confusing.⁶ This is

⁶Note that we are scaling the spatial forces, Y_i , which are combined as a product in equation 6.11. Thus, the scale factor is linear in the load. The resulting input power spectrum, S_{ij} , will contain the square of the scale factor.

especially true since enforced acceleration cannot be used to apply the force . The example above uses United States customary units. The **wtmass** parameter has been applied. In SI units, **wtmass**= 1, and the force would need to be multiplied by g to apply the input as acceleration in g 's.

The input acceleration may be examined by evaluating the output PSD at the input degree of freedom. This is done by putting the applied load set into the **frequency** section (7.4), and adding the **acceleration** keyword. The output is in the native units of analysis. For the example above, the output will be in $(in \cdot lbm/s^2)^2/Hz$, and must be divided by $(386.4)^2$ to convert to g^2/Hz .

6.4. Initial Conditions

Initial conditions are specified via the **initial-conditions** section. The initial conditions are used as the initial state for transient analysis. Both linear and nonlinear transient are supported. The variables supported for initial conditions are given in Table 6.4.

Keyword	Field Names	Description
displacement	dispX, dispY, dispZ, displacement_x, displacement_y, displacement_z	Translational displacement
velocity	velX, velY, velZ, velocity_x, velocity_y, velocity_z	Translational velocity
rotation	rotX, rotY, rotZ, rotational_displacement_x, rotational_displacement_y, rotational_displacement_z	Rotations for beams, shells, etc.
rotational_velocity	velRX, velRY, velRZ, rotational_velocity_x, rotational_velocity_y, rotational_velocity_z	Rotational velocity for beams, shells, etc.
acoustics	acoustics	Primary acoustic variable, usually pressure
acousticsdot	acousticsdot	Derivative of primary acoustic variable, usually particle velocity

Each term may be initialized in one of three ways. Initial conditions may be either

1. Read in from the **Exodus** file,
2. specified globally in the **initial-conditions** section,
3. or specified on a block-by-block basis in the input deck.

6.4.1. Reading Initial Conditions from the Mesh File

Initial conditions can be read from fields present on the mesh file. This method would usually be used when handing of initial conditions from the end state of a previous **Sierra/SD** or **Sierra/SM** run. The **time** command can be used to set which time step from the exodus file the initial conditions are read from. Valid options for time include reading from a specific time, start time of the transient simulation, or from the first or last step of the **Exodus** database. By default, time is set to the first step of the input **Exodus** database.

```
Initial-Conditions
time=<real>time_val|last|first|transient_start_time
end
```

An example of initial conditions input from **Exodus** file is given below. In this case, **Sierra/SD** will read both velocity and displacement initial conditions from the last time step of the **Exodus** file using the field names in Table 6.4.

```
INITIAL-CONDITIONS
  velocity=from_file
  rotational_velocity=from_file
  displacement=from_file
  rotation=from_file
  time = last
END
```

6.4.2. Setting Initial Conditions in the Input Deck

An example of setting initial conditions globally is given below.

```
INITIAL-CONDITIONS
  velocity=1 0 0
END
```

In this case, the entire model is given an initial velocity of 1 in the x direction, and 0 for the y and z directions.

An example of the third option (block-by-block specification) follows.

```
INITIAL-CONDITIONS
  velocity=by_block
END

BLOCK 1
  velocity = 1 0 0
```

```
END
```

```
BLOCK 2
```

```
velocity = 0 1 0
```

```
END
```

In this case, the velocity is read in from the input deck on a block-by-block basis. If two blocks share nodes and are given different initial conditions, then the results may be unpredictable, since the common nodes on the blocks would have conflicting initial conditions. Thus, it is recommended to verify that blocks are disjoint before specifying different initial conditions on a block-by-block basis.

Initial conditions are currently only implemented for transient analysis. They can also be used in multcase solutions, but they will only have an effect on the transient analysis that are in the multcase solution. For multiple transient analysis in a multcase, only the first transient analysis will use the initial conditions. The subsequent transient cases will get initial conditions from the previous case.

Usage Guidelines Care must be taken when setting initial conditions. If initial conditions are set from a **Sierra/SM** analysis, then that analysis must in general have small deformations. For example transferring an initial condition displacement from **Sierra/SM** with significant rigid body rotation would yield large stress in a linear **Sierra/SD** analysis due to the incompatible linear and nonlinear deformation state. Additionally, for elements like shells the nodal rotations must be compatible with the overall translational displacement.

6.5. *Use cases for initial acceleration*

All these cases can be applied as a spring mass system.

1. Apply force, displacement and velocity as an initial condition for transient dynamics. Compute,

$$A_n = (F_{applied} + C v_n + K d_n) / M$$

With a constant force, we should get a constant acceleration.

2. Apply a static load to the mass on the spring. Then apply (in a second case) the same load as a transient load. This should also result in a constant displacement.
3. TSR read. Read a static stress to the body. This is usually followed by a subsequent transient run.

Ring down is expected. Internal stresses are not treated the same as applied forces. Applied external forces are not carried forward. Internal stresses are carried forward

in the solution. When you start the transient run, you have an internal force only. Compute

$$A_n = (F_{internal})/M$$

4. **receive_sierra_data.** Read undeformed coordinates, X_o , and a nodal displacement, U_o . Compute *initial* coordinate, $X_1 = X_o + U_o$. Following this, no load is required to maintain the X_1 configuration. Initial stresses may cause a deformation from that initial system. We may get velocity, force, and acceleration from file.
5. Prescribed acceleration. $A_o = a(t)$. This is applied only on a surface or node set. We convert this to a prescribed force, and treat it identically to an applied force.
6. Restart. $A(0)$, $v(0)$ and $d(0)$ are prescribed at all locations. No solve needed.
7. There is a jump in force at some time other than $t=0$. Do nothing. We assume that the analyst wants it to ring.

6.5.0.1. Summary

- Use **vel**, **disp** from state change (whether **read_sierra**, user input, or multicase). Eigen may store displacement on the database. It is never used in a subsequent transient or statics analysis. It is always used in modal superposition cases. We get this by calling them PHI.
- Retain stresses across state change (multicase or read). Some cases retain stresses (statics, TSR, transient). Eigen may compute stresses, but does not store them on the database.
- Do NOT retain accelerations or forces across state change. We should use the forces calculated at that time (in the next state) for computing the balance.

7. OUTPUTS

7.1. Exodus

The **outputs** section identifies the data to write to the corresponding output files. Geometry-based finite element results are written to an output **Exodus** file. The name of this file is generated by taking the base name of the input **Exodus** geometry file, and inserting **-out** before the file extension. For example, if the input **Exodus** file specification is **example.exo**, output will be written to **example-out.exo**. When using a multicas solution (Section 3.2), the case identifier is used in place of **out**. More details are available in the **file** section (2.2).

Non-geometry-based finite element data, including system matrices and tables may be output in a format that is compatible with MATLAB. These text files have the extension **m**. The base file names describe the nature of the output. These files are generated in the current working directory.

Option	Description	Section
acceleration	acceleration at nodes	7.1.15
acousticlighthill	acoustic Lighthill source term	N/A
aforce	acoustic forces	N/A
block_energies	block kinetic and strain energies	7.1.31
constraint_info	nodal constraint information	7.1.48
disp	displacements at nodes	7.1.13
EForce	element forces for beams	7.1.38
ElemEigChecks	first eigenvalue, seventh eigenvalue, and the largest eigenvalue	7.1.11
ElemQualChecks	on off (<i>any boolean</i>), default is on	7.1.12
energy	element strain energy and strain energy density	7.1.28
eorient	element orientation vectors	7.1.43
faa	force vector in the a-set	7.1.8
fatigue	Fatigue Damage Estimates	7.1.23
force	the applied force	7.1.34
GEnergies	global sum of energies	7.1.29
globals	ensure global output	7.1.30
Kaa	stiffness matrix in the a-set	7.1.7
line_weld	line weld elements	7.1.39
locations	nodal coordinates and DOF to node map	N/A
Maa	mass matrix in the a-set	7.1.4
material	material parameter element output	7.1.5
material_direction_1	Local element coordinate system R vector	7.1.6
material_direction_2	Local element coordinate system S vector	7.1.6
material_direction_3	Local element coordinate system T vector	7.1.6

MLumped	nodal lumped mass in eigen solution	7.1.9
MPhi	Mass $\times$ Displacement in eigen solution	7.1.10
mesh_error	mesh discretization errors	7.1.32
MFile	MATLAB MFiles	7.1.33
nodal_charge	applied charge at nodes	N/A
pressure	pressure load vector	7.1.44
rainflow	Rainflow cycle counter output	7.1.22
reaction_force	the Dirichlet boundary reaction force	7.1.36
relative_disp	1D element nodal separation	7.1.40
rhs	RHS of system of equations to be solved	7.1.37
statistics	Mean and Std deviation of some variables	7.1.49
strain	element strain	7.1.16
stress	element stress	7.1.18
stress = gp	element stresses at Gauss points	7.1.17
principal_stresses	element principal stress vectors and magnitudes	7.1.19
signed_vonmises	element signed von Mises	7.1.21
temperature	element centroid temperature	7.1.53
velocity	velocity at nodes	7.1.14
voltage	voltage at nodes	7.1.54
vonMises	element von Mises stress	7.1.20
vrms	RMS quantities (random vibration only)	7.1.25
WarningLevel	Control of warning messages	N/A

Table 7-141. – OUTPUT Section Options.

In the following example, the mass and stiffness matrices will be output in MATLAB format, but the displacement variables, stresses and strains will not be output. All the various options of the **outputs** section are shown in Table 7-141. The next sections describe each of the options and their results assuming an input file named **example.inp** and a geometry file named **exampleg.exo**.

```

Outputs
  Maa
  Kaa
  Faa
// displacement
// stress
// strain
// energy
end

```




Figure 7-68. – Convergence of maximum stress at element centroids and surfaces.

7.1.1. Surface Projection of Element Variables

Element output such as stress, strain or temperature is evaluated at the element centroid. Hexahedra, tetrahedra and prisms support the projection of some element quantities to element faces. We do this by sampling the gradients of the element shape functions at the local coordinates associated with that face's centroid.

Notes:

1. Surface Projection is triggered by the `output_sideset_data` parameters flag.
2. Surface Projection results are stored as sideset variables on the **Exodus** mesh
3. Only some outputs have been enabled for surface projection, most notably the strain tensor, stress tensor, and VonMises stress.
4. Due to the nature of bubble shape functions when sampled away from the element centroid, the default Hex8 element (Hex8b) will report skewed surface projection results, and is not expected to give the same benefit as other volumetric elements, given the same displacement field.
5. Even when using the bubble hex element, surface projection is expected to give a more accurate representation of maximum stress, given that the maximum stress occurs on a sideset. See figure [7-68](#)

7.1.2. Database Name

Option **database name** in the **outputs** section allows the user to specify an output file name instead of the default behavior as described above. As before, **-out** (or the case name for multi-case solutions) will be inserted before the file extension.

For instance, where the geometry filename **exampleg.exo** would previously result in the output file **exampleg-out.exo**, the user could instead specify **database name = new_example.exo**, which would result in output written to **new_example-out.exo**.

7.1.3. Properties



Exodus Output Properties is currently BETA release.
Enable with the “**-beta**” command-line option.

A number of options are available to control the contents of **Exodus** output files. These options can be used to shrink the output size or change the format of the output file. All of these properties are preceded by the **property** keyword in any input deck. Available properties are given in Table 7-142.

Table 7-142. – Exodus Property Output Options.

Property Type	Description
integer_size_db	Number of bytes for ints in Exodus output. Valid sizes are 4 or 8. By default, the size is the same as used in the input mesh file. 8 byte integers are required if node/element IDs greater than about two billion are used.
real_size_db	Number of bytes for reals in Exodus output. Valid sizes are 4 or 8. Default is 8, double precision. Use of 4 byte single precision output will roughly halve the output size but reduce output accuracy. Of particular note, Do not use four byte reals for restart.
compression_level	Compression level can be given values from zero to nine. Zero is no compression and nine is maximum compression. The default is generally zero compression. The compression level provides only a suggestion to the Exodus library, the actual amount of compression will vary. Potentially compression can reduce the size of the file by a factor of three or more. Compressed files may be incompatible with some post-processors.

The below example parameters would provide minimal output file size for a run.

OUTPUT

```
property integer_size_db = 4
property real_size_db = 4
property compression_level = 9
END
```

7.1.4. Maa

Option **Maa** selects the analysis-set mass matrix (if it exists) for output to the file named **Maa.m**. These are the file names for the serial version of **Sierra/SD**. In the parallel version, an underscore and the processor number will precede the **.m**, and a separate file will be written for each processor.

7.1.5. Material

The **material** keyword will output material properties for each element. Currently, this capability is only enabled for elasticity calculations, and is also not enabled for Lamé materials or **isotropic_viscoelastic_complex** materials.

7.1.6. material_direction

Local coordinate systems may be defined to orient directional materials (sections 2.7 and 4.7.1 and section 4.7.2.3). The **material_direction_1** , **material_direction_2** , and **material_direction_3** outputs provide the element local coordinate system $(\hat{r}, \hat{s}, \hat{t})$ vectors at each element. Visualizing these vectors can help inform if material coordinate systems have been setup as intended. Additionally, these coordinate vectors are the rows of a 3×3 rotation matrix for transforming quantities between the global (X, Y, Z) and local $(\hat{r}, \hat{s}, \hat{t})$ coordinate systems within the element.

7.1.7. Kaa

Option **Kaa** will output the analysis-set stiffness matrix to a file named **Kaa.m**. These are the file names for the serial version of **Sierra/SD**. In the parallel version, an underscore and the processor number will precede the **.m**, and a separate file will be written for each processor.

7.1.8. Faa

Option **faa** will output the analysis-set force vector (if it exists) to a file named **Faa.m**. These are the file names for the serial version of **Sierra/SD**. In the parallel version, an underscore and the processor number will precede the **.m**, and a separate file will be written for each processor.

7.1.9. MLumped

Option **MLumped** will output the lumped mass matrix to the **Exodus** mesh as a nodal variable. **MLumped** is only implemented for the **Eigen** solution case. The lumped mass output is based on the analysis-set reduced mass matrix. Thus, mass on fixed degrees of freedom will be zero.

MLumped Variable Names
M_X (x translation)
M_Y (y translation)
M_Z (z translation)
M_RX (x rotation)
M_RY (y rotation)
M_RZ (z rotation)
M_A (acoustic)
M_V (voltage)
M_T (temperature)

7.1.10. MPhi

Option **MPhi** triggers computation and output of the mass matrix product $M\Phi$ to the **Exodus** file for the mode shapes, Φ , computed in the previous **eigen** solution. **MPhi** is only implemented for the **Eigen** solution case. A consistent mass matrix is used. Mass matrix output is explained in Section 7.1.4. Like the mode shapes, the mass matrix is defined only on the analysis-set dofs. An **ASetMap** is also provided. The $M\Phi$ vanishes on fixed dofs. The names of the different dofs at a node are specified in Table 7.1.10.

MPhi Variable Names
MPhi_X (x translation)
MPhi_Y (y translation)
MPhi_Z (z translation)
MPhi_RX (x rotation)
MPhi_RY (y rotation)
MPhi_RZ (z rotation)
MPhi_A (acoustic)
MPhi_V (voltage)
MPhi_T (temperature)

7.1.11. ElemEigChecks

Option **ElemEigChecks** will turn on the element output of the lowest eigenvalue, the 7th eigenvalue (commonly the first flexible eigenvalue), and the largest eigenvalue for the element. The output will be stored in the **Exodus** output file. The element variable names for the 1st eigenvalue, the 7th eigenvalue, and the maximum eigenvalue are *ElemEig_1st*, *ElemEig_7th*, and *ElemEig_max*, respectively. Note: All 3-d and 2-d elements have this capability. The **Beam2**, **TiBeam**, **Spring**, **Truss**, **Spring3**, and **RSpring** elements are also supported. All remaining elements will output values of zero. Finally, if $ElemEig\_1st < -1e-12 ElemEig\_max$, a negative eigenvalue warning will be printed.

7.1.12. ElemQualchecks

Option **Elemqualchecks** takes a boolean, e.g. the default **on** or **off**. Unless this option is **off**, all the elements in the input file are checked for quality using the SIERRA element quality metrics.⁶⁵ The following table shows the elements currently checked and their acceptable ranges. The element quality reporting may also be modified by the *condition_limit* parameter specified in the *Parameters* section (2.3).

Element Type	Full Range	Acceptable Range
Hex8	$1 - \infty$	1 – 8
Tet4	$1 - \infty$	1 – 3
Tria3	$1 - \infty$	1 – 1.3
TriaShell	$1 - \infty$	1 – 1.3
Quad4	$1 - \infty$	1 – 4
Wedge6	$1 - \infty$	1 – 5

If the option **on** is selected and the element's condition numbers falls outside the acceptable range, a warning message is printed. The value output with the warning is normalized by the maximum number of the acceptable range for that element. A summary is also printed, reporting the maximum condition number of all elements in the mesh.

In addition to these checks, solid elements are checked for negative volumes. This can occur if the node ordering for the element establishes a “height” vector using the right-hand rule that is in the opposite direction of the actual element height. In other words, the nodes should normally be ordered in a counter clockwise direction on the bottom surface of the element. Some codes such as NASTRAN, are insensitive to this ordering. If element checks are run, then **Sierra/SD** will correct (and report) any solid elements found to have negative volumes. Without these corrections, the code will continue, but results that depend on these elements are suspect.

It is strongly recommended that any **Exodus** file with negative volumes be corrected.

7.1.13. Displacement

Option **disp** will output the displacements calculated at the nodes to the output **Exodus** file. The output file has the following nodal variables.

Variable	Description
dispX	X component of displacement
dispY	Y component of displacement
dispZ	Z component of displacement
rotX	Rotation about X
rotY	Rotation about Y
rotZ	Rotation about Z

In addition, if the analysis involves complex variables (**ceigen** Section 3.24.2, frequency response analysis such as **ModalFrf** or **sa_eigen**), then the imaginary vectors are also included. The imaginary component of the vector has “*imag*” prefixed to the name. For example, the imaginary component in the *X* direction is “**imagDispX**”.

7.1.14. Velocity

Option **velocity** will output the velocities at the nodes to the output **Exodus** file.

7.1.15. Acceleration

Option **acceleration** will output the accelerations at the nodes to the output **Exodus** file.

7.1.16. Strain

Option **strain** will output the strains for all the elements to the output **Exodus** file. The total strain ϵ , elastic strain ϵ^e , and thermal strain ϵ^θ are defined by:

$$\epsilon = \frac{1}{2} (\nabla u + \nabla u^T) \quad (7.1)$$

$$\epsilon^e = \epsilon - \epsilon^\theta \quad (7.2)$$

$$\epsilon_{ii}^\theta = \alpha \Delta T \quad (7.3)$$

They can be requested separately by using the options

- **strain** : ϵ
- **elastic_strain** : ϵ^e
- **thermal_strain** : ϵ^θ .

Strains will be output for shell elements. The output variable names start **SStrain**. Shell elements output engineering strain. E.g. shells output ϵ_{xx} , ϵ_{yy} , τ_{xy} where $\tau_{xy} = 2\epsilon_{xy}$. This is distinct from solid elements which output strain as ϵ_{xx} , ϵ_{yy} , ϵ_{zz} , ϵ_{xy} , ϵ_{yz} , ϵ_{xz} . This difference in strain definition means that for an equivalent deformation a shell element will report twice the numerical shear strain that a solid element will report even though the two shear states are identical.

SStrainX1, SStrainY1, SStrainXY1 - *top layer of the shell*
 SStrainX2, SStrainY2, SStrainXY2 - *mid-plane of the shell*
 SStrainX3, SStrainY3, SStrainXY3 - *bottom layer of the shell*

For elastic strain,

SElasticStrainX1, SElasticStrainY1, SElasticStrainXY1 - *top layer of the shell*
 SElasticStrainX2, SElasticStrainY2, SElasticStrainXY2 - *mid-plane of the shell*
 SElasticStrainX3, SElasticStrainY3, SElasticStrainXY3 - *bottom layer of the shell*

and for thermal strains we have

SThermalStrainX1, SThermalStrainY1, SThermalStrainXY1 - *top layer of the shell*
 SThermalStrainX2, SThermalStrainY2, SThermalStrainXY2 - *mid-plane of the shell*
 SThermalStrainX3, SThermalStrainY3, SThermalStrainXY3 - *bottom layer of the shell*

Strains are evaluated in the local element coordinate system. The local element coordinate system and also the ordering of the layers both depend on the ordering of the element's nodes.

