

Sandia  
National  
Laboratories

*Exceptional  
service  
in the  
national  
interest*



U.S. DEPARTMENT OF  
**ENERGY**



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

# A Whirlwind Tour Through the Mathematics and Computer Science of Computational Simulation in Industry

Dr. David E. Womble

Senior Manager, Computational Simulation Group

Sandia National Laboratories

# A Quick Overview of Sandia National Laboratories



*Albuquerque, New Mexico*



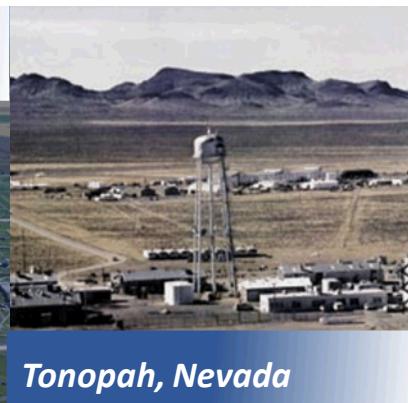
*Livermore, California*



*Kauai, Hawaii*



*Pantex Plant,  
Amarillo, Texas*



# Nuclear Weapons

## Pulsed power and radiation effects sciences



## Warhead systems engineering and integration

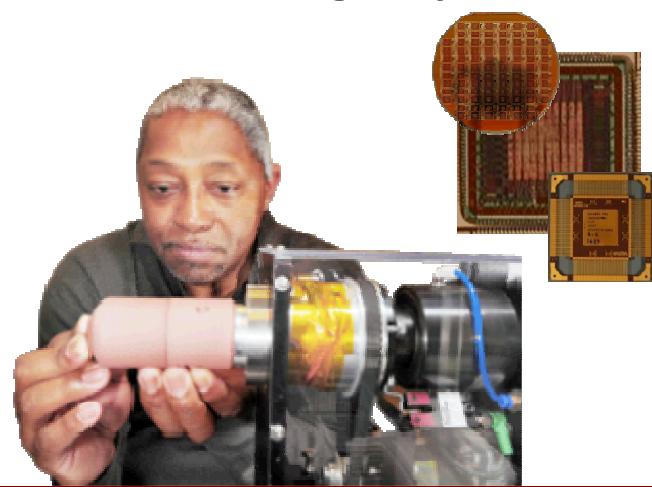


## Design agency for nonnuclear components

- Neutron generators
- Arming, fuzing and firing systems
- Safety systems
- Gas transfer systems

