

Graph Partitioning for Computational Mechanics Simulations with GDSW



PRESENTED BY



David Day: Sandia National Laboratories



Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and 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.



APPLICATION Space





Challenges in Computational Mechanics

- Normal Environments: Structural Dynamics
 - Linear elasticity (this talk)
 - Sparse symmetric generalized eigenvalue problems:
 - Implicit linear transient analysis
 - predict linear/nonlinear response (Civil/Mechanical Engineering)
 - Computational bottleneck: linear system, successive right-hand sides
 - GDSW: Overlapping Additive Schwarz
- Hostile Environments: Computational Mechanics
 - Nonlinear elasticity
 - Implicit transient and quasi-static analysis:
 - shares many of the computational challenges of structural dynamics
 - FETI-DP now,, GDSW longer term
 - Explicit transient: Just works, not this talk



SUMMARY



Summary: Graph Partitioning and Linear Solvers

Graph Partitioning:



- finite element mesh partitioned into subdomains
- expect constant nodes per subdomain
- well defined linear problem

Domain Decomposition Linear Solvers

- unsurpassed scalability for 20 years, highly cited
- (not this talk) other linear solvers
- challenges to developing robust linear solvers for computational mechanics
- simplified requirements on the partition limit scalability
- deep dive: partition requirements
- 3 examples



TECHNICAL CHALLENGES



Example Problems

- severely reduced numbers of time steps or modes
- practically all of the compute time is spent in the linear solver
- linear solver initialization is not counted in the linear solve time

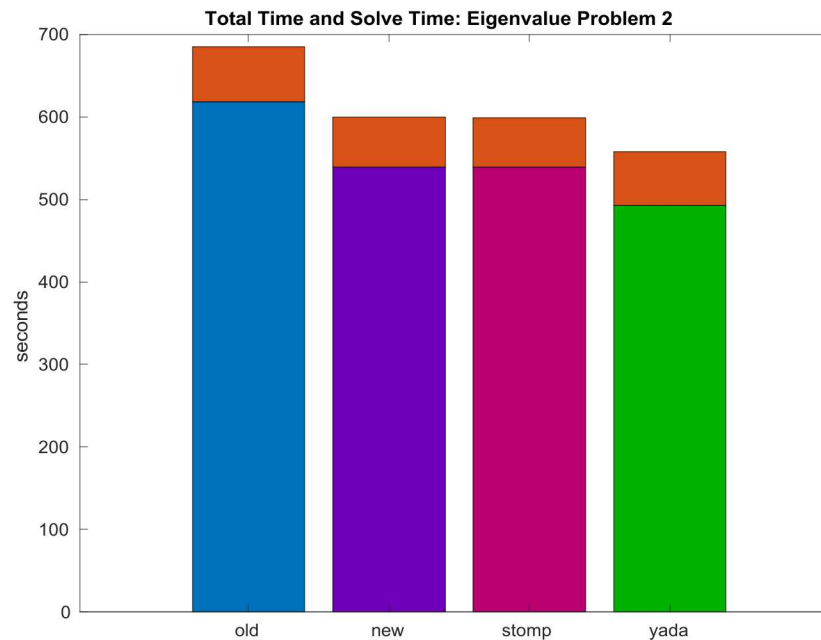
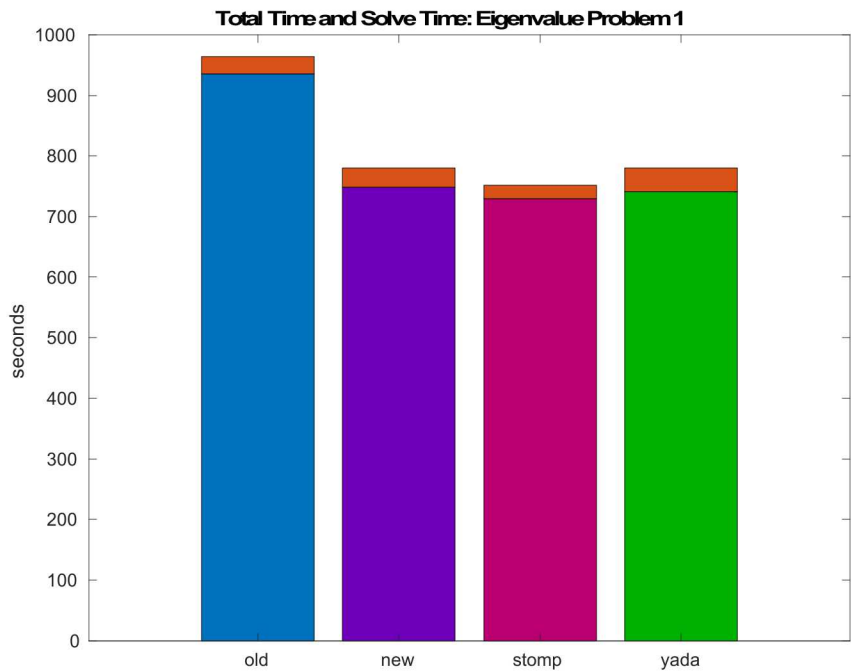
	Eigenvalue 1	Eigenvalue 2	Transient 3
Unknowns	4437565	5212584	5173605
Processors	252	216	216
Linear Solves	58	82	100