The following strains will be output for volume elements:

VStrainX, VStrainY, VStrainZ,
 VStrainYZ, VStrainXZ, VStrainXY

Likewise, for elastic and thermal strains, we have

VElasticStrainX, VElasticStrainY, VElasticStrainZ,
VElasticStrainYZ, VElasticStrainXZ, VElasticStrainXY
VThermalStrainX, VThermalStrainY, VThermalStrainZ,
VThermalStrainYZ, VThermalStrainXZ, VThermalStrainXY

Note: These strains are in the global coordinate system, not the local coordinate system.

For more information on stress/strain recovery, see Section 7.6.

7.1.17. Strain = GP

An output specification of `strain = GP` reports strain at the Gauss points of volumetric elements. For more information, see section 7.1.24.

7.1.18. Stress

Option `stress` will output the stresses for all supported elements to the output **Exodus** file. Only shell and volume elements are supported, i.e. there is no stress output for beams.

This is the Cauchy stress $\sigma^C = C : \epsilon^e$; see equation 7.2.

7.1.18.1. Shell Stresses

The following stresses will be output for shell elements.

SStressX1, SStressY1, SStressXY1, SvonMises1 - *top layer of the shell*
SStressX2, SStressY2, SStressXY2, SvonMises2 - *mid-plane of the shell*
SStressX3, SStressY3, SStressXY3, SvonMises3 - *bottom layer of the shell*

*Note: the top layer of the shell is determined by the ordering of the nodes of the shell, and can be output by using the **eorient** output options (see Section 7.1.43). Also, the stresses are in the local element coordinate system defined by the ordering of the nodes.*

7.1.18.2. Volume Stresses

For volume elements, the stress is always output in the global coordinate system, not the local coordinate system. The following stresses will be output for volume elements:

Variable	Value
VStressX	σ_{xx}
VStressY	σ_{yy}
VStressZ	σ_{zz}
VStressYZ	σ_{yz}
VStressXZ	σ_{xz}
VStressXY	σ_{xy}
VonMises	von Mises stress

For more information on stress/strain recovery, see Section 7.6.

7.1.19. Principal Stresses

Option `principal_stresses` will output the three principal stress vectors and magnitudes in order of descending signed value. `Principal_stresses` are only supported for volume elements.

The magnitudes are output as:

`Max_Principal_Stress`, `Intermediate_Principal_Stress`, `Min_Principal_Stress`

The eigenvectors are output as:

`Max_Principal_Stress_x`, `Intermediate_Principal_Stress_x`, `Min_Principal_Stress_x`,
`Max_Principal_Stress_y`, `Intermediate_Principal_Stress_y`, `Min_Principal_Stress_y`,
`Max_Principal_Stress_z`, `Intermediate_Principal_Stress_z`, `Min_Principal_Stress_z`

7.1.20. Von Mises stress

Option `vonMises` will output the von Mises stress for all the elements to the output **Exodus** file. For volume elements, the output will be the von Mises stress of the element. Surface elements define stresses on the top, center and bottom layers. The output will be the maximum of these 3 values.

Note that the von Mises stress is computed and output as a portion of the output if full stress recovery is requested. This option provides a mechanism for reducing output. Thus, if full stress output is requested, then the `vonMises` will provide no additional output. In other words, specifying both `vonMises` and `stress` in the outputs section is redundant, but does not result in an error.

7.1.21. Signed Von Mises Stress

Option **Signed_VonMises** will output the magnitude of Von Mises stress, given the sign of the principal stress with the largest magnitude. **Signed_VonMises** is only supported for volume elements.

The output variable is:

Signed_VonMises

7.1.22. Rainflow Cycle Counting

Option **rainflow** triggers the rainflow cycle counter to track stress cycles encountered by each element over time. **Rainflow** relies on **Signed_VonMises** to convert the stress tensor to a scalar signal, and is intended as a preprocessing step to time domain fatigue calculations. **Rainflow** is only supported for volume elements, and only in transient analyses.

The output variables are:

NumCycles, LastCycleAmplitude, LastCyclePeak

These outputs contain insufficient information to be useful on their own except in simple verification exercises. **Rainflow** is intended to be a silent dependency of **fatigue**, rather than a standalone output option.

7.1.23. Fatigue Damage

Option **fatigue** will output a damage estimate for each element using the stress history of that element. This process is supported for transient solutions and for modal random vibration, but in very different ways.

In a transient analysis, **fatigue** damage is calculated using stress cycles identified by the **rainflow** algorithm, and applying those cycles to the Walker damage function:

$$\log_{10}(N) = A_1 + A_2 * \log_{10}(S_{max} * (1 - R)^{A_3} - A_4)$$

Where S_{max} is the peak stress of the cycle, S_{min} is the minimum Stress of the cycle, and $R = S_{min}/S_{max}$. The number of cycles to failure N is then related to damage D by:

$$D = 1/N$$

A_1 , A_2 , A_3 , and A_4 are material constants. Note that a cycle is ignored if $S_{max} \leq 0$ or $S_{max} * (1 - R)^{A_3} \leq A_4$, because purely compressive cycles are assumed to cause no damage, and because cycles below the endurance limit causes no damage.

The output variable is:

Damage

In modal random vibration, **fatigue** is its own solution case. See section [3.14](#) for more details. The output variables are:

NarrowBandDamageRate, WirschingDamageRate, ZeroCrossingRate,
PeakFrequency, Damage, Vrms

7.1.24. Stress = GP

An output specification of **Stress = GP** reports stress at the Gauss points of volumetric elements. It is currently only available for Hex20, Tet10, and Wedge15 elements. Note that for a Hex20 there are 27 Gauss points with 6 stresses, for a total of 162 outputs per element.

The Gauss point ordering follows the description in the paper by Thompson.^{[71](#)} For the convenience of the reader, that order is reproduced here in Table [7-143](#).

number	label suffix	X	Y	Z
1	111	0	0	0
2	112	0	0	A
3	110	0	0	-A
4	121	0	A	0
5	122	0	A	A
6	120	0	A	-A
7	101	0	-A	0
8	102	0	-A	A
9	100	0	-A	-A
10	211	A	0	0
11	212	A	0	A
12	210	A	0	-A
13	221	A	A	0
14	222	A	A	A
15	220	A	A	-A
16	201	A	-A	0
17	202	A	-A	A
18	200	A	-A	-A
19	011	-A	0	0
20	012	-A	0	A
21	010	-A	0	-A
22	021	-A	A	0
23	022	-A	A	A
24	020	-A	A	-A
25	001	-A	-A	0
26	002	-A	-A	A
27	000	-A	-A	-A

Table 7-143. – Hex20 Gauss Point Locations. The constant $A=0.77459666924148$. The unit element is $2 \times 2 \times 2$, with a volume of 8 cubic units.

7.1.25. **Vrms**

Option **vrms** will output computed root mean squared (RMS) quantities from a random vibration analysis. These quantities are written to a separate output file. Quantities output include the RMS displacement, acceleration and von Mises stress. With the SVD option, the D matrix terms⁵⁹ which contribute to the von Mises stress are also output (see Section 3.20).

7.1.26. **Rotational_displacement**

Option **rotational_displacement** will output computed root mean squared (RMS) quantities for rotational displacement in a random vibration analysis to the output file. In the **frequency** section, it will output the corresponding Power Spectral Densities to the frequency file. (See section 3.20).

7.1.27. **Rotational_acceleration**

Option **rotational_acceleration** will output computed root mean squared (RMS) quantities for rotational acceleration in a random vibration analysis to the output file. In the **frequency** section, it will output the corresponding Power Spectral Densities to the frequency file. (See section 3.20).

7.1.28. **Energy**

Option **energy** will place strain energies and strain energy density in the output **Exodus** file. Note that the current implementation of strain energies requires updating the element stiffness matrix, which can be expensive.

7.1.29. **GEnergies**

Option **GEnergies** in either the **echo** or **outputs** will trigger computation of global energy sums for the results or output **Exodus** file, respectively.

strain energy The strain energy is computed from $u^T K u / 2$ where u is the displacement and K is the current estimate of the tangent stiffness matrix.

kinetic energy Computed as $v^T M v / 2$. Here v is the velocity and M is the mass matrix.

work As a particle moves along $x(t)$ in the force field F , it does **work**

$$\begin{aligned} W(t) &= \int_{x(0)}^{x(t)} F(x) dx \\ &= \int_0^t F(\tau) \dot{x}(\tau) d\tau \end{aligned}$$

The simplest possible approximation is used,

$$W_n \sim \sum_{i=0}^n F_i \dot{x}_i \Delta t.$$

Integral approximation errors may introduce inconsistencies with the other energies. For the **outputs** case, the total energy is written out at each time step.

Known Issues:

- The strain energy may not be complete for nonlinear solutions. Linear viscoelastic materials have contributions that will not be included in this sum.
- The strain energy and work are currently not computed correctly in models with non-zero displacement, velocity, or acceleration boundary conditions.
- When used in statics, GEnergies only works with **echo**, not results output.

7.1.30. **Globals**

Option **Globals** in the **outputs** section will ensure that any global data is output, even if no other outputs are requested. Without this option (and with an empty output block), no output file would be written. If no global data is defined for a given solution type, or if any other outputs are requested, this option will have no effect. This option is also available for history, frequency, and statistics output, and behaves similarly.

7.1.31. **Block_Energies**

Option **block_energies** in the **echo** or **outputs** will trigger block-wise energy sums for the results or output **Exodus** file, respectively. The energy computations are done as described in 7.1.29, where the displacement and velocity vectors have been restricted to the element block. Kinetic and strain energies are computed. Global variables in the **Exodus** file are “KineticEnergy_” and “StrainEnergy_” with the block name appended.

7.1.32. Mesh_Error

The **mesh_error** keyword causes mesh discretization error metrics to be computed. These are computed as output quantities, but the overhead associated with the metrics is not negligible. Mesh discretization quantities depend upon the solution type, and are not available for all solutions. Output is typically available as element quantities (usually in the *mesherr* field). For some mesh discretization errors, a global quantity is also output.

Output	Description
ErrExplicitLambda	Relative error in λ .
ErrExplicitFreq	Frequency error estimate (Hz)

We note that for eigenvalue analysis, *relative* errors are reported for the eigenvalue when using the **mesh_error** keyword. Thus, for a given eigenvalue λ , the reported error is

$$\text{ErrExplicitLambda} = \frac{\lambda_h - \lambda}{\lambda} \quad (7.4)$$

This is more convenient since the analyst does not have to divide by the eigenvalues to see the percent error. The global variable **ErrExplicitFreq** provides an absolute estimate (useful in plots for example).

7.1.33. MFile

Option **MFile** will cause **Sierra/SD** to output various MFiles like *Ksrr.m*, *Mssr.m*, etc. Using MFile with a model with a large numbers of elements is challenging due to the large size of the corresponding MATLAB files. A partial index of the files written using this option is provided in Table [7-144](#).

Table 7-144. – Data Files Written Using the MFile Option.

Filename	Description
Stiff.m	Unreduced stiffness matrix including all active dofs
Kssr.m	Reduced stiffness matrix
Mass.m	Unreduced mass matrix
Mssr.m	Reduced mass matrix
LumpedMass.m	unreduced lumped mass matrix
xxx_gid.m	global IDs of the nodes
ASetmap_a.m	Map to convert from G-set to A-set The right-hand side is the equation number. The lhs index is $8*(\text{node index}) + \text{coordinate}$
Dampr.m	unreduced damping matrix (real components)
Dampi.m	unreduced damping matrix (imaginary components)
xxx_accelNN.m	G-set acceleration output of step NN
xxx_accel_aNN.m	A-set acceleration output of step NN
xxx_afNN.m	G-set applied force output of step NN
xxx_af_aNN.m	A-set applied force output of step NN
xxx_dispNN.m	G-set displacement output of step NN
xxx_disp_aNN.m	A-set displacement output of step NN
xxx_presNN.m	G-set nodal applied pressure of step NN
xxx_pres_aNN.m	A-set nodal applied pressure of step NN
xxx_velocNN.m	G-set velocity output of step NN
xxx_veloc_aNN.m	A-set velocity output of step NN
modal_amp.m	modaltransient output of mode amplitude vs time

- Above the **xxx** refers to the input file name root.
- G-set output has dimension 8 (number of nodes).
- The file names above are for the serial version of **Sier-ra/SD**. In the parallel version, an underscore and the processor number will precede the “.m”. For example, the reduced stiffness matrix becomes **Kssr_0.m**. Matrices are never globally assembled, and so cannot be output.
- Some solution methods will not write all files. For example, there are no mass matrices output in the solution of statics. Generally, matrices are output in sparse symmetric row format.
- The 1 to N node ordering of the input **Exodus** file defines the **ASetMap**. Output file ordering may be different if there is a node order map.

7.1.34. Force

Option **force** will output the applied force vector to the output **Exodus** file, and the net force applied to active degrees of freedom, and the net moment about the origin of forces applied to active degrees of freedom. The net force is calculated from the right-hand side given to the solver, so implementations that modify the right-hand side (cavitation) may display non-physical net forces and moments. Net forces are available for static and direct transient solution cases.

If rigid body filtering is requested via the **FilterRbmLoad** option (section 6.3.19), this option will also output the filtered force to the output **Exodus** file, and the net force and moment about the origin. The naming convention for both nodal and net (global) values are **force_inertia_relief** for forces, and **moment_inertia_relief** for moments.

7.1.35. Constraint force

Option **constraint_force** will output the forces required to tie blocks together due to multipoint constraints, RBE3s, Rbars, or rigid elements. Additionally, it will output the net constraint force on the active degrees of freedom, and the net moment about the origin of constraint forces on the active degrees of freedom. This output is currently only active for the static and transient solution cases and for the GDSW solver.

7.1.36. Reaction Force

Option **reaction_force** will output the Dirichlet boundary reaction force vector to the output **Exodus** file. Additionally, it will output the net reaction force on the active degrees of freedom, and the net moment about the origin of reaction forces on the active degrees of freedom. This output is currently only active for the static and transient solution cases and for the GDSW solver.

7.1.37. Right-hand side

This output is useful primarily for verification and debugging purposes. Option **RHS** selects the right-hand side vector for the analysis type. For statics and dynamics, the **RHS** vector is the applied forces, pressures, inertial forces, or any pseudo forces introduced in preload (say by TSR).

7.1.38. EForce

Option **eforce** will output the element forces for line elements (such as beams and springs) to the output **Exodus** file. Each two node, 1-dimensional element will have 6 force entries for each node, for a total of 12 element forces per element, and an additional 3 variables describing the difference in displacement across the element.

The element force is not a stress or a strain, and should not be used as such. If you want beam stresses, you may want to mesh that portion of the structure either as a shell or a solid. Only limited stress output is available for beams. EForce is used primarily to help understand the behavior of nonlinear line elements such as the Joint2g element (see Section 5.21). The output is the direct output of our internal force routine (which is a nonlinear routine). It can be confusing to output these nonlinear forces in a linear analysis. ⁷

Eforce variable names are different for each solution case. When requested in typical analyses, such as **statics**, **transient**, or **eigen**, **eforce** produces the variables in table 7-145. For **FRF** solutions, “Imag” is prepended to the existing names 7-146. For **modalranvib**, only translational forces are output, and are different between **outputs** and **frequency** files. In an **outputs** result file, you will see the RMS of the force in the translation directions 7-147. In **frequency**, the Spectral Densities of the element force are represented by a Hermitian tensor at each frequency 7-148.

Variable Names
eforce1_x
eforce1_y
eforce1_z
emoment1_x
emoment1_y
emoment1_z
eforce2_x
eforce2_y
eforce2_z
emoment2_x
emoment2_y
emoment2_z
e_dx
e_dy
e_dz

Table 7-145. – Typical Output.

⁷Confusion arises because of the transformation to the element coordinate frame. For finite length elements, we perform a transformation of the element coordinate frame based on the displacements. After the coordinate frame is transformed, we express the element force in the new coordinate frame. This is done for both linear and nonlinear analyses. The resulting element force is no longer linear in displacement. Zero length elements do not have a rotated coordinate frame by default. Forces for zero length elements are linear in the displacement.

Real Names	Imaginary Names
eforce1_x	Imageforce1_x
eforce1_y	Imageforce1_y
eforce1_z	Imageforce1_z
emoment1_x	Imagemoment1_x
emoment1_y	Imagemoment1_y
emoment1_z	Imagemoment1_z
eforce2_x	Imageforce2_x
eforce2_y	Imageforce2_y
eforce2_z	Imageforce2_z
emoment2_x	Imagemoment2_x
emoment2_y	Imagemoment2_y
emoment2_z	Imagemoment2_z
e_dx	Image_dx
e_dy	Image_dy
e_dz	Image_dz

Table 7-146. – FRF Output.

Real Names	Imaginary Names
eforceRMS_x	ImageforceRMS_x
eforceRMS_y	ImageforceRMS_y
eforceRMS_z	ImageforceRMS_z

Table 7-147. – ModalRanVib Exodus Output.

Real Names	Imaginary Names
eforce_xx	Imageforce_xx
eforce_yy	Imageforce_yy
eforce_zz	Imageforce_zz
eforce_xy	Imageforce_xy
eforce_yz	Imageforce_yz
eforce_xz	Imageforce_xz

Table 7-148. – ModalRanVib Frequency Output.

7.1.39. Line_Weld

Option **Line_Weld** outputs line-weld-specific outputs as described in Section 5.22). The outputs include whether a given beam is in an active weld, and the force per unit length produced by the weld

NOTE: The force returned is in the element (not global) coordinate frame.

7.1.40. Relative_Displacement

Option **relative_disp** in the **outputs** , **history** , and/or **frequency** section(s) will output the relative displacement between the two nodes of 1-D elements in the model. That is, it will output the difference in displacements across the 1-D element. Currently, **relative_disp** is only supported for the Joint2G element type, and is not supported for **FRF** solutions. Relative rotations are not available. For solutions other than **modalranvib** , including **FRF** , this data is also accessible through **eforce** .

relative_disp is output in the element-local coordinate system if a coordinate system has been defined for the Joint2G block, and in the global coordinate system otherwise. This is consistent with **eforce** outputs, but differs from **displacement** outputs. Using the **coordinate** keyword in the **outputs** , **history** , or **frequency** will not effect this output.

relative_disp is particularly useful for the **modalranvib** solution case, where **Sierra/SD** does not typically calculate or output enough information to post-process the difference in displacement between two points. The names of **relative_disp** outputs also change in **modalranvib** results. See tables 7-149, 7-150, and 7-151.

Variable Names
RelDispX
RelDispY
RelDispZ

Table 7-149. – Typical Output.

relative_disp outputs are available in PSD form if requested in the **frequency** section during **modalranvib** (Table 7-150). Note that the diagonal **relative_disp** terms are output for both real and imaginary components. Hermitian symmetry is assumed in the calculation of **relative_disp** outputs, and incorrect results will be reported if this assumption is broken. Neither the input load, nor the **relative_disp** output should have non-zero imaginary terms along the diagonal. That is, **ImagRelDispGxx** , **ImagRelDispGyy** , and **ImagRelDispGzz** should all be zero. The warning message “Correlation Matrix is negative” indicates you may have this problem.

Real Names	Imaginary Names
RelDispGxx	ImagRelDispGxx
RelDispGyy	ImagRelDispGyy
RelDispGzz	ImagRelDispGzz
RelDispGxy	ImagRelDispGxy
RelDispGyz	ImagRelDispGyz
RelDispGxz	ImagRelDispGxz

Table 7-150. – ModalRanVib Frequency Output.

Real Names	Imaginary Names
RelDispRMSX	ImagRelDispRMSX
RelDispRMSY	ImagRelDispRMSY
RelDispRMSZ	ImagRelDispRMSZ

Table 7-151. – ModalRanVib **Exodus** Output.

relative__disp outputs are available in RMS form if requested in the **outputs** section during **modalranvib** (Table 7-151). Only diagonal terms are reported. The imaginary diagonal terms are set to 0.

7.1.41. Residuals

For most solution types, a linear solver is used to compute systems of the form $Ax = b$. For direct serial solvers, these systems are typically solved to numerical precision. However, with iterative solvers the solution is only approximate. Sometimes it is advantageous to evaluate the performance of the solver. For example, regions with large residuals may be candidate areas for mesh refinement, or may point to other mesh problems.

Eigen. For eigen analysis, the residual is $(K - \lambda_i M)\phi$. The vector is *not* normalized by the norm of ϕ , or any other quantity. A nodal residual work is also output. This is the product $\phi^T(K - \lambda_i M)\phi$ summed to the nodes, i.e. on a given node we sum the contributing degrees of freedom. Again, the value is *not* normalized. With mass normalized eigenvectors (which do not have units of length), the units of the residual work are not energy, and the term may well be negative. The residual is output for each mode.

Transient Dynamics. For transient analysis the residual reported is $Au - b$, where A is the dynamics stiffness matrix defined in subsection Linear transient analysis section Solution Procedures of the Theory Manual.

With a displacement based Newmark-Beta integrator the dynamic stiffness is $K + \frac{2}{\Delta t}C + \frac{4}{\Delta t}M$. The residual is output at each time step.