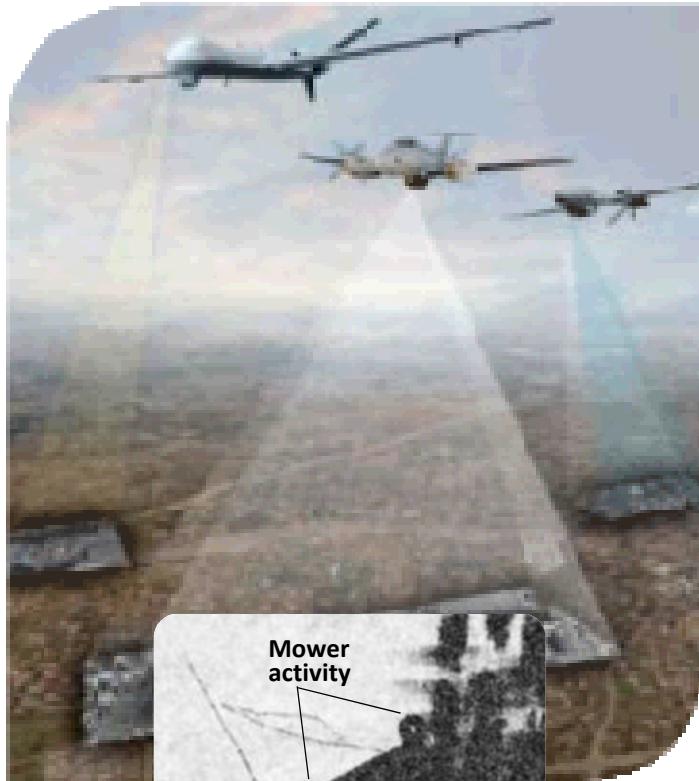


## Production agency



# Defense Systems and Assessments

## Synthetic aperture radar



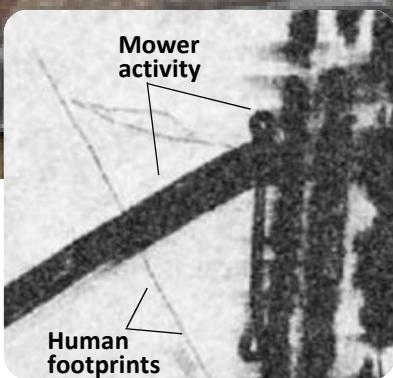
## Support for NASA



## Support for ballistic missile defense



## Ground sensors for future combat systems

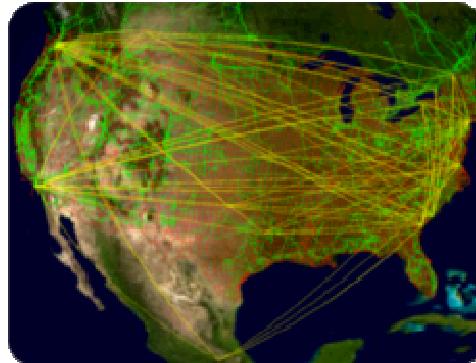


# Energy, Climate, and Infrastructure Security

## Energy



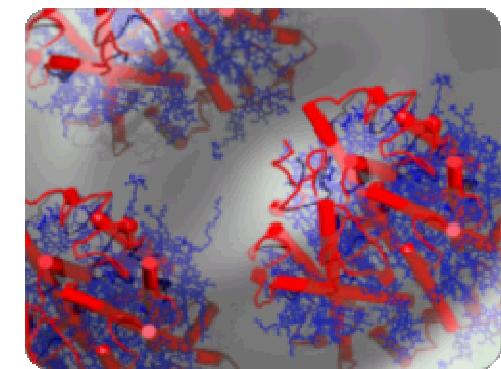
## Infrastructure



## Crosscuts and enablers



## Climate



# International, Homeland, and Nuclear Security

## Critical asset protection



## Homeland defense and force protection



## Homeland security programs



## Global security

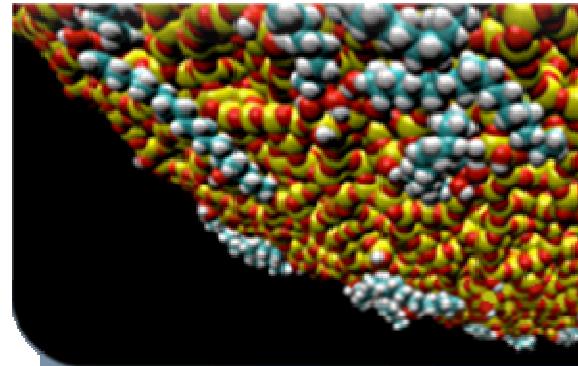


# Science and Engineering Foundations

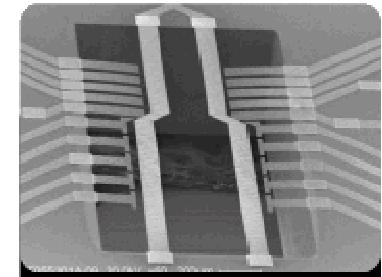
## Computing and information science



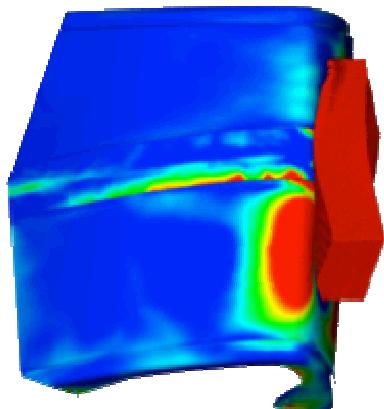
## Materials science



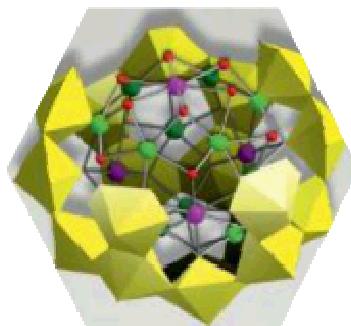
## Nanodevices and microsystems



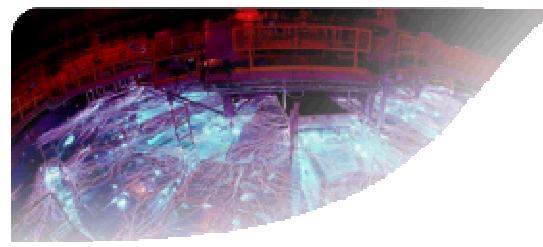