All examples have constraints

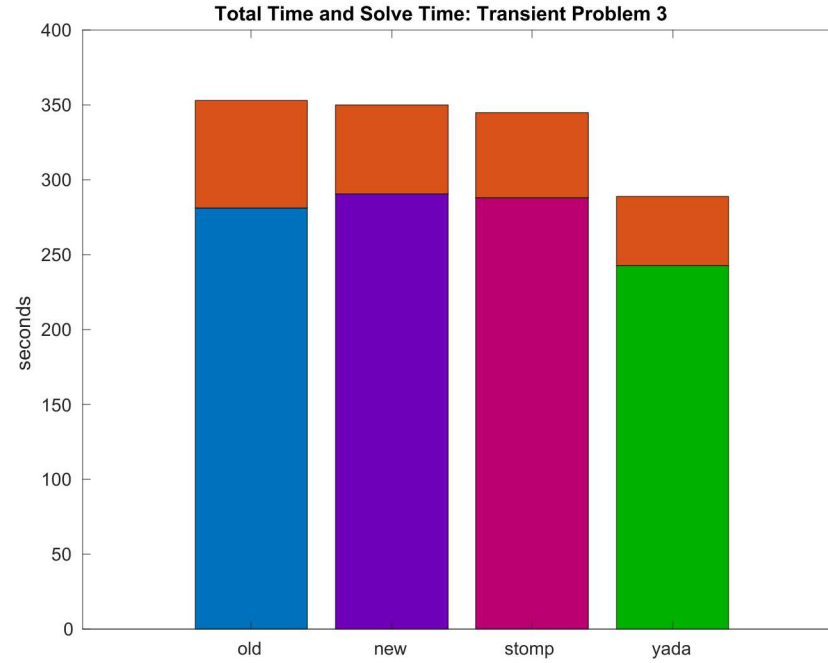
Constraints: First Layer of Complexity

1. Graph Partitioner: sees $K u = f$
2. Homogeneous constraints $C u = 0$
3. Explicit Elimination: independent and dependent unknowns
4. Node-Face Tied Data: $C u = 0$ transformed to $[B \ I] [u_i; u_d] = 0$
5. Linear Solver: sees $K_{ii} u_i = f_i$

9 Linear Solve Time as a Fraction of Total Time



Linear Solve Time as a Fraction of Total Time



11 Different Partitioners

Hypothesis: All the partitioners are solving an over-simplified problem

Zoltan2 mesh interface, Sierra Tool Kit, STK Balance

1. Old: old version of stk balance
2. New: my version of stk balance
3. Stomp: new with “spider stomper” (whole other talk)
4. yada = Chaco (serial)

This talk spurred the current version of stk-balance

current stk-balance not yet configured for computational mechanics

yada is used in legacy models (and I use it)