In addition to the residual vector, the norm of the residual is output as a global variable.

7.1.42. TIndex

It is occasionally useful to examine the residual after each iteration or solve. In the cases of nonlinear transient or nonlinear statics analysis, there may be many solves per output. Because of limitations in the output database format, it is difficult (or impossible) to intersperse the residuals from each solve with the usual solution output. However, it is possible to select between the standard time step and an “iteration time step”. Note that the **Exodus** database writes output for each “time step”. It uses the step number as an index to the data, and only one such index is supported. When we substitute the iteration number for the time step we can write the data properly, but once iteration has completed, we may not write data using the other index (time step, or mode number). Should that occur, we would have residuals from one iteration sharing the same time axis index with transient data. The parameters for the option are listed in Table 7-152.

Keyword	Application
standard	use time step or mode number as index
iteration	use the iteration count as index

Table 7-152. – TIndex parameters.

```
OUTPUT
  disp
  residuals
  Tindex=iteration // output on each iteration
END
```

Input 7.1. TIndex example

TIndex makes sense only in solutions that require multiple iterations per solve, such as nonlinear solutions. In other solutions, it is ignored, and output is provided at the standard time step.

NOTE: TIndex is a debugging function. As such, we do minimal checks. In some solutions, it might be possible to output data using both steps.

7.1.43. EOrient

Option **EOrient** in the **outputs** will output the element orientation vectors for all elements. The element orientation is a design quantity that normally does not change significantly through the course of an analysis. This output is provided to help in model construction and debugging.

The orientation vectors are output as nine variables that collectively make up the three vectors required for element orientation. The output variables and the associated meanings for various elements are shown in tables 7-153 and 7-154 respectively.

Table 7-153. – Element Orientation Outputs.

Name	Description
EOrient1_X	first orientation vector
EOrient1_Y	
EOrient1_Z	
EOrient2_X	second orientation vector
EOrient2_Y	
EOrient2_Z	
EOrient3_X	third orientation vector
EOrient3_Y	
EOrient3_Z	

Table 7-154. – Element Orientation Interpretation.

Element	EOrient1	EOrient2	EOrient3
Beam2	axial	first bending (I1)	2nd bending (I2)
Shells	Element X	Element Y	Normal
Solids	Element X	Element Y	Element Z
Hexshell	Element X	Element Y	thickness
ConM	null	null	null

7.1.44. Pressure

Option **pressure** in the **outputs** selects applied pressure output to the **Exodus** file as both a sideset variable and a new nodeset variable. The addition of nodeset pressure output enables restarts using the output pressure as an input load. For most applications this also provides a useful tool for checking input loads.

7.1.45. NPressure

Option **npressure** in the **outputs** will output the *nodal* pressure to the output **Exodus** file as a nodal variable. This output is only available for solutions that introduce nodal pressure (currently only the random pressure loading).

7.1.46. **APressure**

Option **APressure** in the **outputs** will output the acoustic pressure to the output **Exodus** file as a nodal variable. For purely acoustic elements, this will result in one degree of freedom per node. At the wetted interface, the nodes of node-face interactions inherit the degrees of freedom from the face, this will result in four degrees of freedom per node in the output **Exodus** file.

7.1.47. **APartVel**

Option **APartVel** in the **outputs** will output the acoustic particle velocity to the output **Exodus** file as an element variable. This is the velocity of the fluid particles. It is computed in **Sierra/SD** as the gradient of the velocity potential. For purely acoustic elements, this will result in three degrees of freedom per element.

7.1.48. **Constraint_Info**

Linear system solvers are sensitive to redundant or inconsistent constraints. Keyword **constraint_info** selects constraint information on the nodes. The information is useful in preparing models that are less sensitive to redundant constraints. Redundancy can be generated in any of the constraint types or intersections between types.

Constraint information has several fields.

MPC_Status Indicates if a given node is involved in any MPC equation. A value of '1' indicates the node is used in at least one equation, '0' otherwise.

MPC_Touched Indicates how many times a given node shows in MPC equations.

Node_Face_MPC_Count For only node-face contact constraints, this indicates how many times that a node is used as the node of a node on face constraint. A value greater than one can be problematic as a given node can only be correctly tied to a single face without introducing over-constraint.

Node_Face_MPC_Redundancy For only node-face contact constraints this highlights nodes that may be over-constrained. Pay attention to values greater than one. The **Node_Face_MPC_Redundancy** is typically one less than **Node_Face_MPC_Count** unless there are more than 3 independent constraints for a specific sideset pair, as may occur if a sideset pair is used in both a tied constraint and a slip contact, or Tied Joint.

Node_Face_MPC_Gap indicates the distance the node of a node-face constraint must be moved to be placed on the face. Many problems with constraints stem from surfaces that do not properly match up geometrically.

Node_Face_MPC_Both_Node_and_Face For node face constraints only this indicates if a particular node acts as a node in one node-face constraint and also is attached to the face of separate node-face constraint. A '1' means such situation occurs. In some cases such nodes may be part of a problematic cyclic constraint. However, in other cases this situation may be expected and cause no problems.

MPC_Origin Indicates what capability created the MPC. For example sliding contact vs. tied contact. The index number output matches a table in the rslt log file. Currently, this output is only available from node face contact constraints.

7.1.49. Statistics

For transient dynamics solutions only, summary statistical information may gathered and output for the time history of variables listed in Table 7-155. We gather information about the mean, the min/max, and the standard deviation. Data is gathered at each time step, independent of the frequency of output (e.g. **nskip** is ignored).

Because this is summary data, it is not convenient to append this data to the file used for output of the time history. Another file is generated with a **.stat** extension to store that data.

Statistical data requires *two* keywords for output. Both “statistics” and the keyword associated with that output quantity must be selected. For example, to output statistics of the force, the following output section is required.

```
Outputs
  statistics
  force
end
```

As with output in general (section 7.1.2), users can define a custom statistical output filename via the **statistics** block:

```
Statistics
  database name = <string>
end
```


Keyword	Section	Comment
Displacement	7.1.13	
Velocity	7.1.14	
Acceleration	7.1.15	
Force	7.1.34	applied force
RHS	7.1.37	Right-Hand Side vector at each load.
vonMises	7.1.18	von Mises stress, supported for echo output only.

Table 7-155. – Supported Statistical Data types for Transient Dynamics. Selection of these quantities along with “statistics” results in an addition **Exodus** file containing mean, min/max and standard deviation data.

7.1.50. KDiag

Option **kdiag** in the **outputs** will output the maximum and minimum values of the diagonal of the stiffness matrix as nodal variables KDiagMax and KDiagMin. These are the max and min of the 7 variables associated with the 3 translational, 3 rotational and 1 acoustic degree of freedom on each node. These values are primarily useful for diagnostics purposes, where they may help identify stiff regions of a model. All 7 terms may be seen by outputting KDiag in the **echo** section.

Figure [7-69](#) illustrates the use of this option. Note how the center sections of the model are highlighted by their stiffness terms. This tool is especially important for analyzing some collections of beams. Beam stiffness is proportional to $1/L^3$. A common mistake is to generate stiff beams, which can ruin the numerical solution. See Section [7.1.51](#) for a related diagnostic on the dynamics matrix. ¹

¹The stiffness diagonal and dynamic matrix diagonal depend to some extent on the linear solver used. Domain decomposition solvers generally use Lagrange multipliers to eliminate constraints, while some sparse solvers remove constraints through reductions of rows and columns of the matrices. Because the matrices to be solved are different, the diagonals and conditioning of the matrices are also different.




Figure 7-69. – Example *KDiag* output.

7.1.51. **ADiag**

Option **ADiag** in the **outputs** will output the maximum and minimum values of the diagonal of the dynamics matrix as nodal variables ADiagMax and ADiagMin. Refer to the *KDiag* section, (7.1.50), for format information.

The “dynamic matrix” is the matrix which is solved by the linear solver. The “ADiag” diagnostic can help identify regions of the model that may contribute to poor matrix conditioning. Summary of a few of the dynamics matrix terms are listed in Table 7-156. Refer to the theory manual for details of the matrix to be solved. Dynamics matrix output is available for most solvers (including GDSW), and for some solution methods.

Solution	Matrix	Comment
eigen	$K - \sigma M$	real eigen problem
transient	$K + \frac{4}{\Delta T^2} M + \frac{2}{\Delta T} C$	standard Newmark-Beta
statics	N/A	dynamics matrix is stiffness matrix
qevp	N/A	unimplemented

Table 7-156. – Selected Dynamic Matrix Definitions.

7.1.52. **ddamout**

The **ddamout** keyword is a single keyword that can be used to trigger **Exodus** output (if used in the OUTPUT section), or history file output (if used in the History section). In

Ddam analysis some of this data is written to text files by default. It is provided in **Exodus** files too for the convenience of the user.

Table 7-157 lists the nodal and element variables that are output when the **ddamout** keyword is selected. There are 8 nodal variables and 2 element variables.

Table 7-157. – Variables that are output from ddam analysis.

Option	data type	Description
ddam_mdisp	nodal variable	ddam modal displacements
ddam_mvel	nodal variable	ddam modal velocities
ddam_macc	nodal variable	ddam modal accelerations
ddam_mfor	nodal variable	ddam modal forces
ddam_nrl_sdisp	nodal variable	ddam nrl-summed displacements
ddam_nrl_svel	nodal variable	ddam nrl-summed velocities
ddam_nrl_sacc	nodal variable	ddam nrl-summed accelerations
ddam_nrl_sfor	nodal variable	ddam nrl-summed forces
ddam_mvmstr	element variable	ddam modal von Mises stresses
ddam_nrl_svmstr	element variable	ddam nrl-summed von Mises stresses

Finally, we note a couple of additional details for output of ddam data.

- In parallel runs, the text file output will not include nodal variables, since that data would not be usable in that form. Instead, that data could be written to the **Exodus** file with the **ddamout** keyword.
- History output of ddam data will only write the nodal variables, not the element variables. Element variable history output for ddam analysis is currently not in place.

7.1.53. Temperature

The **Temperature** keyword is a single keyword that can be used to trigger **Exodus** output of element temperature at the centroid. Element temperature can either be read in as input temperature or calculated from energy deposition. The element centroid temperature can originate from the reference or block temperature, and the centroid, nodal or Gauss point **Exodus** data.

7.1.54. Voltage and Charge

The **voltage** keyword is a single keyword that can be used to trigger **Exodus** output of the nodal voltages for electro-mechanical coupled models. The **nodal_charge** keyword is a single keyword that can be used to trigger **Exodus** output of the applied nodal charges.

Volume Keyword **Volume** selects **Exodus** output of element volume in the undeformed state. Volume is the integral of shape functions for solid elements. For shell elements the

volume is the area times the thickness of the elements. For bar and beam elements the volume is length times the cross-sectional area of the element.

7.2. *User Output*

7.2.1. **Block Statistics**

```
User Output
  block <list>blockNames/Ids

    compute global <string>Name1 as average|avg
      ↪ of element <string>Variable1 (weighted by volume)

    compute global <string>Name2 as max|min|maxabs|minabs
      ↪ of element <string>Variable1

    ...
end
```

The **user output** block selects the volume-weighted average of any element variable within a user specified list of blocks for output file or history file as a global variable. Multiple blocks are defined in the same way as the subdomain selection in the echo block (Section [7.7.4](#)).

Each user output can then be included by name in the **outputs** or **history** block to turn on that user output in the corresponding output file, e.g.

```
History
  Name1
  Name2

  nodeset tail
    disp
    ...
end
```

The format of the global output will then be **NameVariable**, i.e. for **Name="StressGauge "**, and **Variable="stress"**, your outputs for a solid element would be of the form **StressGaugeVStressX**, **StressGaugeVStressXY**, **StressGaugeVStressXZ**, ...

7.2.2. Closest Distance

```
User Output
  block <list>blockNamesA/Ids
  (OR) surface <list>sidesetNamesA/Ids
  (OR) nodeset <list>nodesetNamesA/Ids

  compute global <string>name1 as closest distance
    ↪ to block <list>blockNamesB/Ids
    ↪ (search node_face|node_node|mixed)

  compute global <string>name2 as closest distance
    ↪ to surface <list>sidesetNamesB/Ids
    ↪ (search node_face|node_node|mixed)

  compute global <string>name3 as closest distance
    ↪ to nodeset <list>nodesetNamesB/Ids

  compute nodal <string>name4 as closest distance
    ↪ to surface <list>sidesetNamesB/Ids

  compute nodal <string>name5 as closest distance
    ↪ to block <list>blockNamesB/Ids

  ...
end
```

The **global** option computes the closest distance between the “A” set blocks, sidesets, nodesets and the “B” set blocks, sidesets, or nodesets as a single global variable. This is the minimum separation distance between these two sets anywhere in the model in the current displaced shape. The optional search command applies to the block/sideset from the "compute global" line, the default (mixed) is `node_face`. The node face search calculates the distance from the nodes in set “A” to the closest face in set “B”. The node-node search option computes the closest distance from the nodes in set “A” to the nodes in set “B”. If the “B” set contains only nodesets only the `node_node` search can be used, so any search option there will be ignored (with a warning).

The **nodal** option computes the closest distance between each node in the “A” set to the faces in the “B” set. This output is a full field result for the closest distance calculation.

As with block statistics output (section 7.2.1), the results of the user output calculation will be output as global or nodal variables, and are enabled by including the variable name in the `outputs` or `history` section, e.g.

```

History
  nodeset tail

  Name1
  Name2

  disp
  ...
end

```

7.2.3. Derived Variables

Derived `user output` variables support the same use case as the `Statistics` file, with more granular control over variable selection, and broader support for variables to be derived. Derived `user output` variables are intended as a replacement for the `Statistics` file.

Statistical `user output` values do not match values in the `Statistics` file. The cause of this is the inclusion of the model's initial state in `user output` calculations, and its exclusion from `Statistics` file calculations. The `OutputInitialTime` parameter has no effect on either calculation.

```

User Output
  compute element <string>name1 as average|standard deviation|rms|min|max
    ↪ over time of element <string>var1

  compute nodal <string>name2 as average|standard deviation|rms|min|max
    ↪ over time of nodal <string>var2
end

```

This `user output` block computes two new derived variables `name1` and `name2` by applying an operator over time to the built-in variables `var1` and `var2`. These derived outputs require time histories, and are only enabled for transient solution cases.

Supported operators include minimum(`min`), maximum(`max`), mean estimate(`average`), the Root-Mean-Square (`rms`), and the `standard deviation`. Note that the `standard deviation` operator uses Bessel's correction in its calculations. That is, we define the variance as:

$$\frac{1}{N-1} \sum_{i=1}^N (x_i - \bar{x})^2$$

not

$$\frac{1}{N} \sum_{i=1}^N (x_i - \bar{x})^2$$

.

The source variables `var1` and `var2` must be built-in variable request strings; "chaining" derived variables together is not currently supported. Additionally, these variable request strings must match variable requests in other contexts; `Stress` and `Displacement` are valid, but `StressXX` and `DispX` are not.

Most built-in variable requests represent several variable fields in the results file. Therefore, derived variable fields are defined by prepending the user string `name1` to the built-in field names `StressXX`, `StressYY`, `StressZZ`, etc.. The result in this case would be `name1StressXX`, `name1StressYY`, `name1StressZZ`, etc..

Every built-in element request and *most* built-in nodal requests are available for use in derived outputs. Some very specialized nodal requests are not supported for derived outputs, due to the unusual nature of their calculations. For example, variables derived from `AcousticLighthill` will always report 0. This is a known deficiency which can be addressed as needed.

7.3. History

All the results from the "OUTPUT" section can be output to a limited portion of the model using a history file. Only those outputs described in Table 7-141 are supported. If the output is also specified in the OUTPUT section, there is little need to write the data in the history file. The following output section options are ignored in the history section because all history file output will be in the **Exodus** format.

- MFile
- Kaa, Maa, Faa
- vrms

In addition to the **outputs** selection options, the history file section contains information about the regions of output. The default is NO output selection. Selection may be for node sets, side sets, a node list file (see Section 6.1.5), or element blocks. Virtual blocks can be included in this section. For example, one could output the element force in a virtual Joint2g element. If side sets are selected, the side set selection is for the nodes associated with that side set, not for the elements themselves. All nodal variables selected in the history file will be output for all selected nodes. Selecting an element block automatically selects the associated nodes in that block. The format for the selection of multiple regions uses a series of MATLAB concatenated ranges, as with subdomain selection in the echo block (Section 7.7.4). Additionally, the keyword *all* can be used to select all nodesets, sidesets, or blocks for output. For example,

```
History
  nodeset tail,1:10,17
  sideset all
  nodeset '8,15' coordinate 4
```

```

block '5,6,nose'
stress
disp
nskip 10
element centroid stress nearest
    location 0 1 2 as StressGauge
element strain nearest location 1 2 1 as StrainGauge
element stress at element 1 as el1StressGauge
element strain nearest center of
    sideset wing as SS1StrainGauge
node acceleration at node 7 as nd7AccGauge
node displacement nearest location 2 2 10 as dispGauge
end

```

Any number of nodeset, block, and/or sideset selections can be specified in the history section. Nodeset and sideset specifications may be followed by an optional coordinate entry. If a **coordinate** is specified (see section 2.7), all nodal results for the nodes in the nodeset/sideset are transformed to the specified coordinate system before output to the file. If a particular node is identified in more than one specification, the last specification is used for the output. Note that the **coordinate** keyword for history section output will only work with nodesets and sidesets: it is not supported for blocks or node list files.

The coordinate *name* of nodes in the history file may be printed out in the echo file by specifying **nodes** in the **echo** section of the input. The coordinate *ID* will also be written to the history file (as a nodal variable *CID*) provided any nonzero coordinate frames have been specified. Note that if the coordinate name is not convertible to an integer id (i.e. if it is a string), the *CID* output will be -1.

Note that for output specifications, the corresponding nodeset, block, and sideset ids/names should be specified as shown in the above example, with either a list of ids/names, or the keyword “all”. The rules for specifying these lists are the same as for integer lists, and are detailed in Section 2.1.

In transient dynamics solutions, user control of output step interval “nskip” and output buffer “flush” operations are provided to increase efficiency of output. See Section 3.32 for examples. The history file respects the **nskip** and **flush** parameters set in the solution block, but additional user control is provided for history files by inserting the **nskip** and **flush** parameters in the history block. In that case, history files for all multicaselutions will have output and buffer flushing at the intervals specified in the history section, and the entries in the solution section will be ignored for history files.

Unlike subdomains, node set and side set IDs need not be contiguous in the **Exodus** file. If the selection criteria is a range, it may identify nonexistent sets. These will be silently ignored. In example 7.3 above in which nodesets **tail**, 1:10 and 17 are selected, the **Exodus** mesh must have nodeset number 10. Nodeset, sideset, and block names/IDs in the history file must be consistent with the corresponding **Exodus** input file.

As with output in general (section 7.1.2), users can define a custom history output filename via the **database name** keyword. Additionally, **Exodus** properties can be defined as shown in Section 7.1.3.

Only one history file will be written per analysis. The name of the history file is derived from the name of the **Exodus** output file, except that the extension is “.h”.

While the history file provides a convenient means for transforming coordinates, its applicability may be somewhat limited when output in many coordinate frames is desired. In particular, only a single history file is written in each analysis, and only one coordinate frame may be output per node. The history file will display variables as Cartesian regardless of coordinate choice. Table 2-23 shows the corresponding values for cylindrical and spherical coordinates.

7.3.1. Global History Output Near a Location

The history data of a mesh entity near a location can be written to the output file as a global variable using any of the following syntax:

```
NODE|ELEMENT (optional, for element)CENTROID
  <string>variable_name
  NEAREST LOCATION <real>global_x,
    <real>global_y, <real>global_z
  AS <string>history_variable_name
```

```
NODE|ELEMENT <string>variable_name
  AT NODE|ELEMENT <int>global_id AS
  <string>history_variable_name
```

```
ELEMENT <string>variable_name
  NEAREST CENTROID OF SIDESSET
  <int(/string)>sideset_id(/sideset_name)
  AS <string>history_variable_name
```

Outputs at multiple locations can be requested from a history block. Each output location needs a unique **history_variable_name**. This capability is currently only supported for outputting data at the closest location for 3D elements or the closest centroid for 2D and 3D elements (using the optional **CENTROID** flag). If the closest location to the centroid of a specified sideset is requested, the data is obtained at the nearest location on a 3D element *surface* that touches the sideset. If the data at a specified element is requested, the data is obtained from that element centroid for 2D and 3D elements. Trying to specify a 1D element is not supported. If a 1D element is specified directly, or if no elements can be found that satisfy the search conditions, the .rslt output summary will contain a statement that no element was found.

Note: User-specified element, nodal, and sideset IDs should be base-1 to match the **Exodus** convention.