## Engineering sciences



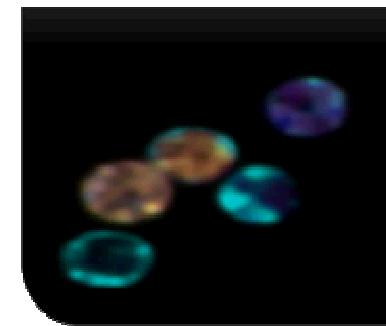
## Geoscience



## Radiation effects and high-energy density science



## Bioscience



# Now back to computational simulation

## What type of computing am I talking about?

### Physics-based computing

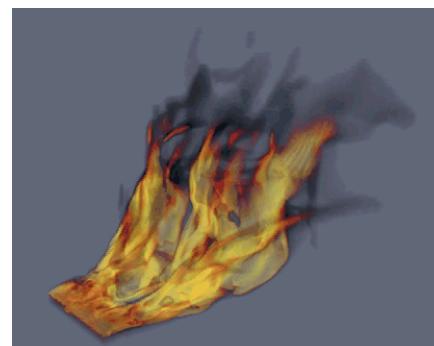
- Starts with a model, often a PDE or conservation law
- And a physical system
- And an initial/boundary conditions
- Then evaluates the model for that system.

### Information-based computing

- Starts with data
- And searches for meaningful patterns or conclusions in the data
- Models exist, but are most often implicit and indirectly evaluated. (An analysis of Google FluTrends highlights the problems.)

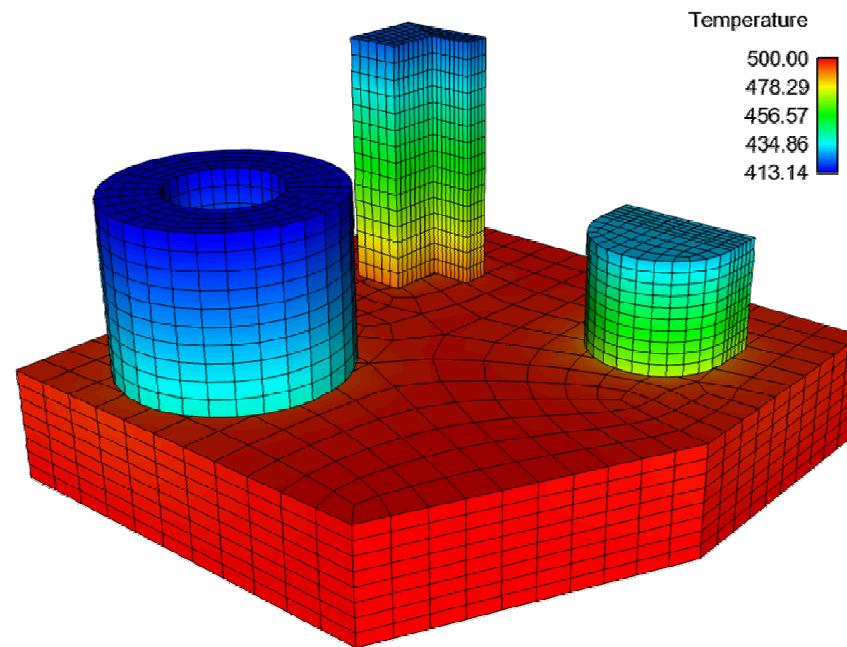
# As an example, suppose we want to know the response of a system in a fire.....

- Fire is one source of a thermal “insult”
- There are many types of fires, and each creates a different environment.
- The physics of a fire is interesting, but let's take that as a boundary condition for now.