GDSW Bottlenecks

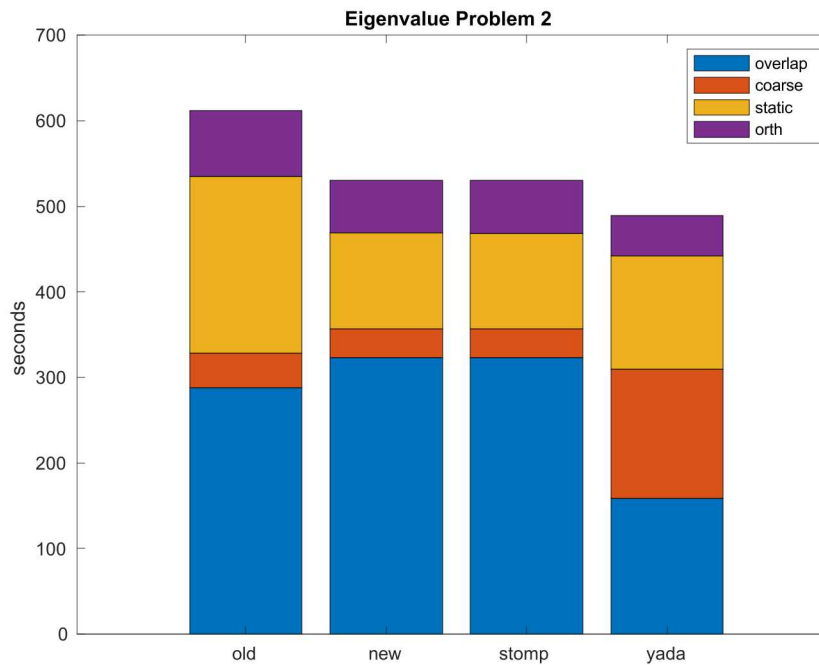
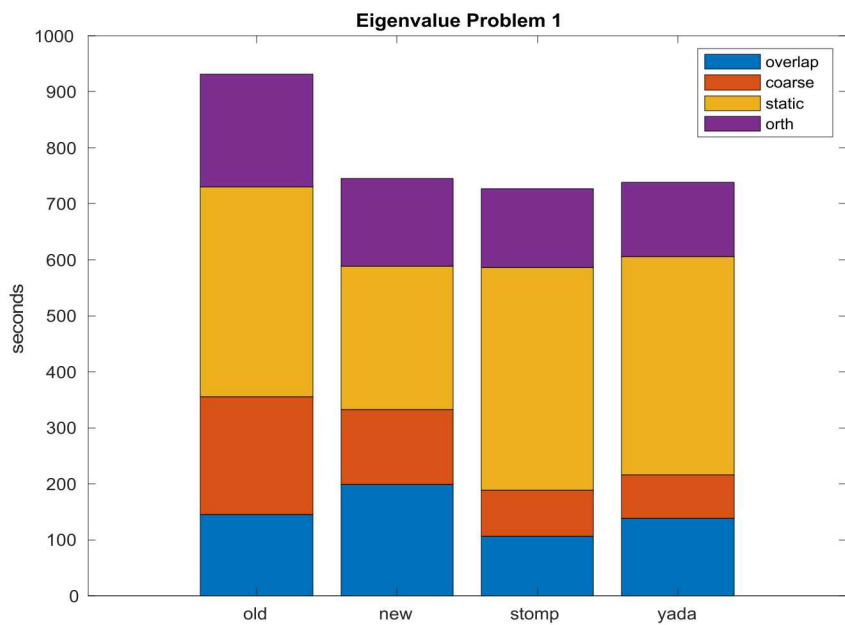
Overlap operator (overlap non-zeros)

Static condensation (interior non-zeros)

Coarse correction (num processors)

Orthogonalization (overlap unknowns)

Linear Solver Time Itemization



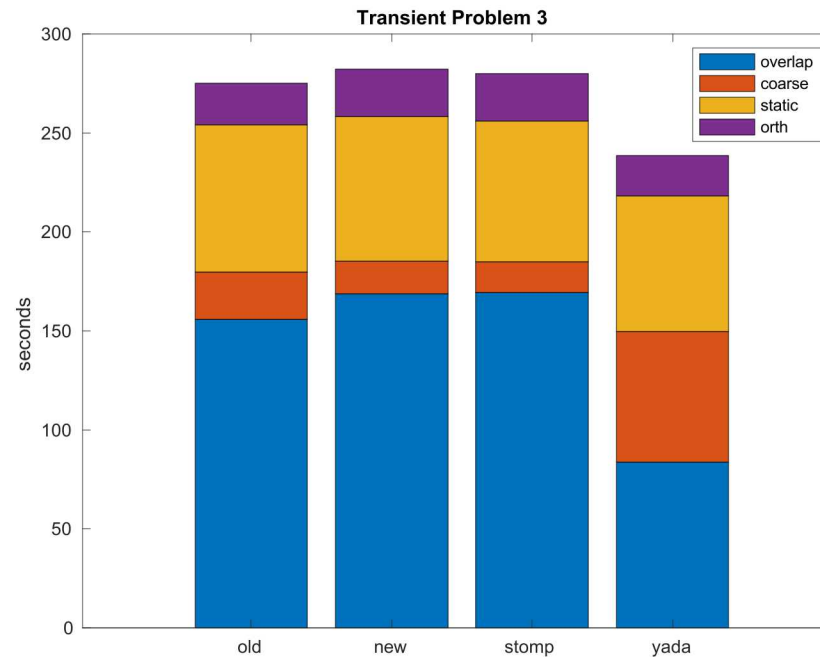
Linear Solver Time Itemization continued