All element output cases are currently only supported for stress, strain, and von Mises stress output. Likewise, displacement, velocity, acceleration, and force outputs are the currently supported nodal outputs. A similar capability and syntax is used in **Sierra/SM**. For the example given in the history block from the previous section, a summary of the global history data output will be written to the .rslt file in the following format:

===== GLOBAL HISTORY OUTPUT =====					
Prefix	Block	Entity, GId(: Type)	Variable	Near	At
StressGauge	1	element 1: Hex8b	stress	(0,1,2)	(0.5,0.5,0.5)
StrainGauge	1	element 5: Hex8b	strain	(1,2,1)	(1,2,1)
StressGage	1	element 1: Hex8b	stress	N/A	(2.5,1.5,0.5)
StrainGage	1	element 6: Hex8b	strain	(2,3,1)	(2,2,1)
AccelGauge	N/A	node 7	accel	N/A	(1,1,0)
dispeGauge	N/A	node 38	disp	(2,2,10)	(2,2,1)
=====					

In the **Exodus** history file output, the global variable's prefix or `history_variable_name` will be prepended onto the element variables output name (see sections 7.1.18 and 7.1.16 for stress and strain variable names). For example, if you had two requested outputs near a shell centroid and a solid, with prefixed names "shellStrain" and "solidStress" respectively, the global history output would produce the following global variables in the history file:

```
shellStrainSStrainX1, shellStrainSStrainY1, shellStrainSStrainXY1,
shellStrainSStrainX2, shellStrainSStrainY2, shellStrainSStrainXY2,
shellStrainSStrainX3, shellStrainSStrainY3, shellStrainSStrainXY3,
solidStressVStressX, solidStressVStressY, solidStressVStressZ,
solidStressVStressXY, solidStressVStressXZ, solidStressVStressYZ
```

Variables requested for global history output will be output on every other element block in the history block.

7.4. Frequency

The frequency section provides information for data output from the ModalFrf, directFRF, shock, modalshock, and random vibration solution methods. One frequency file is written per analysis. The name of the frequency file is derived from the name of the **Exodus** output file, except that the extension is ".frq". The section format follows that of the history section. As in the case of the history section, data can be written to a sideset, nodeset, node_list_file, or a block. In the case of output to a block, the block can be a virtual

block. Thus, one could output element force on a Joint2g element. Solution methods that do not write frequency domain output silently ignore the Frequency section.

The frequency section also includes the definitions of the frequency values for calculation. A frequency section (with some output selection region) must be selected for any solution method requiring frequency output. To fail to do so is an error, since the solution would be computed and no output provided.

As with output in general (section 7.1.2), users can define a custom frequency output filename via the **database name** keyword. Additionally, **Exodus** properties can be defined as shown in Section 7.1.3.

The frequency values may be specified using the methods specified in Table 7-158. The methods are mutually exclusive, i.e. do not mix keywords from the “linear” method with those of the “table” method. An example follows.

Table 7-158. – Frequency Value Specification Methods.

Method	Keyword	Description
method=linear	freq_min	minimum frequency (typically in Hz)
	freq_max	stop frequency
	NF	number of frequency intervals.
	freq_step	frequency increment (or use NF)
method=log	freq_min	minimum frequency (typically in Hz)
	freq_max	stop frequency
	NF	number of frequency intervals.
method=table	table	name of a 1D table (see Section 2.8.19)

```

FREQUENCY
  nodeset '1:10,17'
  sideset '3:88'
  block 5,6,3
  disp
  acceleration
  freq_min=10    // starting frequency in HZ
  freq_step=10   // frequency increment
  freq_max=2000  // stop freq. Outputting 201 freq pts.
END

```

For the “linear” and “log” methods, the frequencies are obtained by the following equation.

$$F_k = \begin{cases} F_{min} + k \cdot F_{step} & \text{for method=linear} \\ F_{min} \exp(kD) & \text{for method=log} \end{cases}$$

Here $D = 1/N_F \log(\frac{F_{max}}{F_{min}})$. If both **freq_step** and **NF** are specified, **NF** is used.

Output Region:

The controls in the **frequency** section also affect data written to the results (or echo) file. In particular, the echo file contains data only for those nodes in the selection region of the **frequency** section. Selection of a specific output (such as displacement or acceleration) is independent. For example, you may echo only displacements, but write displacements and accelerations to the **Exodus** frequency output file. The history section (7.3) has more information on specification of the output region.

The SIERRA translator *exo2mat* may be used to translate the output into MATLAB format for further manipulation and plotting.

7.5. Linesample

The line sample (**LineSample**) section of the input file provides a means of evaluating and outputting fields or internal variables at sampling points within a structure. These sampling points are defined on a series of lines.

Section 3.34 discusses the primary application of line sample, verification of stress field input to **Sierra/SD** from TSR. Line sample is used for energy deposition (see *Two Element Exponential Decay Variation Hex20* in the Verification manual⁵⁴). Energy deposition is interchangeable with supplying an applied temperature. Line sample is also used for far-field processing in acoustics problems (see 6.1.9.1 or How To⁵³), for example with infinite elements.⁶⁸

Keywords for the line sample input are listed in the table below. An example follows.

Keyword	Arguments
samples per line	<i>integer</i>
endpoint	<i>6 real numbers</i>
format	<i>string</i>
nskip	<i>integer</i>

samples per line The number of sample points on each line. All lines will have the same number of samples.

endpoint The endpoints of the line. There should be 3 real numbers for the XYZ location of the beginning of the line, followed by 3 real numbers at the end. There can be any number of endpoint entries.

format The format of the output file. Two output formats are supported: **Exodus** and MATLAB **MFile**. The default is **MFile**.

nskip Results output frequency; defaults to the value specified in the solution section; see Section 3.32.0.4.

Output from the line sample is written to `linedata.m` for MFile output, or to `linedata.exo` for **Exodus** format. There is no need to join this data for parallel runs. In those output files, a nodal variable called **Displacement** will be created. The entries in this array correspond to 3 displacement variables, 3 rotation variables, acoustic pressure, voltage, and temperature. For transient data, the time values are also output for each of these arrays.

```

LineSample
    samples per line 5
    endpoint 0. 0. 0.    1. 1. 1.
    endpoint 0.0 0.5 0.5    1. 0.5 0.5
    format exodus
END

```

7.6. *Stresses and Strains*

Stress and strain values at element centroids are available. Solid and beam element stress is evaluated in the global or basic frame. However, shell element stress and strain are evaluated in element space.

The total strain (7.1), elastic strain (7.2), and thermal strain (7.3) can all be requested separately; see Section 7.1.16. The stress reported is always the Cauchy stress $\sigma^c = C : \epsilon^e$; see Section 7.1.18.

7.6.1. **Stress/Strain Truth Table**

The **Exodus** data format provides an element truth table. Element variables are defined globally (for all element blocks), but output data is stored only for those blocks that have entries in the truth table. Thus, in **Sierra/SD** if stress output is requested (see Section 7.1.18), then stress variables are defined for solids and shells.² Space is allocated in the output **Exodus** file, and data is written only if it is applicable. Table 7-159 illustrates this for stresses. A similar table can be generated for strains. Note that volume stresses always start with “V” and surface stresses start with “S”. Note that “vonMises” is the only entry that applies to both solids and shells.

²The variables are defined for solids and shells even if only one or the other occurs in the model

Table 7-159. – Element Stress Truth Table.

Variable Name	Element		
	Solid	Shell	Beam
SStressX1		σ_{xx}^{top}	
SStressY1		σ_{yy}^{top}	
SStressXY1		τ_{xy}^{top}	
SvonMises1		σ_{vm}^{top}	
SStressX2		σ_{xx}^{mid}	
SStressY2		σ_{yy}^{mid}	
SStressXY2		τ_{xy}^{mid}	
SvonMises2		σ_{vm}^{mid}	
SStressX3		σ_{xx}^{bottom}	
SStressY3		σ_{yy}^{bottom}	
SStressXY3		τ_{xy}^{bottom}	
SvonMises3		σ_{vm}^{bottom}	
VStressX	σ_{xx}		
VStressY	σ_{yy}		
VStressZ	σ_{zz}		
VStressYZ	σ_{yz}		
VStressXZ	σ_{xz}		
VStressXY	σ_{xy}		
VonMises	σ_{vm}	$\max(\sigma_{vm})$	
Signed_VonMises	$\pm\sigma_{vm}$		
Max_Principal_Stress	$\ \sigma_1\ $		
Intermediate_Principal_Stress	$\ \sigma_2\ $		
Min_Principal_Stress	$\ \sigma_3\ $		
Max_Principal_Stress_x	σ_{1x}		
Max_Principal_Stress_y	σ_{1y}		
Max_Principal_Stress_z	σ_{1z}		
Intermediate_Principal_Stress_x	σ_{2x}		
Intermediate_Principal_Stress_y	σ_{2y}		
Intermediate_Principal_Stress_z	σ_{2z}		
Min_Principal_Stress_x	σ_{3x}		
Min_Principal_Stress_y	σ_{3y}		
Min_Principal_Stress_z	σ_{3z}		
ElemForce			forces

7.6.2. Solid Elements

If stresses are requested, solid elements will output the values of stress at the element centroid. The values reported are the engineering stresses in the global coordinate system.

If principal stresses are requested, solid elements will output the three principal stress vectors and magnitudes in order of descending signed value. The vectors are of unit length, defined in the global coordinate system.

7.6.3. Shell Elements

Shell elements introduce two complexities to stress/strain recovery. First, it is often important to recover data from the virtual surfaces of the elements (where the stresses are highest). This requires data recovery at the top, mid-plane and bottom surfaces. Second, there are no stresses or strains normal to the surface. Thus, stresses are naturally reported in the surface of the element. This can also introduce confusion about the in-plane coordinate frames. As shown in Figure 7-70, the stresses and strains are recovered in the physical space x_1, x_2 coordinate frame, which has been mapped from the η_1, η_2 frame in element space. Note that the direction of the x_1 vector depends on the ordering of the mesh, and may vary from element to element in the same surface mesh. Element orientation vectors are selected with the **eorient** keyword described in Section 7.1.43. Von Mises stress, an invariant, is independent of the element orientation.

Stress recovery for the **TriaShell** is interesting. A **TriaShell** is a shell element created by combining Allman's triangle² with a DKT element.⁸ Its stress vector $\vec{\sigma} = (\sigma_x, \sigma_y, \sigma_{xy})^T$ is the sum of stresses $\vec{\sigma}_{at}$ for Allman's triangle and $\vec{\sigma}_{dkt}$ for the DKT element,

$$\vec{\sigma} = \vec{\sigma}_{at} + \vec{\sigma}_{dkt}. \quad (7.5)$$

Allman's triangle represents the membrane dof, i.e., (u, v, θ_z) . If the element lies in the x-y plane, then β_x and β_y are rotations of the normal to the undeformed middle surface in the x-z and y-z planes, respectively.

$$\{\kappa\} = \begin{bmatrix} \beta_{x,x} \\ \beta_{y,y} \\ \beta_{x,y} + \beta_{y,x} \end{bmatrix} \quad (7.6)$$


$\{\epsilon\}$ is the strain vector, and $[D]$ is the elasticity matrix for Allman's triangle. The stresses through the three surfaces of the shell element are the same. Therefore,

$$\vec{\sigma}_{at} = [D]\{\epsilon\}. \quad (7.7)$$

For the DKT element, z is the coordinate direction normal to the element, with $z = 0$ representing the mid-plane. $[D]$ is the elasticity matrix.

$$\vec{\sigma}_{dkt} = z[D]\{\kappa\} \quad (7.8)$$

Figure 7-70. – Tria3 Stress Recovery. Stresses are output in the orthogonal x_1, x_2 coordinate frame in physical space, which has been mapped from the η_1, η_2 frame in element space.



$\bar{\sigma}_{dkt}$ does vary with the thickness of the element. Note, the above stress equations are written with respect to a local element coordinate system as shown in Figure 7-70.

Combining the stress vectors from Allman's triangle and the DKT element above yields the stress vector for the element which is output in the local element frame.

For composite elements (such as QuadT, Quad8T and Tria6), the stresses are computed from the underlying Tria3 element and then transformed to the element orientation of the composite element. For the quad elements, the stress of the two central triangles is averaged. Figures 5-34, 5-35 and 5-36 describe these composite elements.

7.6.4. Beam Elements

Reporting stresses for line type elements (Beams, Rods, Springs, etc) is even more problematic than it is for shells. For many of these elements an axial stress could be reported. But, for beam elements that stress could not include the effects of beam bending unless details of the beam cross section were available. For some elements (such as a spring) no concept of stress is even correct. As a consequence, we do not report stresses for most line type elements. However, some recovery may be obtained using the element force output (see Section 7.1.38).

For Beam2, Nbeam, and TiBeam elements (sections 5.9 to 5.11), axial stress will be reported if **stress** is requested in the **outputs** section. Additionally, stress recovery points may be requested using the **stress recovery point** keyword. For each stress recovery point, the bending and 2 shear stresses will be reported, up to a maximum of 6 points (see input 7.2). Each stress recovery point represents a point on the beam cross-section and is defined by a pair of offsets in physical coordinates from the neutral axis of the beam. In

either case, the maximum von Mises stress over all stress recovery points will also be reported. If no stress recovery points are requested, this is equivalent to the absolute value of the axial stress.

```
BLOCK beam_block
  material=1
  beam2
    Area=8
    I1=2
    I2=10
    J=7
    stress_recovery_point  0.5  0.5
                           -0.5 -0.5
END
```

Input 7.2. Beam Stress Recovery Points Example

7.7. *Echo*

Results, in ASCII format, from the various intermediate calculations may be output to a results file, e.g. `example.rslt`, where the filename is generated by taking the basename of the text input file (without the extension) and adding the extension `rslt`. Output to the results file is selected in the **Sierra/SD** input file using the **echo** section. An example is given below, and the interpretation of these keywords is shown in Table 7-160.

```
echo
  materials
  elements
  Jacobian
  mesh
  input on
  nodes
  MPC
end
```

We remark that for virtual blocks, element variables such as element force are also written to the results file. Since only Joint2g elements are currently supported as virtual blocks, the only element variable that can be written at this time is the element force, **eforce**.

Table 7-160. – Echo Section Options.

Option	Description
ADiag	diagonal of dynamics matrix
acceleration	nodal accelerations (better in output section)
block	block wise mass properties (used only following mass)
debug	debug output
displacement	nodal displacements (better in output section)
EForce	element force for beams
ElemEigChecks	element eigenvalues
elements	element block info, i.e. what material, element type, etc
Elmat	element material properties
energy	element strain energy and strain energy density
eorient	element orientation
fatigue	fatigue related parameters
force	applied forces (better in output section)
GEnergies	global kinetic and strain energy sums
block_energies	block kinetic and strain energies
input (<bool>)	echo of post-Aprepro input (<i>for parse errors</i>) – default = input on
input_summary	summaries of many sections
Jacobian	block summary of Jacobians
KDiag	diagonal of stiffness matrix
line_weld	Line-weld-specific output variables
mass	mass properties in the basic coordinate system
materials	material property info, e.g. E, G
memusage	prints per processor per task memory use to results file and an external text file
mesh	summary of data from the input Exodus file
mesh_error	mesh discretization error metrics
ModalVars	modal force and amplitude for modal solutions (echo section)
MPC	MPC equations
NLresiduals	turns on residual output per iteration of the Newton loop for nonlinear solution methods
nodes	nodal summary
residuals	residual vectors
rhs	Right Hand Side vector (better in output section)
strain	element strains at centroids
stress	element stresses at centroids
subdomains “0:3:6,10”	Controls which processor will output results file
threading_summary	threading summary table
timing_summary	timing and threading summary tables
velocity	nodal velocities (better in output section)
vonMises	von Mises stress only
vrms	RMS quantities (random vibration only)

7.7.1. Mass Properties

The mass properties may only be reported in the **Echo** section (i.e. at this time there is no mass property report in the **outputs** section). The mass properties are the total mass, the center of gravity and the moments of inertia of the system. All are reported in the basic coordinate system. Note that moments are about the origins, not about the center of gravity. Masses are reported in a unit system consistent with the input, with or without the **wtmass** parameter (see Section 2.3).

An additional option of **block** may be used in the echo section to output the block wise mass properties to the results file. Please note that the block wise mass properties, though summed for all processors (if running on a parallel machine), are only output to the result file from the first processor (processor 0). The block wise mass properties option, called **block**, reports the number of blocks, the mass of each block, and the center of gravity of each block along the x, y, and z axis. Please note that **block** may only be used in the **echo** section following the **mass** option as shown below.

```
echo
  materials
  elements
  mass=block
  nodes
end
```

If the keyword **mass** does not directly precede **block** in the **echo** section, then **Sierra/SD** will abort with the following error.

```
Unrecognized "echo" option "block".
Aborting.
```

Requesting both **massblock** and **mass=block** causes output to the result file of only the global mass properties. If only block-level mass properties are desired, then specifying **mass=block** will suffice, as follows

```
echo
  mass=block
end
```

7.7.2. Multipoint constraints

Text descriptions of the MPC equations are output to the result file using **MPC**. This is a check on the input deck. An example of the output format is as follows

```

MPC
  coordinate 0
    25 P 1
    106 P -1
  // G = 0.000000
  // the source is global
END

```

In this case, the MPC equation is constraining the acoustic pressure in nodes 25 and 106 to be equal in the global (default) coordinate system.

7.7.3. ModalVars

modalvars text output which contains modal forces and modal amplitudes for modal based superposition solutions including “modaltransient” and “ModalFrF”. Two text files are written: **Qdisp.txt** and **Qforce.txt**. Each line of the file contains data for a solution increment (a time or frequency step). For transient solutions, each column corresponds to a mode in the solution. Because FRF solutions are complex, two adjacent columns describe the complex modal amplitude (or force) for a mode. In terms of the physical force at time t_n , $F(t_n)$, and the i th eigenvector ϕ_i , the modal force and displacements are

$$f_{q_i}(t_n) = \phi_i^T F(t_n), \quad u(t_n) = \sum_i \phi_i q_i(t_n).$$

The expressions in the frequency domain are similar.

The text files are readable by either MATLAB or MS/excel.

7.7.4. Subdomains

In parallel calculations, one results file is written per subdomain. Only data associated with that subdomain are written to the file. By default, results are only written to subdomain 0 (the root processor), and in the results for subdomain 0 will *always* be output. Use the “subdomains” option to specify additional subdomains for which data will be written. The **subdomains** specification is made using a MATLAB like string, as detailed in Section 2.1. For example,

```
subdomains '0:2:8'
```

selects subdomains 0, 2, 4, 6 and 8 (again, subdomain 0 will always be output, and so it is redundant here). The following selects subdomains 0, 2, 3, 4, 6, 8, 9 and 15.

```
subdomains 0:2:8,3:3:9,15
```

In addition, the keyword “all” selects all subdomains.

7.7.5. Memusage

The **Memusage** keyword selects memory usage information output for the results file. In addition, it also requests that a per processor, per task break down be written to a separate text file. The memory in the text file is shown in Mbytes. The output for the text file will be formatted as follows:

```
1   2   3   4
40  41  41  42  "initialization"
23  12  15  15  "assembling matrices"
45  56  65  54  "initializing solver"
10  25  12  52  "writing output"
```

The primary use of the memory usage printing is to diagnose where or why memory is being exhausted by an analysis.

Elmat The bulk modulus, mass density and shear modulus of each element are reported if the **Elmat** keyword is added to the echo section.

8. CONTACT

8.1. Tied Surfaces

Tied surfaces provide a mechanism to connect surfaces in a mesh that will always be in contact. Because the surfaces are always tied, the constraints may be represented by a set of linear multipoint constraints [2.9](#). Tied surfaces are also known in the literature as *glued surfaces* or as *tied contact*. They are used almost exclusively to combine two surfaces of a mesh that have not been meshed consistently.

There are a number of ways of combining surfaces that have not been consistently meshed. The simplest method constrains the nodes of one surface (node-surface) to lie on the faces of another surface (face-surface). In this method, the constraint is called *inconsistent* because the mesh does not ensure that linear stress will be maintained across the boundary. The stress and strain in the region of the constraint will be wrong. However, loads are properly transferred across the boundaries, so a few element diameters away from the boundary, the stresses and strains should be approximately correct.

Tied surfaces can currently be specified for structural-structural interfaces, acoustic-acoustic interfaces, and structural-acoustic interfaces (i.e. wet interfaces). The syntax in the **tied data** section is the same. For structural-structural interfaces, the nodal displacements on the node-surface are constrained to lie on the faces of the face-surface. In the last case, the nodal acoustic pressures on the node-surface are constrained to match the interpolated value of pressure on the face-surface.

For tied structural-acoustic interfaces, it is necessary to ensure a weak continuity of both stress and displacement (velocity) across the wet interface. [72,52](#) Also, we recommend that the acoustic surface be defined as the face-surface (and hence should have its sideset number listed first in the input deck). Defining the structural surface as the face-surface sometimes causes an error related to singular subdomain matrices.