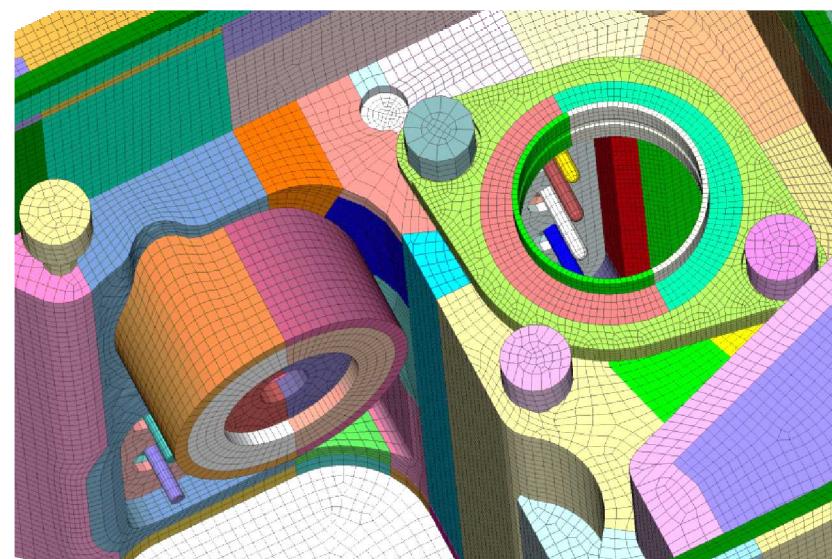
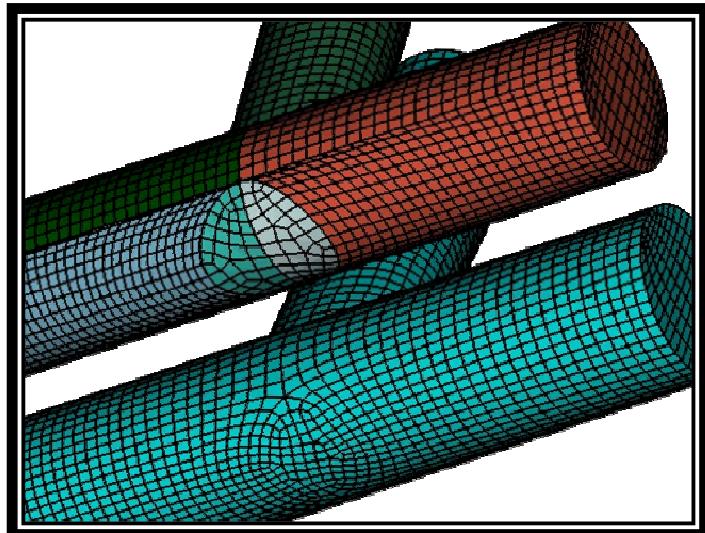
# Let's start with thermal transport

- What are the equations
  - Diffusion
  - Convection
  - Radiation
- What discretization should we use?
  - Finite difference? Finite volume? Finite element?
  - What type of element?
  - What order?



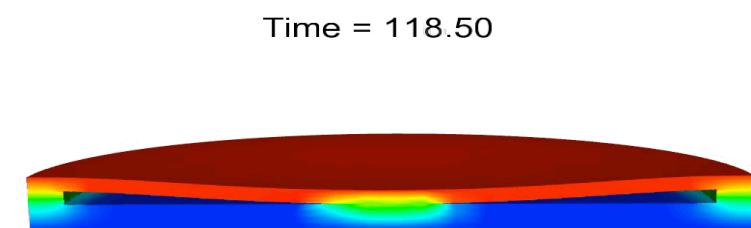
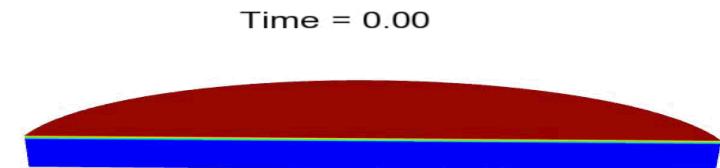
# What are some of the other decisions that must be made?

- How do we mesh?
  - Geometry cleanup
  - Defeaturing
  - Meshing and mesh quality
- What solvers will be used?
  - Direct
  - Iterative
  - Preconditioned CG
- Are we ready to solve the thermal problem