Overlap operator is significant

Static condensation is too

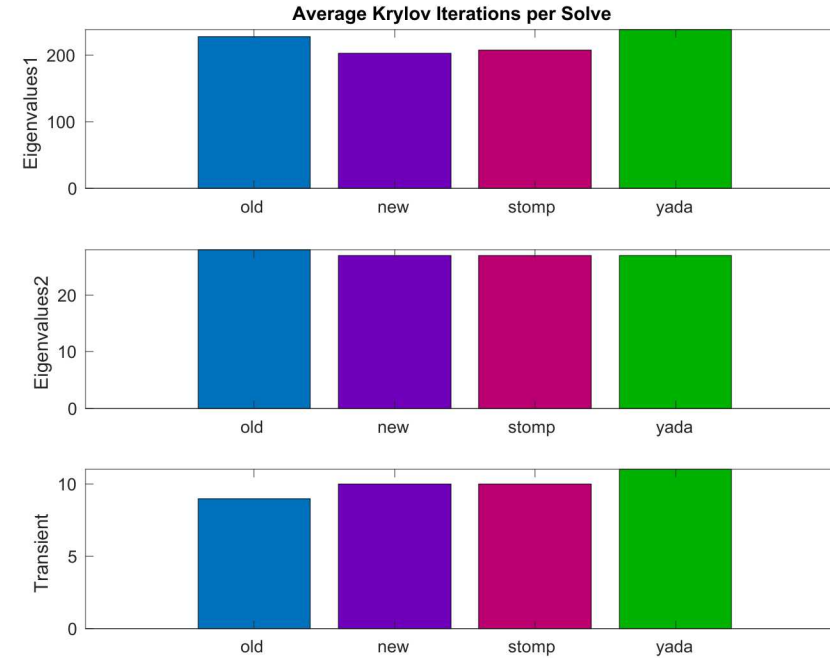
Information overload

Not yet predictable a priori- That's acceptable



Partitions influence preconditioner strength

not critical here
 information overload
 significant

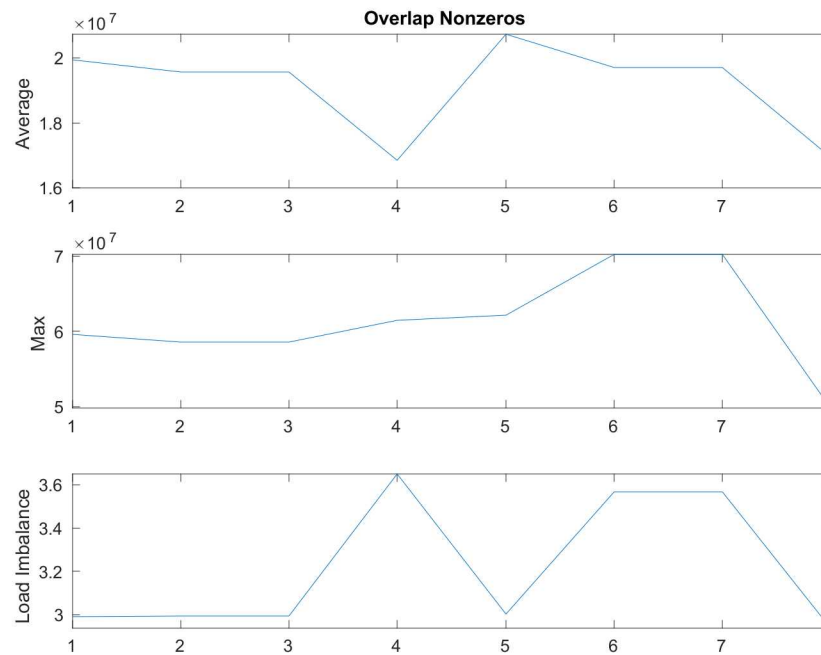


16 Load Imbalance: Overlap

Load imbalance = max/ave

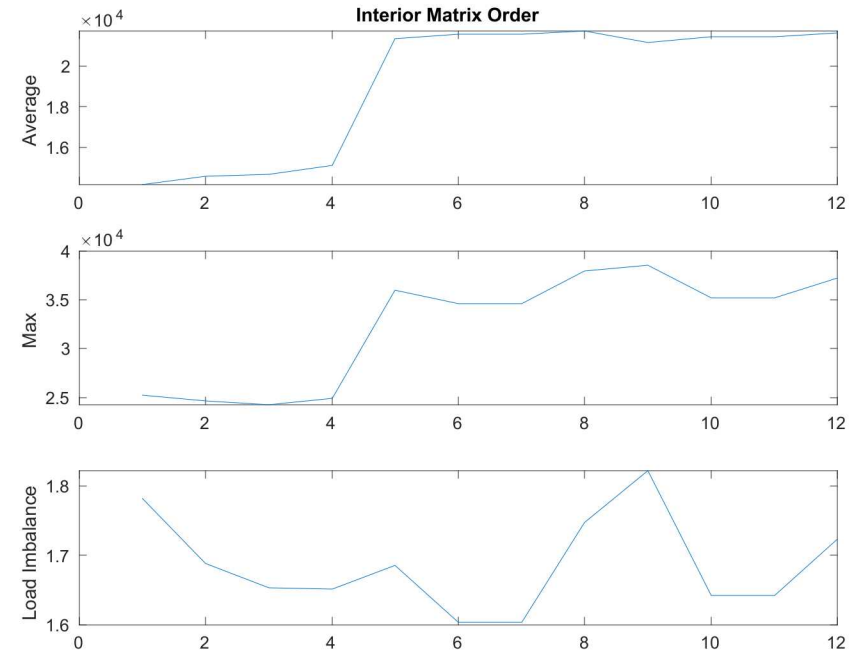
$3 < \text{Load imbalance} < 4$

Application specific “elements”
influence overlap size



Load Balance: Static Condensation

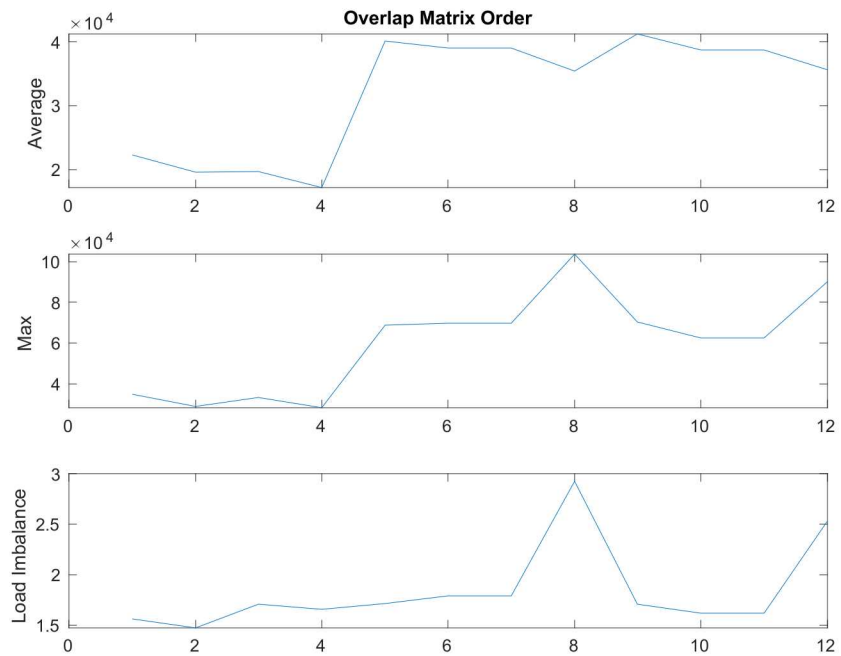
Non-zeros in the factorization of the subdomain interior matrix



Load Imbalance: Orthogonalization

Unknowns in overlap matrix

$$2 < LB < 3$$



Conclusion

1. Load imbalance inhibits scalability in computational mechanics.
2. There are unsolved problems between graph partitioning and scalable iterative linear solvers
3. For GDSW the partitioning problem has multiple objectives and is nonlinear.
4. Potential impact: faster simulations on the same number of processors
5. Potential impact: rehabilitate the parallel processing route to performance

Acknowledgements

Sierra SD Team:

Nate Crane, Clark Dohrmann, Sean Hardesty, Payton Lindsey, Lynn Munday, Brian Stevens

Sierra STK Team:

Dave Glaze, Don Hun Lee, Tolulope Okusanya, Tim Smith, Jonathan Vo, Alan Williams

Zoltan2 Developers, Sandia Analysts