A simulation may have multiple tied surfaces as long as certain requirements are met. For example acoustic-acoustic and tied structural-acoustic data blocks in the same input file are supported. However, it is necessary that each sideset be exclusively attached to either structural elements or acoustic elements. A sideset containing both acoustic and structural elements is not supported. This does not restrict the possible types of analysis. It can increase the number of **tied data** blocks. However, this extra input will reduce confusion and likely also reduce potential modeling errors.

The keyword **transverse** controls the constraints on transverse displacements. Transverse displacements can be **tied** or **slip**. Its default is **tied**. The tied option is the standard inconsistent tied surface approach. The slip option only constrains normal degrees of freedom between the surfaces. In this option, the tangential degrees of freedom are free to slide. This would be the case if there was no friction between the surfaces. The **friction** option for specifying a simple friction model is not currently supported.

In a **tied data** block, if the tied contact is inconsistent, i.e., the method is not mortar or other options, then by default the gap is removed. Gap removal only works with **tied data**. In the tied data section, the order of the surface ids is important. The first surface id becomes the face surface, and the second becomes the node-surface. Gap removal moves node-surface nodes to the face surface. Set **gap removal** to *off* to skip gap removal. ACME's gap and push back vector quantities provide the gap. Updated coordinates instead of the original coordinates appear in the output **Exodus** file, and are also used by the system matrices.

Debugging tied contact is easier using the gap removal solution case [3.38](#).

8.1.1. Contact Normal Vectors

For all the contact type interactions, including tied surfaces, tied joints, and contact definition the algorithms used restrict the search to matching faces that have opposing normal vectors. For solids, this is seldom an issue. The normal vectors for a solid are always outward from the solid, so two interacting solids (unless they occupy the same volume), will naturally have opposing normal vectors. However, the situation for shell-shell or shell-solid interactions can be more complicated.

Sidesets may be created from the top or bottom surfaces of the shells. Thus, the shell surface has a natural normal direction determined by its connectivity, and the sidesets generated from the shells have a direction too. The sideset direction may align *or oppose* the direction normal of the shell itself. If the shell normal does not oppose the normal of the mating surface, no interactions will be found, and the surfaces cannot be tied. See [Figure 8-71](#).




Figure 8-71. – Shell Normal in Contact or Tied Interactions.

8.1.2. Mortar Methods

Mortar methods may also be used to tie the surfaces. This is currently under development, but some capability is available. Large tied surfaces using the mortar methods may have many fully coupled constraints which can overwhelm most parallel solvers. The cost in computing the mortar contribution is higher than the *inconsistent* method, but the solution will typically be much better in the region of the constraint.

Two different mortar methods are available. Both constrain surfaces together in an integral (or weak) sense. Standard mortar methods are somewhat simpler, but can result in a constrained system which fully couples all the nodes of both surfaces together in a single constraint. Dual mortar methods are much more friendly to the linear solver, as the constraint system decouples the constraints similarly to what is seen in node-face contact. The dual mortar method is the default.

Mortar methods are specified by adding **mortar** to the **tied data** block. To select the type of method, standard or dual, in the **parameters** block specify **MortarMethod=standard** or **MortarMethod=dual** respectively.³

8.1.3. Node to Face

Tied surfaces are specified by a listing of face-surface and node-surface side sets. Any number of tied data blocks may be specified in the input. Each tied data block represents a *single* logical pairing of constraint side sets.

```
TIED DATA
  tied faces 12
  tied nodes 18
  name "tying_12-18"
  transverse slip
  search tolerance = 1e-7
  edge tolerance = 1e-8
  gap removal = on
END
```

In the example above, sideset 12 is the face-surface. Side set 18 is the node-surface. Each node in the node-surface may be tied to the nearest face in the face-surface by a constraint equation. The transverse degrees of freedom are allowed to slip in this example. If the **transverse** keyword were omitted, standard tied surfaces would be used.

Tied surfaces use a node-to-face search algorithm. In this algorithm, the “search tolerance” represents the normal distance from a node on one surface to a corresponding face on the other. Thus, the search tolerance will typically be small and represents the amount the two

³There is no means of applying standard mortar methods to some interactions, and dual mortar methods elsewhere.




Figure 8-72. – For node-to-node searches the *search tolerance*, must be large enough to capture nearby nodes. For node-to-face searches (as used in tied surfaces), it should only capture the nearby surface.

surfaces may not be coincident. This is in contrast to a node-to-node search, where the “search tolerance” represents a search radius. See Figure 8-72.

Special care should be used when using the “edge tolerance”. If this tolerance is too large, non-intuitive interactions can be created.

Note: The current implementation ties a face-surface that consists of the two-dimensional faces of shell or solid elements. It is not possible to tie a node to the one dimensional edges of shell elements.

The relevant parameters for tied surfaces are shown in Table 8-161.

Smoothing parameters may be needed to control smoothing of the normal. Figure 8-73 illustrates the normal definitions on a faceted surface. The discontinuity in normal vectors can be an important consideration on curved surfaces where faceting affects tangential sliding. Smoothing parameters are illustrated in Figure 8-74 and include the following .

smooth angle If an angle between two faces exceeds this value (in degrees), then the angle is considered to be “sharp”, and no smoothing is done. Default is 30°.

smoothing resolution The resolution method can be either node based, or edge based. This may be needed to control smoothing on edges that include both a sharp and a non-sharp edge. Default=node.




Figure 8-73. – Normal Definitions on Faceted Geometry. When low order elements are used to describe a curved boundary, the normal is poorly defined at the edge of the facets.

Table 8-161. – Tied Surface Parameters

Parameter	type	description
Name	<i>String</i>	name of tied data block, useful for diagnosing error messages, defaults to tied
Surface	<i>integer pair</i>	face-surface and node-surface sidesets separated by comma or space
Tied Faces	<i><sideset></i>	face-surface sideset
Tied Nodes	<i><sideset></i>	node-surface sideset
Search Tolerance	<i>Real</i>	face normal of search tolerance defaults to 1e-8
Edge Tolerance	<i>Real</i>	search tolerance beyond an edge facet defaults to 1/10 search tolerance.
Method	<i>String</i>	inconsistent (<i>default most solvers</i>) mortar
Transverse	<i>String</i>	tied (<i>default</i>) slip (<i>transverse displacements can slip</i>)
Gap Removal	<i>String</i>	Yes (<i>default: for inconsistent only</i>) No
smooth angle	<i>Real</i>	maximum angle for smoothing (def=30)
smoothing resolution	<i>String</i>	“node” or “edge” based




Figure 8-74. – Smoothing Parameters for Surface Normal Vectors. No smoothing occurs for faces that are misaligned by more than the specified “smooth angle”. Within the “smooth distance”, normal vectors vary linearly with relative distance from the node.

8.2. Contact Definition

The contact definition block provides flexible syntax to define tied MPCs. The contact definition provided a similar capability to Tied Data (see Section 8.1) but with a more powerful and flexible ways to define contact surfaces. The contact definition capability leverages the same Dash contact algorithms used in Sierra SM. Both the syntax and capability of the Dash contact definition is compatible with SM recommending Dash contact for both SM/SD **hand-off** analyses and analyses with a lot of general contact.

A brief description of contact definition commands is given here. More detailed descriptions of how contact surface and interaction definition commands function can be found in the Sierra Solid Mechanics User Manual.⁶⁹

The full set of available contact commands are as follows:

```
Begin Contact Definition <name>
  contact [surface|nodeset] <surfName> contains <strings>
  skin all blocks = on [exclude <names>]
  gap removal = off|on(on)

  Begin Interaction Defaults
    general contact = off|on(off)
    self contact = off|on(off)
    normal tolerance = <Real>
    tangential tolerance = <Real>
    constraint formulation = Node_Face|Face_Face(Node_Face)
    friction model = tied|frictionless(tied)
  End Interaction Defaults

  Begin Interaction
    normal tolerance = <Real>
    tangential tolerance = <Real>
    side a = <string_list>
    side b = <string_list>
    surfaces = <string>
    interaction behavior = No_Interaction
    constraint formulation = Node_Face|Face_Face
    friction model = tied|frictionless
  End Interaction

  Begin Dash Options
    separate disconnected mesh components = true|false(false)
    ignore_shells = true|false(false)
    cutoff variable <nodal_variable> >|< <Real>
  End Dash Options
End Contact Definition
```

Information is presented in the format `command = option (default) format`. For example gap removal is on by default. Debugging contact is easier using the gap removal solution case 3.38. Gap removal is another name for initial overlap removal.

For a surface in the `begin interaction` section to define interacting surfaces, it must first be assigned a local name. To define a surface named `surfName`, use

```
contact surface <surfName> contains <entityStrings>
```

Here `'entityStrings'` may reference blocks, sidesets or nodesets in the input mesh. It may reference a block, surface, or node set by its name. Or mesh entities can be accessed using the **Sierra/SM** convention, by index using `'block_##'`, `'surface_##'` or `'nodelist_##'`.

General contact is a far-reaching capability. If there are 2 (i.e. more than one) interaction, and general contact is on, then the contact module automatically detects all interactions between any pair of surfaces (except self interactions, and assuming that self contact is off).

8.2.1. Defining Contact Surfaces

Each contact definition block is a self-contained description of contact surfaces and how those contact surfaces interact. Several options are available for defining contact surfaces as shown in the following examples.

```
// Create contact surface from exterior skin of blocks
contact surface fixture contains block_1 block_9 block_12

// Create contact surface from union of sidesets
contact surface boltFlange contains surface_1 surface_10

// Short cut to create contact surface based on the
// exodus part names.
contact surface block_7
contact surface surface_10
contact surface bearing contains bearing1 bearing2

// Create contact nodeset from union of node sets
contact nodeset ns1 contains nodelist_7 nodelist_11

// Create contact nodeset containing all nodes in a
// set of blocks
contact nodeset ns2 contains block_15 block_19

// Generate exterior skins for all blocks
skin all blocks = on
```

```
// Generate exterior skins for most blocks
skin all blocks = on exclude block_7 block_11
```

The skin of a finite element block contains all the exterior faces. When using 'skin all blocks' one contact surface is created for each block in the mesh, the contact surface is given the same name as the block. The block skinning algorithm is described in the Sierra SM Users Manual. For a model using only solid elements block skinning relieves the user from having to set an extensive number of sidesets in the input mesh.

Note the commands that create contact surfaces from blocks only work on solid and shell elements (hexes, tetrahedra, wedges, quads, triangles, etc.). If "ignore_shells" is set to true in the Dash options block, then all-to-all contact will ignore shells. Sidesets can be used to define contact surfaces on either solid or shell blocks. Contact nodesets can be defined on any element type, solid, shell or beam. When defining interactions a contact nodeset must be paired with a sideset or block skin in order to find node/face constraints. Two contact nodesets cannot directly interact.

In rare situations poorly posed cyclic/self contact constraints are created that are problematic for the linear solver. First the model is divided into disjoint components (without shared nodes). Pairs of disjoint components then tie to each other with one sided node-face interactions. Multipoint constraints tie nodes on one side of the interface to faces on the other side of the interface. Parts sharing nodes may have no unique node-to-face pairing. A conformally meshed part that contacts itself has no unique pairing.

If no unique pairing exists, nodes on side A of the interface generate MPCs with faces on side B and also nodes on side B generate MPCs with the faces of side A . The constraints may be over-determined. Some redundant constraints are removed. Determining unique constraints is an open problem.

The Interaction Weight Matrix shows the potential node-face interactions specified in the input deck. It is in the log file. A pair of surfaces must be in proximity to actually interact.

$$\left\{ \begin{array}{ll} \text{disjoint} & \text{if } IJ = JI = 0 \\ \text{dependnent node } I & \text{if } IJ = 1 \text{ and } JI = 0 \\ \text{symmetric} & \text{if } IJ = JI = H \\ \text{error} & \text{otherwise} \end{array} \right. \quad (8.1)$$

Ideally the weight is either 1 or 0. This indicates one-way node face pair where nodes on one side of the interface interact with faces on the other side. Self contact is denoted by H for half. For self contact nodes on side A of the interface interact with faces on side B and also nodes on side B interact with faces one side A . Such self contact constraints are often redundant and cyclic. The attempt to remove and make the contact constraints uniquely determined creates messages and warnings about removed and redundant constraints. Even after the redundant constraint removal step the self contact constraints can cause solver robustness and accuracy issues. Thus, self contact should be avoided.

The disconnected component finder is available for setting up contact surfaces. Say a mesh contained a flashlight with four batteries, and the four batteries were all in the same element block in the mesh. The disconnected component finder would split this battery block into four separate contact surfaces. The disconnected component finder is useful for setting up interactions in a way that avoids self contact. See the **Sierra/SM** User Manual⁶⁹ for more details on use of the disconnected component finder.

```
Begin Dash Options
  separate disconnected mesh components = TRUE|FALSE(FALSE)
End Dash Options
```

Surfaces may be defined to be in or out of contact solely based on proximity. The **cutoff variable** Dash option enables users to further filter contact based on any nodal input variable. Fine-grained contact information determined by **Sierra/SM** may be passed to **Sierra/SD** to refine its contact constraints. Including the following lines in a contact definition will ignore contact where the **Sierra/SM** `celement` field is below 0.55 on the nodes of the node-face constraints.

```
Begin Dash Options
  cutoff variable celement < 0.55
End Dash Options
```

Multiple such lines may be defined in a single Dash Options block. This would enable the filtering of constraints based on the union of multiple nodal variables, or only retaining contact where a variable is inside a range (a, b) (i.e. `cutoff variable var_name < a` and `cutoff variable var_name > b`).



`cutoff variable` is currently BETA release.
Enable with the “`--beta`” command-line option.

8.2.2. Setting up Contact Interactions

Once the contact surfaces are defined, the next step is to set up the interactions between those contact surfaces. The **interaction defaults** block can be used to define both which surfaces will interaction with each other and how those surfaces interact. One and only one interaction defaults block may be present in a contact definition. The **interaction** block can enforce contact between specific surface pairs and set the interaction parameters for that pair (overriding the interaction defaults behavior for the surface pair.) Any number of **interaction** command blocks may be present in the contact definition.

Contact is used to tie structures together that are in adjacent. The normal and tangential tolerance define how far away from a face a node can be and still find contact. By default, a reasonable normal and tangential tolerance is automatically computed in the Dash contact library based on a small fraction of the characteristic element size. Thickness of shell elements are also considered in the default tolerance. For shell blocks the default

search tolerance is set to at least sixty percent of the maximum element thickness of the elements in the block.

```
BEGIN CONTACT DEFINITION <name>
  BEGIN INTERACTION DEFAULTS
    GENERAL CONTACT = OFF|ON(OFF)
    SELF CONTACT = OFF|ON(OFF)
    NORMAL TOLERANCE = <real>
    TANGENTIAL TOLERANCE = <real>
    CONSTRAINT FORMULATION = NODE_FACE|FACE_FACE(NODE_FACE)
    FRICTION MODEL = TIED|FRICTIONLESS(TIED)
  END INTERACTION DEFAULTS

  BEGIN INTERACTION
    SIDE A = <string_list>
    SIDE B = <string_list>
    SURFACES = <string_list>
    NORMAL TOLERANCE = <real>
    TANGENTIAL TOLERANCE = <real>
    CONSTRAINT FORMULATION = NODE_FACE|FACE_FACE
    FRICTION MODEL = TIED|FRICTIONLESS
    INTERACTION BEHAVIOR = NO_INTERACTION
  END INTERACTION
END CONTACT DEFINITION
```

Command options:

- **general contact** : On means that every surface will contact every other surface. By default, the Dash contact library will pick which surface is used for nodes and which for faces in each surface-to-surface pairing automatically. Generally finer meshed surface will provide the nodes and the coarser the faces. Additional considerations may also be taken into account for selection of node and face surfaces such as avoiding cyclic constraints.
- **self contact** : Self contact on indicates a surface may contact itself, this could occur if a structure folds over on itself. Self contact should generally be avoided in SD as it leads to over-constraint problems.
- **normal tolerance** and **Tangential Tolerance** : The default face-sized based search distance can be overridden by manually specifying the normal and tangential tolerances. These tolerances have units of length.
- **constraint formulation** : The constraint formulation line defines the constraint type to be used. Only the node_face option should be used with **Sierra/SD**. The face_face option is experimental.

- **friction model** : The friction model line selects the type of contact constraint, either **tied** or **frictionless**. Tied contact ties all translational DOFs together at the interface preventing any normal or tangential motion. For structural problems each of the three translational degrees are constrained together. Rotational DOFs are never tied. For acoustic-acoustic contact the acoustic degree of freedom is tied together. For structure-acoustic contact the normal-motion of the structure is tied to the acoustic degree of freedom. The frictionless keyword selects sliding contact that is tied in the normal direction only. For frictionless structural contact only the surface-normal motion of the contacting surfaces are constrained together, the surfaces are free to slip in the tangential directions. The normal direction for the frictionless constraint is taken from the normal of the contacting face. For acoustic-acoustic or structural-acoustic contact frictionless contact is equivalent to tied.
- **SIDE A SIDE B** : For node/face contact the nodes are defined by the *B* surface and the faces by the *A* surface. If multiple surfaces are given for side A or side B, then the faces of each side A surface will be constrained to then nodes of each side B surfaces.
- **surfaces** defines a set of surfaces in contact. The Dash library picks the node face pairing automatically based on relative mesh density and other considerations such as avoiding cyclic constraints. If more than two surfaces are given a contact interaction will be formed between each surface in the list and each other surface in the list. Using *SIDEA* or *SIDEB* and *SURFACES* commands in the same interaction section will cause an error.
- **interaction behavior** The special interaction behavior command allows turning off contact between specific surface pairs. The `no_interaction` option would generally be paired with interaction defaults general contact on and used to turn off specific pairings where contact should not occur, such as slide lines.

8.2.3. Gap removal and element quality

As with tied surfaces [8.1](#), [3.38](#) the Contact Definition supports gap removal.

`GAP REMOVAL = OFF|ON(ON)`

The rewards and risks of gap removal are described in [Section 3.38](#).

Gap removal will trigger output of an element quality table listing the worst element in each block before and after gap removal. The element quality metric used in this table is not consistent with Sierra Solid Mechanics. Element condition number is used instead. This maintains consistency with other features within Sierra Structural Dynamics, and with other Sierra applications such as Cubit.

The table will begin with the header shown below:

===== ELEMENT CONDITION QUALITY INFORMATION =====

A value of 1.0 is an ideal element. As the value approaches INF the quality is decreasing. Values less than or equal to zero indicate the element cannot compute condition number. For example a beam, conmass, rbar, etc. A value of N/A also means no shape metric exists for that element topology. Element quality can be plotted on the mesh with the ElementQuality variable, triggered by the gap_removal solution case.

-----+-----+-----+-----+-----

The table contents below the header are too wide to fit in this document.

8.2.4. Examples

```
// Most basic contact definition to tie everything that touches
begin contact definition
  skin all blocks on
  begin interaction defaults
    general contact = on
  end interaction defaults
end contact definition
```

```
// Tie a few specific surfaces, like in 'tied data' block
begin contact definition
  contact surface s2 contains surface_2
  contact surface s3 contains surface_3
  contact surface s4 contains surface_4
  begin interaction
    side A = s2
    side B = s3 s4
    normal tolerance 0.25
  end interaction
end contact definition
```

```
// All-to-all contact with some custom tolerances
begin contact definition
  skin all blocks = on
  begin interaction defaults
    general contact = on
    normal tolerance 1e-3
    tangential tolerance 1e-6
  end interaction defaults
end contact definition
```

```
// Tie the nodes of the two bolt flanges to a fixture block
begin contact definition
  contact nodeset boltFlange1 contains nodeset 901
  contact nodeset boltFlange2 contains nodeset 902
  contact surface fixture contains blocks_1 block_3 block_4
  begin interaction
    side A = fixture
    side B = boltFlange1 boltFlange2
  end interaction
end contact definition
```

```
// All-to-all contact, but turning off tying of a specific
// surface pair and using sliding contact for different
// surface pair.
begin contact definition
  skin all blocks = on
  begin interaction defaults
    general contact = on
  end
  begin interaction
    surfaces = piston piston_housing
    friction model = frictionless
  end-interaction
  begin interaction
    surfaces = drive_shaft drive_bearing
    interaction behavior = no_interaction
  end interaction
end contact definition
```

8.2.5. Notes and Usage Guidelines

- Multiple 'begin contact definition' blocks may be included in a single analysis. However, as with **Sierra/SM** using multiple contact definition blocks is generally discouraged. If multiple contact definitions are used, then the surfaces used in the contact definitions can not overlap. If the same surfaces are used in multiple contact definitions duplicate and/or incompatible constraints may be found between the contact definition blocks causing solver difficulty.
- Contact constraints are ultimately enforced by translational MPCs. The MPCs generated by contact tie only translational degrees of freedom.
- Self contact (contact constraints generated from a surface folding over on itself) are often over-determined and cannot be enforced accurately. Though the contact

definition can find such self contact constraints, the use of self contact should be avoided.