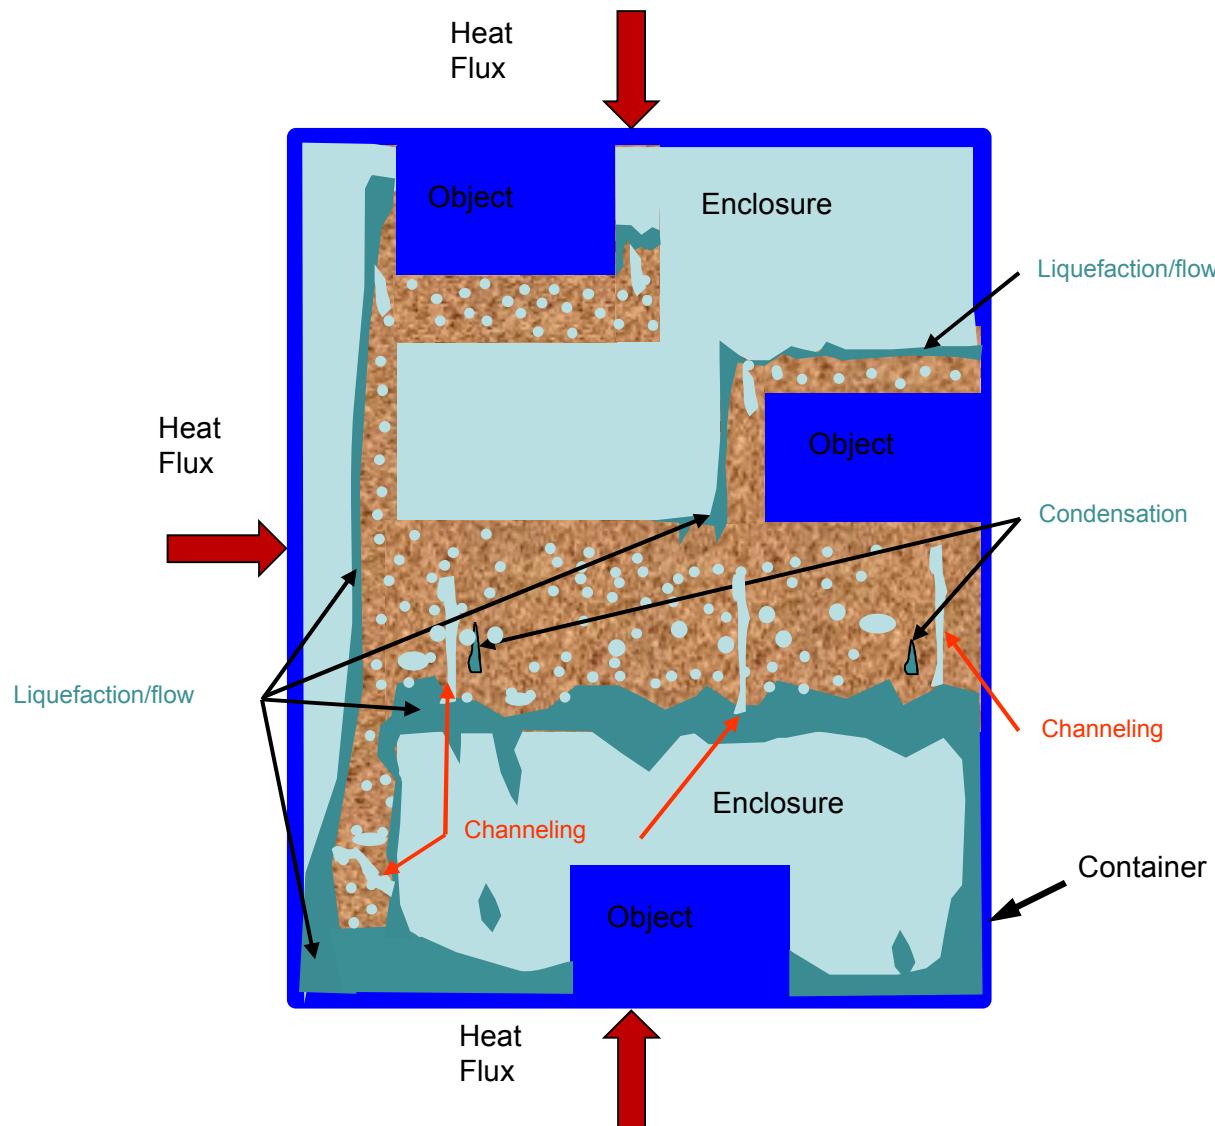


# But we also know there are structural changes in thermal environments

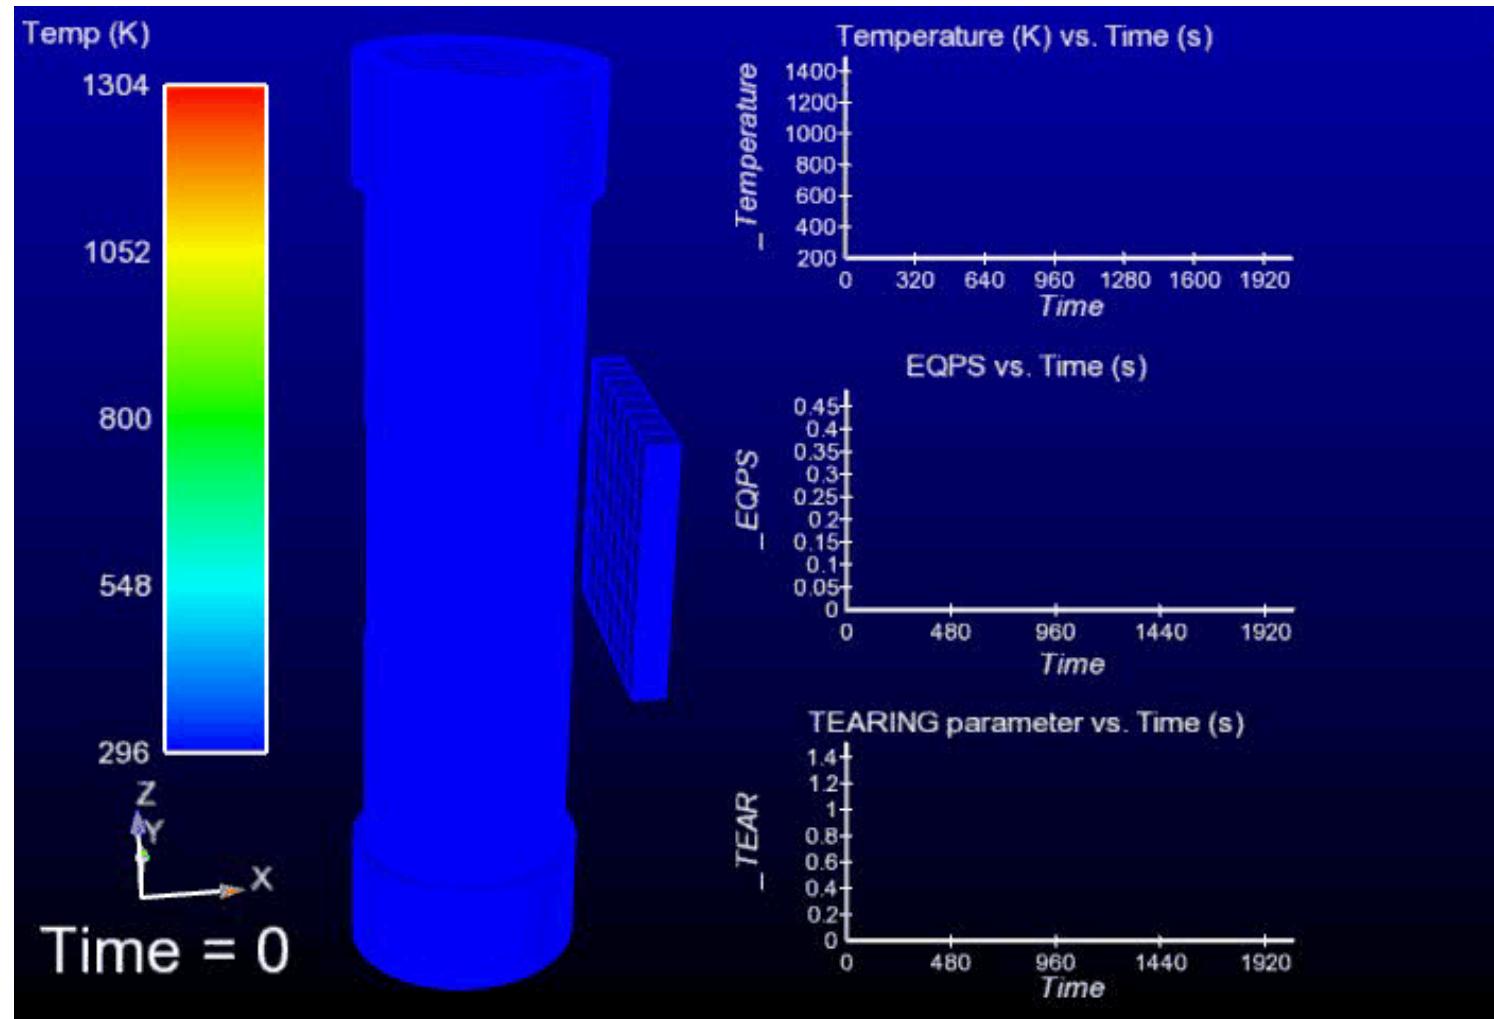
- What new physics is introduced?
  - Solid mechanics, including thermal stress
  - Contact can create or eliminates thermal paths
  - Phase change
- We also have to go back and ask the same questions as before
  - Finite difference?, Finite volume? Finite element?
  - What type of element?
  - What order?
- And we need to think through the mathematics of coupling
  - Implicit or explicit
  - Time steps and errors
  - Data transfers