- If contact constraints are defined at a significant gap those constraints will artificially impede the rotation of the model. Gap removal can help with this issue, but it is recommended to only use contact to tie objects that are in close physical proximity.
- Use the 'outputs' option [constraint_info] to visualize more information about the generated contact constraints. The nodes of a node/face contact constraint will have a '1' for [node_face_mpc_count] output. The [mpc_status] nodal field will be painted with '1' for any node involved in a contact constraint on either the node or face side. If any nodes have the [node_face_mpc_both_node_and_face] flag this could indicate an issue of over-constraint.

8.2.6. Differences Between SM and SD Defaults

- By default, SD uses node/face constraints. SM uses face/face constraints by default. Face/face constraints may be optionally used within SD. However, face/face constraints are considered experimental for SD at this time and not recommended.
- By default, SD will try to remove the gap from contact constraints. SM does not have a gap removal option at this time.
- In SD the friction model defaults to 'TIED'. In SM the friction model defaults to 'FRICTIONLESS'.
- The SD friction model is a linearization of the SM frictionless model. A key difference is that in SD the frictionless constraint is free to slide on the face but can have no motion in the normal direction of the face. In SM the sliding along the face is also unconstrained, the node is prevented from penetrating the face, but differing from SD the frictionless constraint can open gap and separate the surfaces. This different behavior for the positive and negative normal directions is a fundamentally nonlinear behavior not applicable to linear structural dynamic analysis.

8.3. Lofted Surfaces and Gap Removal

Lofted surfaces are important because analysts often build meshes with an initial gap between the surfaces. If standard methods are used to tie the surfaces, but the separation (or *lofting*) is not taken into account, then the constraints are no longer consistent with rigid body motion. Generally this means that the rotational rigid body motion introduces strain into the system.

There is a gap removal [3.38](#) solution.

8.3.1. Example

This example illustrated in Figure 8-75 uses the coordinates listed in Table 8-162. In this figure, the face (represented by nodes 1-4) constrains each of the 3 nodes (nodes 5, 6, and 7.) Node “6” is on the face. It will be clear that standard methods apply the proper constraints. However, nodes “5” and “7” are offset from the face. As a consequence, constraint equations written for a node constrained to the face introduce errors when applied to the lofted node.




Figure 8-75. – Lofted Constraint Example.

Node	Coordinates		
1	0	4	0
2	0	2	0
3	2	0	0
4	4	0	0
5	1	3	0
6	1	1	0
7	3	1	0

Table 8-162. – Coordinates of Face (red) and Nodes (blue).

8.3.2. Projection Approach

The constraint equations from a conventional approach (meaning that the constraints are written by projecting the node location to the plane of the face, but not adjusting for the lofting) are shown in Table 8-163. These equations are not orthogonal to rigid body modes, and as a consequence, there are only two zero energy modes for this system rather than the 3 we anticipate. ¹

¹All motion out of the xy plane has been eliminated.

$$\begin{aligned}
u_x(1) + u_x(2) - 2u_x(5) &= 0 \\
u_y(1) + u_y(2) - 2u_y(5) &= 0 \\
u_x(2) + u_x(3) - 2u_x(6) &= 0 \\
u_y(2) + u_y(3) - 2u_y(6) &= 0 \\
u_x(3) + u_x(4) - 2u_x(7) &= 0 \\
u_y(3) + u_y(4) - 2u_y(7) &= 0
\end{aligned}$$

Table 8-163. – Conventional Constraint Equations.

These constraints are represented by the matrix C , where the x and y dofs are grouped together.

$$C = \begin{bmatrix} 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -2 \end{bmatrix}$$

The three rigid body vectors (in this 2D frame) are,

$$R = \begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ -3 & -1 & -1 & -1 & 1 & 1 & 1 & 3 & -1 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

where the first two vectors represent translations, and the last is a rigid body rotation about point “6”. The product of $C * R'$ can be computed.

$$CR^T = \begin{bmatrix} 0 & 0 & -2 \\ 0 & 0 & -2 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 2 \\ 0 & 0 & 2 \end{bmatrix}$$

Each row of this matrix corresponds to a constraint equation from Table 8-163. Each column is associated with one of the three rigid body vectors. The translational rigid body vectors are orthogonal to the constraint matrix; the products are zero and no strain is induced. However, rotation about node 6 induces strain. The constraints are not invariant to rotation.

In a transient dynamic analysis with modest rotations and small gaps the effects of these constraint errors are often imperceptible. However, for large rotations or large gaps they may become apparent. They are always observable in modal analysis where they manifest as nonzero rigid body modes.

Mitigation Strategies As demonstrated in the previous section, constraint errors can introduce resistance to either translational or rotational rigid body motion. There are several strategies to mitigate these issues.

1. The analyst building the model ensures there are no projection errors.
2. Correct the initial geometry (using gap removal) so there are no projection errors.
3. Modify the constraints through algebraic means to ensure that they are orthogonal to rigid body motion.
4. Use the constraints appropriate to the lofted geometry. There is no software to do this currently.

If these methods cannot be applied, then the analysis must absorb the errors due to the constraint errors introduced by projection.

8.4. *Moving MPCs*

Sierra/SD supports moving contact through the use of moving MPCs. This can be enabled using the solution parameter `nUpdateConstraints = 1` as shown in 2.3. There are several other recommended parameters to be used with moving contact. These will be described in this section.

The example input blocks below show how the moving MPCs may be enabled. The `predictorCorrector` is set to 0 within the `transient` block. In the `Parameters` block, the solver should be set to update constraints without updating the matrices by setting `solverReset = constraints`. There is an inexpensive preconditioner for acoustic problems with moving constraints. To select this preconditioner, set the `preconditioner_type`, `krylov_method`, `orthog`, and `max_numterm_C1` as shown in the example input file below. The final two parameters in the `GDSW` block enable the block diagonal preconditioner with block size defined by the element condition number. To make a more powerful, but also more expensive, preconditioner increase the `max_element_condition`.

```
solution
  solver GDSW
  transient
    time_step 1.e-4
    nsteps 100
    nUpdateConstraints = 1
    predictorCorrector = 0
  end

  parameters
    solverReset = constraints
  end
```

```
GDSW
  preconditioner_type DIAG
  krylov_method PCG
  orthog 0
  max_numterm_C1 0
  identify_low_quality_elements true
  max_element_condition 5
end
```


9. EXAMPLE INPUT FILES

Example input files are shown for several types of analyses. The input file is case-insensitive except for special cases such as file names,

9.1. *Eigenvalue problem*

The following input file will output the first four mode shapes to an **Exodus** output file name *hexplate-out.exo*. A results file, *hexplate.rslt*, will not be created since no results have been selected for output in the **echo** section.

```
Solution
  eigen
  nmodes 4
  title 'Mode Shapes of Lowest Frequency Modes'
end

file // finite element mesh
  geometry_file hexplate.exo
end

Boundary
  nodeset 77
  fixed
end

outputs
  deform
end

block 44 // The default is the Hex8b
  material 3
  hex8
end

Material 3
  name "steel"
  E 30e6 +/- 1 %
  nu .3
  density 0.288
end

sensitivity
```

```

    values all
end

```

9.2. *Anisotropic Material*

The following input file is an example of a hexahedron mesh with anisotropic properties.

```

Solution
    eigen
    title 'Anisotropic Format'
end

file
    geometry_file      anisogump.exo
end

Boundary
    nodeset 4 y = 0
    nodeset 5 x = 0
    nodeset 6 z = 0
end

load
    // sum of forces on surface should be equal to area
    // imposed forces are additive
    nodeset 1 force = 0.0  0.083333  0.0
    nodeset 2 force = 0.0 -0.041666  0.0
    nodeset 3 force = 0.0 -0.020833  0.0
end
outputs
    deform
end

block 1
    hex8
    material 1
end

Material 1
    name "anisotropic gump"
    anisotropic
    Cij
    1.346      0.5769      0.5769  0      0      0
              1.346      0.5769  0      0      0

```

```

                                1.346   0       0       0
                                0.3846  0       0
                                0.3846  0.3846  0
                                0.3846
density 1
end

```

9.3. *Multiple materials*

The next example shows the input for an **Exodus** model with many element blocks and materials. Keyword **lumped** in the **Solution** section selects a lumped (nearly diagonal) mass matrix.

```

Solution
    eigen
    nmodes 1
    title 'Dozen blocks and six materials'
    lumped
end
file
    geometry_file      multi.exo
end
Boundary
    nodeset 1
    fixed
    nodeset 3
    x = 0
    y = 0
    z = 0
    RotY = 0
    RotZ = 0
end
outputs
    deform
end
    // A block is required for each element block
block 1 //in the input Exodus (Genesis) mesh database.
    material 2
    Beam2
end
block 101
    integration full
    wedge6
    material 1

```

```

end
block 2
    material 2
end
block 102
    integration full
    wedge6
    material 2
end
block 3
    material 3
end
block 103
    integration full
    wedge6
    material 3
end
block 4
    material 4
end
block 104
    integration full
    wedge6
    material 4
end
block 5 // Tip. Not capitalizing "material" here
    material 5 // helps to distinguish it
end // from a Material section.
block 105
    wedge6
    integration full
    material 5
end
block 6
    material 6
end
block 106
    wedge6
    integration full
    material 6
end // Each material referenced in a necessary block
Material 1 // must be defined here. Extra materials are ignored.
    name "Phenolic"
    E 10.5E5
    nu .3

```

```

        density 129.5e-6
end
Material 2
    name 'Aluminum'
    E 10.0E6
    nu 0.33
    density 253.82e-6
end
Material 3
    name 'foam'
    E 100.
    nu 0.3
    density 18.13e-6
end
Material 4
    name 'HE'
    E 5E5
    nu 0.45
    density 129.5e-6
end
Material 5 // Tip. Capitalizing material helps to
            // distinguish it from a material in a block.
    name 'Uranium'
    E 30e6
    nu 0.3
    density 1768.97e-6
end
Material 6
    name 'wood'
    E 200.e3
    nu .3
    density 77.7e-6
end

```

9.4. **Modaltransient**

The next example shows the input for a **modaltransient** analysis. Accelerations are output to an **Exodus** file *bar-out.exo*. This example has damping, polynomial and linear functions. Also, sensitivities are calculated.

```

Solution
    modaltransient
        nmodes 10
        time_step .000005
        nsteps 100

```

```

        nskip 1
        title 'Test modal transient on prismatic bar'
end

file
    geometry_file bar.exo
end

outputs
    acceleration
end

Boundary
    nodeset 1
        fixed
    end

damping
    gamma 0.001
end

block 1
    material 1
end

Material 1
    name "aluminum"
    E 10e6
    nu .33
    density 2.59e-4
end

load
    nodeset 3
        force = 1. 1. 1.
        function = 3
    end

function 1
    type linear
    name "test_func1"
    data 0.0 0.0
    data 0.0150 0.0
    data 0.0152 1.0
    data 0.030 0.0

```

```

end

function 3
  type linear
  name "white noise"
  data 0.0 1.0
  data 0.0001 1.0
  data 0.0001 0.0
  data 1.0 0.0
end

```

9.5. **ModalFrf**

In this **ModalFrf** analysis, accelerations are output to an **Exodus** file *bar-out.frq*.

```

Solution
  ModalFrf
    nmodes 10
    title 'Test ModalFrf on prismatic bar'
end

file
  geometry_file bar.exo
end

frequency
  freq_min 0
  freq_step=10
  freq_max=3000
  nodeset 3
  disp
end

outputs
  acceleration
end

Boundary
  nodeset 1
  fixed
end

damping

```

```

    gamma 0.001
end

block 1
    material 1
end

Material 1
    name "aluminum"
    E 10e6
    nu .33
    density 2.59e-4
end

load
    nodeset 3
        force = 1. 1. 1.
        function = 3
end

function 2
// a smooth pulse of duration .05 sec
// peaking near t=.02 sec at 0.945
    type polynomial
    name "poly_fun"
    data 0. 0.
    data 2.0 -8.0e2
    data 0.5 8.9443
end

function 3
    type linear
    name "white noise"
    data 0.0 1.0
    data 10000. 1.0
end

```

9.6. *Directfrf*

A **directfrf** is run with displacements output to the **Exodus** file *bar-out.frq*.

```

Solution
    directfrf

```



```

end

frequency
  freq_min = 1000.0
  freq_step = 7000
  freq_max = 5.0e4
  disp
  block 1
end

file
  geometry_file bar.exo
end

outputs
disp
end

Boundary
  nodeset 1
    fixed
end

block 1
  material 1
end

Material 1
  name "aluminum"
  G 0.8E+9
  K 4.8E+9
  density 2.59e-4
end

load
  sideset 1
  pressure = -1.0
  function=3
end

function 3
  type linear
  name "white noise"
  data 0.0 1.0
  data 10000. 1.0

```

end

9.7. Statics

The following example is a **statics** analysis which will output stresses to the **Exodus** output file *quadt-out.exo*.

```
Solution
    statics
    title '10x1 beam of quadt'
end
file
    geometry_file      quadt.exo
end
Boundary
    nodeset 1
    fixed
end
loads
    nodeset 2
    force = 1000.0 1000.0 0.0
end
outputs
    stress
end
block 1
    material 1
    QuadT
end
Material 1
    name "steel"
    E 30.0e6
    nu 0.25e0
    density 0.7324e-3
end
```

BIBLIOGRAPHY

- [1] J. L. Aklonis and W. L. MacKnight. “Introduction to Polymer Viscoelasticity”. In: Wiley, 1983. Chap. 1-6, pp. 1–316 (cit. on p. 255).
- [2] D. J. Allman. “A Compatible Triangular Element Including Vertex Rotations for Plane Elasticity Problems”. In: *Comput. and Struct.* 19.1-2 (1996), pp. 1–8 (cit. on pp. 289, 451).
- [3] Kenneth F. Alvin. “Implementation of Modal Damping in a Direct Implicit Transient Algorithm”. In: *Presented at the 42nd AIAA/ASME/ASCE/AHS/ASC SDM*. Apr. 2001, p. 1589 (cit. on p. 277).
- [4] Kenneth F. Alvin et al. “Incorporation of Sensitivity Analysis into a Scalable Massively Parallel Structural Dynamics FEM code”. In: *Presented at the 5th U.S. Congress on Computational Mechanics*. Boulder, CO, Aug. 1999 (cit. on p. 75).
- [5] ATA Engineering. *Attune User’s Guide*. URL: http://www.ata-e.com/software/attune/Attune%5C_Users%5C_Guide%5C_v2/index.html (cit. on p. 72).
- [6] Philip Avery, Charbel Farhat, and Garth Reese. “Fast Frequency Sweep Computations Using a Multi-point Pade-Based Reconstruction Method and an Efficient Krylov Solver”. In: *Int. J. Numer. Meth. Engng.* 69.13 (Sept. 2006), pp. 2848–2875 (cit. on p. 137).
- [7] C. G. Baker et al. “Anasazi Software for the Numerical Solution of Large-Scale Eigenvalue Problems”. In: *ACM Trans. on Math. Software* 36.3 (2009), pp. 1–23 (cit. on p. 151).
- [8] Jean-Louis Batoz, Klaus-Jurgen Bathe, and Lee-Wing Ho. “A Study of Three-Node Triangular Plate Bending Elements”. In: *Int. J. Numer. Meth. Engng.* 15 (1980), pp. 1771–1812 (cit. on pp. 289, 451).
- [9] T. Belytschko, CS Tsay, and WK Liu. “A stabilization matrix for the bilinear Mindlin plate element”. In: *Computer Meth. in Appl. Mech. Eng.* 29.3 (1981), pp. 313–327 (cit. on p. 288).
- [10] Martin P Bendsøe. “Optimal shape design as a material distribution problem”. In: *Structural optimization* 1.4 (1989), pp. 193–202 (cit. on p. 243).
- [11] J-P Berenger. “Perfectly Matched Layer for the FDTD solution of wave-structure interaction problems”. In: *Antennas and Propagation, IEEE Transactions on* 44.1 (1996), pp. 110–117 (cit. on p. 370).
- [12] A Bermúdez, L Hervella-Nieto, A Prieto, et al. “An optimal Perfectly Matched Layer with unbounded absorbing function for time-harmonic acoustic scattering problems”. In: *Journal of Computational Physics* 223.2 (2007), pp. 469–488 (cit. on pp. 370, 371).

- [13] D. P. Bertsekas. “Projected Newton Methods for Optimization Problems with Simple Constraints”. In: *SIAM J. Control and Optimization* 20 (1982), pp. 221–246 (cit. on p. 229).
- [14] Matthew R. W. Brake. *A Reduced Iwan Model that Includes Pinning for Bolted Joint Mechanics*. Tech. rep. SAND 2016-0207C. Sandia National Laboratories, 2016 (cit. on p. 324).
- [15] Gregory Bunting. *Strong and Weak Scaling of the Sierra/SD Eigenvector Problem to a Billion Degrees of Freedom*. Tech. rep. SAND 2019-1217. Sandia National Laboratories, 2019 (cit. on p. 15).
- [16] Gregory Bunting et al. “Novel Strategies for Modal-Based Structural Material Identification”. In: *Journal of Mechanical Systems and Signal Processing* 149.107295 (2021) (cit. on p. 148).
- [17] Gregory Bunting et al. “Parallel Ellipsoidal Perfectly Matched Layers for Acoustic Helmholtz Problems on Exterior Domains”. In: *Journal of Computational Acoustics* (2018) (cit. on p. 369).
- [18] J. V. Burke, J. J. Moré, and G. Toraldo. “Convergence properties of trust region methods for linear and convex constraints”. In: *Math. Programming* 47 (1990), pp. 305–336 (cit. on p. 229).
- [19] P. H. Calamai and J. J. Moré. “Projected gradient methods for linearly constrained problems”. In: *Math. Programming* 39 (1987), pp. 93–116 (cit. on p. 229).
- [20] M. Christon. “The influence of the mass matrix on the dispersive nature of the semi-discrete, second-order wave equation”. In: *Computer Meth. in Appl. Mech. Eng.* 173.1 (1999), pp. 147–166 (cit. on p. 52).
- [21] J. Chung and G. M. Hulbert. “A Time Integration Algorithm for Structural Dynamics with Improved Numerical Dissipation - The Generalized Alpha Method”. In: *JAM* 60.2 (1993), pp. 371–375 (cit. on p. 208).
- [22] R. D. Cook and M. E. Plesha D. S. Malkus. *Concepts and Applications of Finite Element Analysis*. 3rd. John Wiley & Sons, 1989 (cit. on pp. 208, 309).
- [23] J. E. Dennis, Jr. and R. B. Schnabel. *Numerical Methods for Nonlinear Equations and Unconstrained Optimization*. Philadelphia: SIAM, 1996 (cit. on pp. 227, 228).
- [24] J. M. Dickens, J. M. Nagawa, and M. J. Wittbrodt. “A critique of mode acceleration and modal truncation augmentation methods for modal response analysis”. In: *Comput. and Struct.* 62.6 (1997), pp. 985–998 (cit. on p. 215).
- [25] Clark R. Dohrmann and Olof B. Widlund. “Hybrid domain decomposition algorithms for compressible and almost incompressible elasticity”. In: *Int. J. Numer. Meth. Engng.* 82 (2010), pp. 157–183 (cit. on p. 50).
- [26] U. S. FAA. *MIL-HDBK-5J. Metallic Materials and Elements for Aerospace Vehicle Structures*. Tech. rep. Department of Defence, 2003 (cit. on pp. 259–261).
- [27] Charbel Farhat, Crivelli, and M. Géradin. “Implicit time integration of a class of constrained hybrid formulations - Part I: Spectral stability theory”. In: *Int. J. Numer. Meth. Engng.* 41 (1998), pp. 675–696 (cit. on p. 208).

- [28] C. A. Felippa. *The SS8 Solid-Shell Element: Formulation and a Mathematica Implementation*. Tech. rep. CU-CAS-02-03. Univ. Colo. at Boulder, 2002 (cit. on p. 297).
- [29] J. D. Ferry. “Viscoelastic Properties of Polymers”. In: Wiley, 1980. Chap. 1-19, pp. 1–590 (cit. on p. 255).
- [30] F. Fuentes et al. “Orientation embedded high order shape functions for the exact sequence elements of all shapes”. In: *Computers and Mathematics with Applications* 70.1 (2015), pp. 353–458 (cit. on p. 2).
- [31] Gregory D. Sjaardema. *A Collection of Exodus Utilities: Exodiff, Epu, Ejoin, and Conjoin*. Tech. rep. SAND2011-5715. PO Box 5800, Albuquerque, NM 87185-5800: Sandia National Laboratories, 2011 (cit. on pp. 17, 18).
- [32] M. F. Hamilton and D. T. Blackstock. *Nonlinear Acoustics*. Academic Press, 1998 (cit. on p. 254).
- [33] M. Heroux et al. “An overview of the Trilinos project”. In: *ACM Trans. on Math. Software* 31.3 (2005), pp. 397–423. ISSN: 0098-3500 (cit. on pp. 226, 248).
- [34] Thomas J. R. Hughes. *The Finite Element Method—Linear Static and Dynamic Finite Element Analysis*. Prentice-Hall, Inc, 1987 (cit. on pp. 194, 208).
- [35] G. M. Hulbert and T. J. R. Hughes. “An error analysis of truncated starting conditions in step-by-step time integration: Consequences for structural dynamics”. In: *Earthquake Engineering & Structural Dynamics* 15.7 (1987), pp. 901–910 (cit. on pp. 43, 209).
- [36] A. Ibrahimbegovic and E. L. Wilson. “A Modified Method of Incompatible Modes”. In: *Communications in Applied Numerical Methods* 7 (1991), pp. 187–194 (cit. on p. 281).
- [37] Conor D. Johnson, David A. Kienholz, and Lynn C. Rogers. “Finite element prediction of damping in beams with constrained viscoelastic layers”. In: *AIAA Journal* 20.9 (1982), pp. 1284–1290 (cit. on p. 122).
- [38] C. T. Kelley and E. W. Sachs. “A Trust Region Method for Parabolic Boundary Control Problems”. In: *SIAM J. Optimization* 9 (1999), pp. 1082–1099 (cit. on p. 229).
- [39] Kinsler et al. *Fundamentals of Acoustics*. John Wiley & Sons, 1982 (cit. on pp. 190, 286).
- [40] D. P. Kouri et al. *Rapid Optimization Library*. 2014. URL: <https://trilinos.org/packages/rol> (cit. on p. 226).
- [41] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK Users’ Guide*. Philadelphia, PA, USA: SIAM, 1998 (cit. on p. 151).
- [42] C.-J. Lin and J. J. Moré. “Newton’s method for large bound-constrained optimization problems”. In: *SIAM J. Optim.* 9.4 (1999), pp. 1100–1127 (cit. on p. 229).
- [43] R. H. MacNeal. *Finite Elements: Their Design and Performance*. Marcel Dekker, 1994 (cit. on p. 281).

- [44] Ch Michler et al. “Improving the performance of Perfectly Matched Layers by means of hp-adaptivity”. In: *Numerical Methods for Partial Differential Equations* 23.4 (2007), pp. 832–858 (cit. on p. 370).
- [45] J. Nocedal and S. Wright. *Numerical Optimization*. Springer, 2006 (cit. on pp. 227, 228).
- [46] O. O. Ochoa and J. N. Reddy. *Finite Element Analysis of Composite Laminates*. Kluwer Academic Publishers, 1992 (cit. on p. 288).
- [47] J. T. Oden. *Mechanics of Elastic Structures*. 1st ed. McGraw-Hill, Inc., 1967 (cit. on p. 307).
- [48] Michael A. Puso. “A 3D mortar method for solid mechanics”. In: *Int. J. Numer. Meth. Engng.* 59 (2004), pp. 315–336 (cit. on p. 43).
- [49] J. N. Reddy. *An Introduction to the Finite Element Method*. 2nd ed. McGraw Hill, 1993 (cit. on p. 288).
- [50] Garth Reese, Rich Field, and Daniel J. Segalman. “A Tutorial on Design Analysis Using von Mises Stress in Random Vibration Environments”. In: *Shock and Vibration. Digest* 32.6 (2000) (cit. on p. 174).
- [51] Brett A. Robertson et al. *Dispersion Analysis of Acoustic Elements in Sierra/SD*. Tech. rep. SAND2014-3870P. PO Box 5800, Albuquerque, NM 87185-5800: Sandia National Laboratories, 2006 (cit. on p. 52).
- [52] S D Team. *Sierra Structural Dynamics - Theory Manual*. Tech. rep. SAND2021-11328. PO Box 5800, Albuquerque, NM 87185-5800: Sandia National Laboratory, 2021 (cit. on pp. 53, 71, 188, 224, 269, 278, 297, 301, 395, 397, 458).
- [53] S D Team. *Sierra Structural Dynamics How To*. Tech. rep. SAND2021-11328. PO Box 5800, Albuquerque, NM 87185-5800: Sandia National Laboratories, 2020 (cit. on pp. 64, 67, 127, 156, 196, 343, 344, 349, 403, 448).
- [54] S D Team. *Sierra Structural Dynamics Verification*. Tech. rep. SAND2021-11330. Sandia National Laboratories, 2021 (cit. on pp. 182, 190, 195, 214, 297, 301, 448).
- [55] Larry A. Schoof and Victor R. Yarberry. *EXODUS II: A Finite Element Data Model*. Tech. rep. SAND92-2137. PO Box 5800, Albuquerque, NM 87185-5800: Sandia National Laboratories, 1994 (cit. on p. 28).
- [56] Daniel J. Segalman. *A Four-Parameter Iwan Model for Lap-Type Joints*. Tech. rep. SAND 2002-3828. Sandia National Laboratories, Nov. 2002 (cit. on p. 319).
- [57] Daniel J. Segalman. *An Initial Overview of Iwan Modeling for Mechanical Joints*. Tech. rep. SAND2001-0811. PO Box 5800, Albuquerque, NM 87185-5800: Sandia National Laboratories, 2001 (cit. on p. 321).
- [58] Daniel J. Segalman and Michael J. Starr. *Relationships Among Certain Joint Constitutive Models*. Tech. rep. SAND2004-4321. PO Box 5800, Albuquerque, NM 87185-5800: Sandia National Laboratories, 2004 (cit. on p. 321).

- [59] Daniel J. Segalman et al. “An Efficient Method for Calculating RMS Von Mises Stress in a Random Vibration Environment”. In: *Journal of Sound and Vibration* 230.2 (2000), pp. 393–410 (cit. on pp. 174, 424).
- [60] Daniel J. Segalman et al. “Estimating the Probability Distribution of von Mises Stress for Structures Undergoing Random Excitation”. In: *Transactions of the ASME* 122 (Jan. 2000) (cit. on p. 175).
- [61] Daniel J. Segalman et al. *Handbook on Dynamics of Jointed Structures*. Tech. rep. SAND2009-4164. PO Box 5800, Albuquerque, NM 87185-5800: Sandia National Laboratories, 2009 (cit. on p. 321).
- [62] Gregory D. Sjaardema. *APREPRO: An Algebraic Preprocessor for Parameterizing Finite Element Analyses*. Tech. rep. SAND92-2291. Sandia National Laboratories, 1992 (cit. on p. 28).
- [63] Gregory D. Sjaardema. *GROPE: A GENESIS/EXODUS Database Examination Program*. Tech. rep. SAND92-2289. Sandia National Laboratories, 1992 (cit. on p. 18).
- [64] Gregory D. Sjaardema. *Seacas Assemblies Wiki*. URL: <https://github.com/gsjardema/seacas/wiki/Assemblies> (cit. on pp. 7, 30).
- [65] C. J. Stimpson et al. *Verdict Library Reference Manual*. Tech. rep. SAND2007-2853P. Sandia National Laboratories, 2007 (cit. on p. 416).
- [66] G. Strang and G. Fix. *An Analysis of the Finite Element Method*. Wellesley-Cambridge Press, 2008. ISBN: 9780980232707. URL: <https://books.google.com/books?id=K5MA0wAACAAJ> (cit. on p. 23).
- [67] R. L. Taylor, P. J. Beresford, and E. L. Wilson. “A Non-conforming Element for Stress Analysis”. In: *Int. J. Numer. Meth. Engng.* 10 (1976), pp. 1211–1219 (cit. on p. 281).
- [68] S D Team. *SD Design Manual*. Tech. rep. SAND2021-4312. Sandia National Laboratories, 2021 (cit. on p. 448).
- [69] S M Team. *Sierra Solid Mechanics 4.56 User’s Guide*. Tech. rep. SAND2020-5362. PO Box 5800, Albuquerque, NM 87185-5800: Sandia National Laboratories, 2020 (cit. on pp. 11, 253, 464, 467).
- [70] Sierra Toolkit Development Team. *Sierra Toolkit Manual Version 5.1.4*. Tech. rep. SAND2018-2856. PO Box 5800, Albuquerque, NM 87185-5800: Sandia National Laboratories, 2018 (cit. on pp. 17, 93).
- [71] D. Thompson, P. P. Pébay, and J. N. Jortner. *An Exodus II Specification for Handling Gauss Points*. Tech. rep. SAND2007-7169. Sandia National Laboratories, 2007 (cit. on pp. 213, 214, 422).
- [72] T. F. Walsh et al. “Finite element methods for structural acoustics on mismatched meshes”. In: *Journal of Computational Acoustics* 17.3 (2009), pp. 247–275 (cit. on p. 458).

- [73] Malcolm L Williams, Robert F Landel, and John D Ferry. “The temperature dependence of relaxation mechanisms in amorphous polymers and other glass-forming liquids”. In: *Journal of the American Chemical society* 77.14 (1955), pp. 3701–3707 (cit. on p. 255).
- [74] Paul H. Wirsching and Mark C. Light. “Fatigue under wide band random stresses”. In: *Journal of the Structural Division, ASCE* 106.7 (1980), pp. 1593–1607 (cit. on p. 173).
- [75] Paul H. Wirsching, Thomas L. Paez, and Keith Ortiz. *Random Vibrations: Theory and Practice*. Courier Corporation, 2006 (cit. on pp. 173, 259).

INDEX

- acoustic
 - dispersion, 52
 - p0, 363
 - Pdot, 363
 - point source, 385
 - scale, 386
- acoustic__accel, 382, 385–387
- acoustic__vel, 382, 385–387
- AcousticFraction, 190
- acousticlighthill, 386
- ADiag, 438
- aeigen, 151
 - anblocksize, 152
 - anverbosity, 152
 - eig_tol, 151
 - shift, 151
- AllowExodusDistFacts, 382
- AllStructural, 43
- Anasazi, 151, 166, 184
- angular__acceleration, 400
- angular__velocity, 400
- anisotropic, 249–251
 - example, 478
- APartVel, 435
- APressure, 435
- ARPACK, 145
- average, 354
- Beam2, 301, 301–305, 307, 308, 416
- begin-periodic, 371
- bending_factor, 291
- block, 26, 140, 212, 238, 240, 254, 268, 268, 330, 338, 346, 381, 455
 - blkalpha, 272–275
 - blkbeta, 272–275
 - coordinate, 272
 - density_scale_factor, 272
 - material identification, 238
 - non-structural mass, 273
 - nonlinear, 272
 - NSM, 273
 - nsm, 272
 - parameters, 271
 - stiffness_scale_factor, 272
- block_energies, 425
- blockwise density scaling, 274
- boundary, 174, 358, 370, 371, 384
 - fixed, 359
 - infinite element, 367
 - displacement1, 369
 - FieldTime, 369
 - P, 359
 - RotX, 359
 - RotY, 359
 - RotZ, 359
 - slosh, 366
 - V, 362
 - X, 359
 - Y, 359
 - Z, 359
- buckling, 156
 - nmodes, 156
 - shift, 156
- case, 119
- CBModel, 124, 125, 126, 126, 128, 131
 - file, 127
 - finite_difference, 131
 - format, 127, 347
 - GlobalSolution, 127
 - inertia_matrix, 127
 - netcdf, 126
 - nodeset, 127
 - OTM, 129
 - OTME, 129
 - OutElemMap, 129

- OutMap, 129
- sensitivity_method, 131
- sideset, 127
- spoint_offset, 127
- CBR *see Craig-Bampton reduction*, 123
- ceigen, 166, 184, 188, 257, 417
- centrifugal, 399
- centripetal, 399
- checkout, 122, 122
- Cij, 251, 266
- CMS *see component mode synthesis*, 123
- command line
 - Aprepro, 14
 - beta, 14
- complex load, 398
- component mode synthesis, 123
- con_tolerance, 25
- condition_limit, 38
 - ElemQualChecks, 38
- ConMass, 310, 310, 311
- ConMassA, 311
- consistent loads, 392
- constrain_rbms, 403
- constraint formulation, 468
- constraint_info, 224, 435
- ConstraintCorrectionFrequency, 208
- ConstraintErrorDiagnostics, 208
- ConstraintMethod, 52
 - Lagrange, 52
- contact
 - normal, 459
- contact definition, 11, 223, 464
 - Dash, 464
 - element quality, 469
 - friction model, 469
 - gap removal, 469
 - gap_removal, 223
 - initial overlap removal, 223
 - interaction, 467
 - interaction defaults, 467
 - side, 469
 - surfaces, 469
- coordinate, 76, 268, 274, 291, 293, 297, 381, 400, 444
- copy, 40
- Coriolis, 399
- Craig-Bampton reduction, 123, 124, 346
 - correction=vectors, 125
 - inertia tensor, 125
 - mass inertia matrix, 125
 - null space correction, 124
 - RbmDof, 125
- cutoff variable, 467
- damper, 278, 314, 327
 - cubic, 327
 - viscous, 314
- damping, 174, 275
 - block, 272
 - CJdamp, 122, 265
 - CJetaFunction, 123, 265
 - Conor Johnson, 122
 - ratiofun, 277
- dashpot, 314, 327
- data_file, 236
- data_truth_table
 - inverse-problem, 232
- data_weight_table, 238
- database name, 413, 445, 447
- dd_solver_output_file, 60
- DDAM
 - add, 133
 - analysis_direction, 133
 - athwartship, 134
 - eigen, 133
 - fore_and_aft, 134
 - load, 133
 - modalfilter, 133
 - nmodes, 133
 - preddam, 133
 - vertical, 134
- ddamout, 438, 439
- dead, 348, 348
- defaultSpecificHeat, 41
- density, 258
 - blockwise scaling, 274
- density_scale_factor, 274
- diagnostics, 20

- ADiag, [438](#)
- beams, [437](#)
- cubit, [21](#)
- e_{pu}, [22](#)
- explore, [21](#)
- kdiag, [24](#), [437](#)
- stk_balance, [21](#)
- dielectric, [267](#)
- directfrf, [136](#), [136](#), [164](#), [243](#), [368](#), [484](#)
 - example, [484](#)
- directfrf-inverse, [138](#), [226](#)
- displacement, [369](#), [431](#)
 - prescribed, [361](#)
- DMIG, [343](#), [345](#), [347](#)
- echo, [32](#), [174](#), [179](#), [224](#), [295](#), [346](#), [424](#), [425](#), [437](#), [444](#), [453](#), [455](#), [477](#)
 - Elmat, [457](#)
 - GEnergies, [424](#)
 - mass=block, [455](#)
 - massblock, [455](#)
 - material, [457](#)
 - memusage, [457](#)
 - modal amplitude, [456](#)
 - MPC, [224](#), [455](#)
 - NLresiduals, [454](#)
 - subdomains, [32](#)
- Eforce, [429](#)
- eforce, [330](#), [429](#), [431](#), [453](#)
- eig_tol, [22](#), [142](#), [186](#)
- eigen, [36](#), [119](#), [142](#), [145](#), [151](#), [152](#), [157](#), [164](#), [173](#), [174](#), [179](#), [243](#), [429](#)
 - example, [477](#)
 - fluidloading, [145](#)
 - nmodes all, [142](#)
 - restart, [49](#)
 - untilfreq, [145](#)
- eigen-inverse, [147](#)
- eigenvalue problem
 - comparison, [183](#)
 - eigen_norm, [41](#)
 - element checks, [416](#)
 - Largest_Ev, [153](#)
 - normalization, [41](#)
 - quadratic, [188](#), [191](#)
 - structural acoustics, [188](#)
 - structural acoustics (modal basis), [191](#)
 - tolerance, [36](#)
- elastic-plastic *see eplas element*, [326](#)
- elastic_strain, [418](#)
- ElemEigChecks, [416](#)
- element
 - Beam2, [301](#)
 - ConMass, [310](#)
 - dashpot, [314](#)
 - dead, [348](#)
 - eigenvalue checks, [416](#)
 - eplas, [326](#)
 - force, [429](#)
 - Ftruss, [309](#)
 - Gap, [331](#)
 - Joint2G, [325](#)
 - Gap2D, [334](#)
 - GasDmp, [335](#)
 - Hex20, [282](#)
 - Hex8, [281](#)
 - HexShell, [297](#)
 - Hys, [317](#)
 - dmax, [317](#)
 - fmax, [317](#)
 - kmax, [317](#)
 - kmin, [317](#)
 - Hysteresis element, [317](#)
 - Joint2G, [319](#)
 - Iwan, [319](#)
 - Property, [319](#), [326](#), [327](#)
 - line_weld, [328](#), [330](#)
 - gap removal, [328](#)
 - line_weld_force_rst, [330](#)
 - line_weld_moment_rst, [330](#)
 - Nbeam, [305](#)
 - Nmount, [336](#)
 - Nquad, [287](#)
 - Ntria, [287](#)
 - offset shell, [295](#)
 - orientation, [433](#)
 - Quad8T, [284](#)
 - QuadM, [286](#)
 - QuadT, [284](#)

- Rigid
 - RBar, [339](#)
 - RBE2, [340](#)
 - RBE3, [340](#)
 - Rrod, [338](#)
 - RSpring, [313](#)
 - Spring, [312](#)
 - Spring3, [313](#)
 - SpringDashpot, [315](#)
 - Stress/Strain, [449](#)
 - Superelement, [342](#)
 - Tet10, [283](#)
 - Tet4, [283](#)
 - TiBeam, [308](#)
 - Tria3, [290](#)
 - Tria6, [284](#)
 - TriaShell, [289](#)
 - Truss, [309](#)
 - Truth Table, [449](#)
 - Wedge15, [283](#)
 - Wedge6, [282](#)
- ElemQualChecks
 - condition_limit, [416](#)
 - limitations, [23](#)
- energy_exo_var, [39](#), [390](#), [392](#)
- energy_load, [383](#), [392](#)
- energy_time_step, [39](#), [392](#)
- enforced acceleration
 - random vibration, [406](#)
- engineering units, [35](#)
- eorient, [297](#), [419](#), [433](#), [451](#)
- Euler force, [399](#)
- evaluate expression
 - in Function
 - examples of, [93](#)
 - rules and options for composing, [93](#)
- Exodus
 - assemblies, [30](#)
 - creation, [31](#)
 - entity types, [30](#)
 - naming
 - limitations, [31](#)
- Exodus Read Functions, [394](#)
- Farhat, Charbel, [1](#)
- fatigue, [154](#), [173](#)
 - material, [258](#)
 - S-N curve, [259](#)
- Fatigue Solution, [153](#)
- Felippa, Carlos, [1](#)
- fiber_orientation, [291](#)
- FilterRbm, [277](#)
- FilterRBMLoad, [54](#)
- FilterRbmLoad, [12](#), [40](#), [43](#), [218](#), [403](#)
- flush, [45](#)
- force, [221](#)
 - constraint force, [428](#)
 - reaction force, [428](#)
- frequency, [138](#), [168](#), [236](#), [424](#), [429](#), [431](#), [448](#)
- FRF, [429](#), [431](#)
- Ftruss, [309](#), [310](#)
- fuego, [40](#)
- function, [84](#), [84](#), [113](#), [137](#), [174](#), [220](#), [362](#), [377](#), [386](#), [405](#)
 - blended, [106](#)
 - exo_var, [375](#), [377](#), [378](#)
 - ExodusRead, [378](#)
 - linear, [85](#)
 - loglog, [89](#)
 - offset, [85](#)
 - Piecewise Linear, [87](#)
 - planar step wave, [101](#)
 - plane wave, [98](#)
 - K0, [98](#), [101](#)
 - tied data, [98](#)
 - plane_wave_freq, [100](#)
 - polynomial, [88](#)
 - random, [90](#), [375](#), [377](#)
 - interp, [91](#)
 - ReadNodal, [375](#), [388](#)
 - interp, [375](#)
 - ReadNodalSet, [376](#)
 - ReadSurface, [377](#), [388](#)
 - velX, [377](#)
 - velY, [377](#)
 - velZ, [377](#)
 - SamplingRandom, [89](#)
 - shift, [85](#)

- SpatialBC, 374
- spherical_wave, 102
- table, 88, 111
 - datafile, 112
 - dataline, 112
 - delta, 112, 399
 - dimension, 112
 - origin, 112, 399
 - size, 112
- tablename, 88
- undex loads, 103
- wet surface pressure, 105

Gap, 326

- ellipsoidal, 334

Gap element, 325, 331, 334

Gap2D, 334, 334

GasDmp, 335

GDSW, 53, 60, 403

- options, 50
- Parameters, 53
- prt_debug, 65, 67
- solver_options, 59
- solver_tol, 44
- SuperLUDist, 53, 59, 64

general contact, 468

Generalized Alpha, 208

- rho, 208

geometry_file, 29, 29, 202

global, 441

global variables, 432

globalHist, 445

grepos, 220

hand-off, 195, 200, 252, 373, 378, 464

Hex20, 282

Hex8, 281, 281, 282

Hex8b, 281

Hex8F, 282

Hex8u, 281

HexShell, 297, 298

- autolayer, 297
- Mass, 301

history, 179

hydrostatic balance, 219

I1, 302

I2, 302

iforce, 398

ignore_gap_inversion, 43

igravity, 398

imaginary_data_file, 236

ImagRelDispGxx, 431

ImagRelDispGyy, 431

ImagRelDispGzz, 431

imoment, 398

impedance_pressure, 366

impedance_shear, 366

info, 22

initial overlap, 223

initial-conditions, 406

- time, 407
- velX, 406
- velY, 406
- velZ, 406

input

- acceleration, 362
- Aprepro, 14, 28
- comment
 - entire section, 27
- comments, 26
- end, 26

input *see hand-off*, 26

interaction behavior, 469

interpolation, 40

inverse solutions

- directfrf, 138
- eigen, 147
- modalfrf, 167
- modaltransient, 181
- transient, 210

inverse-directfrf, 231

inverse-problem, 138, 141, 147, 168, 183, 211, 230, 243, 245

- data, 231
- data_file, 236
- data_truth_table, 232
- data_type, 232
- data_weight_table, 238
- imaginary_data_file, 236
- modal_data_file, 237

- modal_weight_table, 237
- psd_data_file, 236
- real_data_file, 234
- inverse_material_type, 238
- inverseBlock
 - heterogeneous, 238
 - homogeneous, 238
 - known, 238
- inverseMaterial
 - acoustic, 239
 - Aij_bounds, 239
 - bulk, 239
 - c0_bounds, 239
 - density_bounds, 244
 - E_bounds, 239
 - Eij_bounds, 239
 - filter_radius, 245
 - G_bounds, 239, 244
 - Gij_bounds, 239
 - Gim_bounds, 239
 - Greal_bounds, 239
 - impedance_match, 239
 - isotropic, 239
 - isotropic_viscoelastic, 239
 - K_bounds, 239, 244
 - Kim_bounds, 239
 - Kreal_bounds, 239
 - material_parameters, 239
 - Nu_bounds, 239
 - num_material_parameters, 239
 - orthotropic, 239
 - penalizationElasticity, 244
 - penalizationMass, 244
 - rho, 239
 - shear, 239
 - smooth_heaviside_slope, 245
 - smooth_heaviside_threshold, 245
 - sound_speed, 239
- ipressure, 398
- isotropic, 249, 250
- isotropic_viscoelastic, 249
- isotropic_viscoelastic_complex, 257
- itraction, 398
- Iwan, 278, 322
- Joint2G, 326, 327, 342, 353, 357
 - relative_disp, 431
- keepmodes, 173
- krylov_solver_output_file, 60
- layer, 293
- left_stretch, 252
- lfcutoff, 173
- Lighthill, 385, 386
- linesample, 213, 369, 448
- load, 119, 120, 141, 174, 207, 220, **381**, 381, 394, 405
 - body, 381
 - complex, 398
 - consistent, 392
 - electrostatic, 393
 - follower, 382
 - follower stiffness, 384
 - function, 394
 - randompressure, 394
 - scale, 394
 - spatially-dependent, 374–377
 - statics, 393
 - transient, 394
- load identification
 - directfrf, 139
 - limitations, 248
 - load entry, 246
 - modalfrf, 168
 - modaltransient, 181
 - transient, 211
 - PSD modalfrf, 170
- loads, 99, 100, 119, 120, 140, 141, 156, 169, 171, 174, 183, 212, 246, 248, 362, 364, **380**, 380, 381, 384, 402, 405
- Lumped *see mass matrix lumping*, 51
- mass, 455
 - blockwise properties, 455
 - properties, 455
- mass matrix lumping, 51, 289, 295, 479
- material, 140, 212, 239–241, 244, 245, **249**, 252
 - Lamé, **252**

- begin-lame-material, 252
- end-lame-material, 252
- lame_state_hyperfoam, 252
- acoustic, 253
- anisotropic, 251
- complexViscoelastic, 257
 - Gim, 257
 - Greal, 257
 - Kim, 257
 - Kreal, 257
- density, 258
- E, 250
- example, 479
- exodus mesh properties, 263
- G, 250
- isotropic, 250
- K, 250
- layered, 290
- nu, 250
- orthotropic, 250
- specific heat, 264
- temperature dependent, 262
- temperature function, 262
- viscoelastic, 254
- material identification
 - directfrf, 140
 - eigenvalue, 147
 - eigenvector, 148
 - multi directfrf, 141
 - transient, 212
 - block entry, 238
 - material entry, 239
- material_direction_1, 414
- material_direction_2, 414
- material_direction_3, 414
- matrix, 405
 - file names, 427
 - output in MFile format, 427
 - RanLoads parameter, 405
- matrix-function, 173, 174, 405
 - nominalt, 109
 - table, 109
- MatrixFloor, 40
- max_newton_iterations, 195
- MaxmpcEntries, 41
- maxRatioFlexibleRbm, 278
- membrane_factor, 291
- memory diagnostics, 22
- Memusage, 457
- mesh discretization error, 426
- mesh_error, 426
- MFile, 426
- MinimumNodalSpacing, 397
- mksuper, 342, 344, 345
- modal acceleration, 164
- modal effective mass, 160, 161
 - meff, 160
 - MPF, 160
- modal_amp, 427
- modal_data_file, 237
- modal_weight_table, 237
- modalfilter, 145
- modalfiltercase, 145
- ModalFraction, 190
- ModalFrf, 137, 163–165, 276, 483
 - example, 483
 - nrbms, 165
 - usemodalaccel, 165
- modalfrf-inverse, 167
- modalranvib, 173, 173, 174, 276, 429, 431, 432
 - acceleration, 173
 - noSVD, 173
 - RMS von Mises stress, 173
- modalshock, 177
- modaltransient, 48, 49, 119, 178–180, 273, 276, 398, 481
 - example, 481
- modaltransient-inverse, 181, 226
- modalvars, 179, 399, 456
- mortar method, 461
 - dual, 461
 - standard, 461
- MPC *see multipoint constraint*, 41
- mpmd_transfer_sidesets, 40
- mpmd_transfer_type, 40
- mpmd_transfer_version, 40
- multipoint constraint, 38, 113
 - constraint force, 428
 - constraint_correction, 41, 208

- constraint_info, 435
- nUpdateConstraints, 208
- orthogonalization, 41
- problematic, 25

NASTRAN

- autospc, 38
- output4 in CBR, 127

Nbeam, 301, 305, 305–308

NegEigen, 142

neglect_mass, 368

netcdf, 132, 345–347

new, 40

Newmark beta, 208, 209

Ng, Esmond, 2

NLstatics, 193–195

NLtransient, 194, 195

nmodes, 124, 145, 158, 165, 173, 276

Nmount, 336, 336

- stability, 336

nodal, 441

nodal_charge, 439

NodeListFile, 364, 443

nodeset, 126, 376, 381

nodesets_with_disp, 44, 208

nonlinear, 271, 274

nonlinear_default, 36, 271

normal tolerance, 468

npresure, 221, 434

Nquad, 287, 287, 288

nquad, 283

Nquad/Ntria, 287, 288

nquad_eps_max, 289

Ntria, 287, 287, 288

num_newton_load_steps, 193, 195

num_procs, 45

num_rigid_mode, 43

nUpdateConstraints, 44

nUpdateDynamicMatrices, 263, 390

nUpdateTemperature, 263, 390, 391

offset shell, 295, 295

old, 40

optimization, 138, 147, 168, 183, 211, 226

orthotropic, 249, 250, 289

orthotropic_layer, 251

orthotropic_piezoelectric, 266

orthotropic_prop, 249, 250

output

- element force, 429
- error metrics, 426
- history, 443
 - coordinate, 431
 - nodes, 443
 - relative_disp, 431

Matlab, 411, 449

- exo2mat, 448

relative_disp, 431

user output, 440

- block statistics, 440
- closest distance, 441
- derived variables, 442

output_sideset_data, 36

outputs, 38, 47, 127, 137, 154, 165, 174, 179, 224, 330, 346, 386, 410, 410, 411, 413, 424, 425, 429, 431–435, 437, 438, 443

- acceleration, 417
- constraint_force, 428
- constraint_info, 435
- coordinate, 431
- disp, 417
- energy, 424
- faa, 415
- force, 428
- GEnergies, 424
 - work, 424
- Globals, 425
- history, 128
- Kaa, 414
- line_weld, 431
- Maa, 414
- material, 414
- Matlab, 17, 128, 132, 160, 179, 213, 369, 456
 - MFile_Format, 42
- MLumped, 415
- MPhi, 415
- reaction_force, 428

- RHS, [428](#)
- RMS, [424](#)
- rotational_acceleration, [424](#)
- rotational_displacement, [424](#)
- strain, [418](#)
- stress, [419](#)
- velocity, [417](#)
- vonMises stress, [420](#)
- overlap removal *see Gap Removal*, [223](#)
- p, [54](#)
- Padé, [137](#)
- parallel computing, [18](#)
 - e_{pu}, [17](#), [18](#)
 - exodus_file, [17](#)
- parameters, [33](#), [33](#), [54](#), [55](#), [210](#), [339](#), [390](#), [403](#), [461](#)
 - info, [36](#)
 - negeigen, [33](#)
 - reorder_RBar, [38](#)
 - reserved_keywords, [37](#)
 - syntax_checking, [37](#)
 - wtmass, [33](#)
- performance, [15](#), [20](#), [53](#), [113](#)
- periodic boundary conditions, [371](#)
- permittivity_ij, [266](#)
- piezoelectric, [224](#)
 - e_ij, [266](#)
- plane_wave, [101](#)
- point_volume_accel, [385–387](#)
- point_volume_acceleration, [7](#), [378](#)
- point_volume_vel, [385–387](#)
- point_volume_velocity, [378](#)
- power spectral density, [175](#), [406](#)
 - acceleration, [175](#)
 - coordinate, [175](#)
 - displacement, [175](#)
 - load identification, [170](#)
 - scaling, [405](#)
- preddam, [145](#)
- PredictorCorrector, [208](#)
- prescribed acceleration, [362](#)
 - accelX, [362](#)
 - accelY, [362](#)
 - accelZ, [362](#)
 - disp0, [362](#)
 - RotaccelX, [362](#)
 - RotaccelY, [362](#)
 - RotaccelZ, [362](#)
 - vel0, [362](#)
- prescribed frequency, [364](#)
 - DispX, [364](#)
 - DispY, [364](#)
 - DispZ, [364](#)
 - FreqP, [364](#)
 - FreqV, [364](#)
 - RotDispX, [364](#)
 - RotDispY, [364](#)
 - RotDispZ, [364](#)
- pressure, [382](#), [385](#), [434](#)
 - depth dependent, [392](#)
 - nodal, [434](#)
 - prescribed, [361](#)
- pressure_z, [392](#), [393](#)
- Problematic Elements, [23](#)
- Problematic Subdomains, [22](#)
- projection_eigen, [184](#), [185](#), [191](#)
- property, [319](#), [326](#), [327](#), [413](#)
- PSD *see power spectral density*, [175](#)
- qevp, [164](#), [166](#), [183–186](#), [188](#), [190](#)
- Quad8T, [284](#), [284](#)
- QuadM, [286](#), [286](#)
- quadm, [283](#)
- QuadT, [284](#), [284](#)
- quadt, [283](#)
- Random Number Generator, [42](#)
- RandomPressure, [395](#)
- RanLoads, [174](#), [404](#), [404](#), [405](#)
 - acceleration, [404](#)
 - dimension, [404](#)
- rational function, [137](#)
- RBar, [339](#), [339](#), [340](#), [350](#)
- RBE2, [340](#)
- RBE3, [340](#), [341](#)
 - refc, [341](#)
- RbmTolerance, [40](#)
- real_data_file, [234](#)
- Reduced Iwan, [325](#)

- regularization, 230
- relative_disp, 11, 431, 432
- reorthogonalization, 166
- reserved_keywords, 32, 37
- residual, 432
 - global var, 432
 - nonlinear norm, 454
 - vector, 432, 454
- residual work, 432
- residual_vectors, 215
 - nrbms, 216
- restart, 45
 - auto, 45
 - database name, 45
 - eigen, 49
 - from, 45
 - num_procs, 45
 - read, 45
 - restart_consistency_checking, 38
 - solution support, 48
- restart_consistency_checking, 38
- rigid, 354
 - rigidset, 348, 350
 - Joint2G, 349
 - limitations, 350
 - nodeset, 348
 - sideset, 348
 - tied joint, 349
 - voltage, 349
 - Rrodset, 350
- Rigid Body Filter, 403
- RMS, 173, 424
- Rod *see* Truss, 309
- ROL output, 248
- rotation, 252
- rotational frames, 399
- rotational_acceleration, 174, 175
- rotational_displacement, 174, 175
- rotational_type, 272
- rotx, 54
- roty, 54
- rotz, 54
- Rrod, 338
- rrod, 354
- RSpring, 312, 313, 313, 416
- sa_eigen, 166, 184, 185, 188
 - ErrorNorm, 190
 - limitations, 188
- Salinas, 13
- sc_option, 149
- scaleDesignVars, 226
- scattering, 52
- sd_factor, 281, 282, 287
- Section Commands
 - Loads Rigid Body Filter, 403
 - block, 268
 - Block Parameters, 268
 - boundary, 358
 - coordinate, 76
 - damping, 275
 - echo, 453
 - file, 29, 410
 - frequency, 100, 446
 - function, 84
 - history, 443
 - Load, 381
 - loads, 380
 - material, 249
 - matrix-function, 109
 - outputs, 410
 - parameters, 33
 - periodic boundary conditions, 371
 - RanLoads, 404
 - Rrodset, 350
 - sensitivity, 71
 - solution, 116
 - solver_options, 59
 - table, 111
 - tied surfaces, 458
 - user output, 440
 - block statistics, 440
 - closest distance, 441
 - derived variables, 442
- self contact, 468
- sensitivity, 71, 74, 131, 481
 - +/-, 74
 - Attune, 72
 - iterations, 71
 - tolerance, 71
 - values, 71

- vectors, [71](#)
- sensitivity_param, [346](#)
- shear_axis, [320](#)
- shells
 - offset, [295](#)
- shift, [144](#), [157](#)
- sideset, [126](#), [297](#), [377](#), [381](#), [382](#)
- SkipmpcTouch, [38](#)
- smoothing parameters, [462](#)
- Solution, [26](#), [44](#), [51](#), [52](#), [116](#), [116](#), [119](#), [144](#), [157](#), [174](#), [193](#), [201](#), [220](#), [479](#)
 - AEigen, [151](#)
 - Buckling, [156](#)
 - ceigen, [186](#)
 - Checkout, [122](#)
 - CJdamp, [122](#)
 - complex eigen, [186](#)
 - Craig-Bampton reduction, [123](#)
 - DDAM, [133](#)
 - DirectFRF, [136](#)
 - DirectFRF-Inverse, [138](#)
 - Dump, [141](#)
 - Eigen, [142](#)
 - Eigen-Inverse, [146](#)
 - Fatigue, [153](#)
 - Gap Removal, [223](#)
 - elementInversionFlag, [224](#)
 - elementQuality, [224](#)
 - GeometricRigidBodyModes, [12](#), [217](#)
 - LargestEv, [153](#)
 - Modal Participation Factor, [159](#)
 - ModalFilter, [157](#)
 - write_files, [157](#), [159](#)
 - ModalFrf, [163](#)
 - ModalFRF-Inverse, [167](#)
 - ModalRanVib, [172](#)
 - ModalShock, [177](#)
 - ModalTransient, [178](#)
 - ModalTransient-Inverse, [181](#)
 - MPF, [160](#)
 - NLStatics, [192](#)
 - NLTransient, [194](#)
 - Options, [44](#)
 - constraint method, [52](#)
 - GDSW, [50](#)
 - mass matrix lumping, [51](#)
 - restart, [45](#)
 - scattering, [52](#)
 - symmetrize_struc_acous, [52](#)
 - Table, [120](#)
 - Q EVP, [183](#)
 - ReceiveSierraData, [119](#), [195](#), [197](#), [200](#), [251](#), [258](#)
 - include_internal_force, [196](#)
 - no_geom_stiff, [197](#)
 - Residual Vectors, [215](#)
 - Statics, [201](#)
 - Superposition, [202](#)
 - Tangent, [203](#)
 - TranShock, [204](#)
 - srs_damp, [205](#)
 - Transient, [206](#)
 - Transient-Inverse, [210](#)
 - TSRpreload, [213](#)
 - Waterline, [219](#)
- solver
 - parameters, [53](#)
- solver_options, [60](#)
- sparc, [40](#)
- spherical_wave, [102](#)
- Spring, [312](#), [312](#), [313](#), [416](#)
 - cubic, [313](#)
 - Linear, [312](#)
 - Parameter Values, [312](#)
 - Rotational, [313](#)
- Spring3, [313](#), [313](#), [314](#), [416](#)
- SpringDashpot, [315](#), [315](#)
- statics, [201](#), [388](#), [429](#)
 - example, [486](#)
- statistics, [436](#)
 - min/max, [436](#)
 - standard deviation, [436](#)
- stiffness_scale_factor, [275](#)
- strain, [418](#)
 - Gauss point, [419](#)
- stress, [174](#), [252](#), [452](#)
 - Gauss point, [422](#)
- Stress and Strain Recovery, [449](#)
- stress recovery point, [452](#)
- Structural Acoustics

- eigen, [188](#), [191](#)
- StructuralFraction, [190](#)
- subdomain
 - output, [456](#)
- Superelement, [130](#)
- SuperLU, [2](#)
- superposition, [202](#)
- surface
 - description of, [30](#)
- surface_charge, [393](#)
- symmetrize_struc_acous, [53](#)
- syntax_checking, [37](#)
- tangent, [119](#), [121](#), [203](#), [401](#), [402](#)
- Tangential Tolerance, [468](#)
- TangentMethod, [36](#)
- tcoord, [297](#), [298](#)
- Temperature, [439](#)
- Tet10, [283](#), [283](#)
- Tet4, [283](#), [283](#)
- thermal_exo_var, [39](#), [390](#), [392](#)
- thermal_load, [39](#), [383](#), [388](#), [390](#), [392](#)
- thermal_strain, [418](#)
- thermal_time_step, [39](#), [263](#), [390–392](#)
- TiBeam, [308](#), [416](#)
- tied data, [223](#), [458](#)
 - gap_removal, [25](#), [223](#)
 - initial overlap removal, [223](#)
 - transverse, [458](#)
- tied joint
 - normal, [352](#)
 - shear_axis, [353](#)
 - side, [353](#)
 - slip, [353](#)
 - surface, [352](#)
- tied surface, [458](#)
 - gap removal, [458](#)
- TIndex, [433](#)
- tolerance, [192–194](#)
- traction, [382](#)
- transhock, [204](#), [205](#)
- transient, [49](#), [204](#), [206](#), [388](#), [429](#)
- nskip, [206](#), [444](#)
- nsteps, [206](#)
- start_time, [206](#)
- time_step, [206](#)
- transient-inverse, [210](#), [226](#)
- Tria3, [284](#), [290](#), [290](#)
- Tria6, [284](#), [284](#)
- TriaShell, [284](#), [289](#), [289](#), [290](#), [451](#)
- troubleshooting, [20](#)
 - cubit, [21](#)
 - explore, [21](#)
- TruncationMethod, [173](#)
- Truss, [309](#), [309](#), [416](#)
- units of measure, [33](#), [35](#)
- update_tangent, [193](#), [195](#), [331](#)
- UseAnalystNodeMap, [44](#)
- user subroutine file, [29](#), [336](#)
- useTransferMatrix, [231](#)
- viscofreq, [187](#)
- viscous damper, [314](#)
- voltage, [439](#)
 - accelV, [363](#)
 - timeV, [363](#)
- Volume, [439](#)
- volume_acceleration, [375](#)
- vrms, [424](#)
- waterline, [219](#)
 - point_a, [219](#)
 - point_b, [219](#)
 - point_c, [219](#)
- Wedge15, [283](#)
- Wedge6, [282](#), [282](#)
- write_files, [157](#), [179](#)
- wtmass, [33](#), [406](#), [455](#)
- x, [54](#)
- XML, [14](#)
- y, [54](#)
- z, [54](#)

DISTRIBUTION

Hardcopy—Internal

Number of Copies	Name	Org.	Mailstop
1	K. H. Pierson	1542	0845

Email—Internal

Name	Org.	Sandia Email Address
Technical Library	1911	sanddocs@sandia.gov

This page intentionally left blank.



Sandia
National
Laboratories

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.