# Phase change and multiphase flow is another twist



We also get pressure build-up in enclosed volumes resulting in fracture and failure.

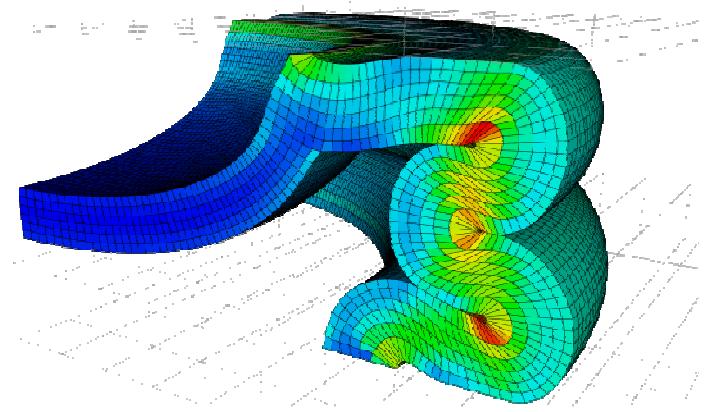


# Let's summarize the situation so far

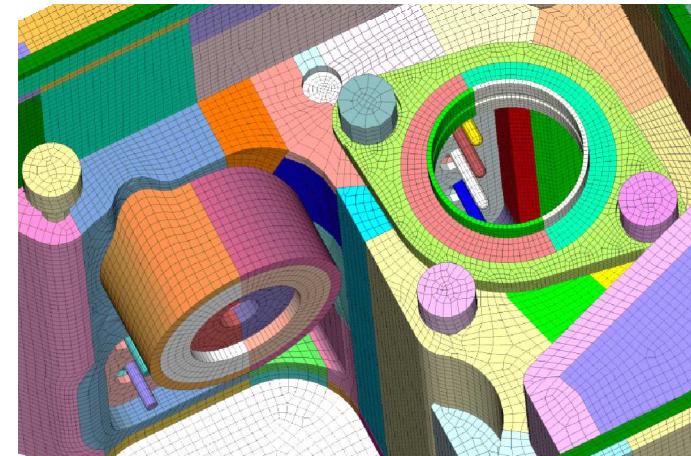
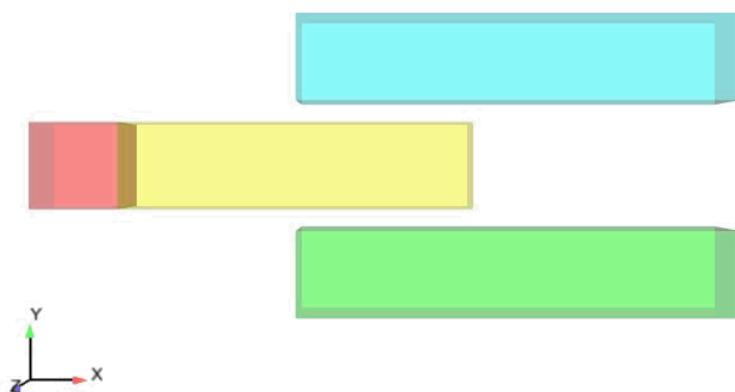
- We are
  - Simulating a system in a thermal environment
  - Coupling mechanical effects due to the thermal environment
  - Pushing the model to the point of system failure
- We are not
  - Modeling the environment itself (fire)
  - Subjecting the system to mechanical or electrical insult
  - Coupling aerodynamics or fluid flow effects
- Where do we need more mathematics or computer science?
  - The basic physics, discretizations and algorithms for the separate physics regimes at these scales is fairly well-known.
  - Need mathematics for coupling, contact, multiscale modeling, solvers, etc.

# One challenge is contact.

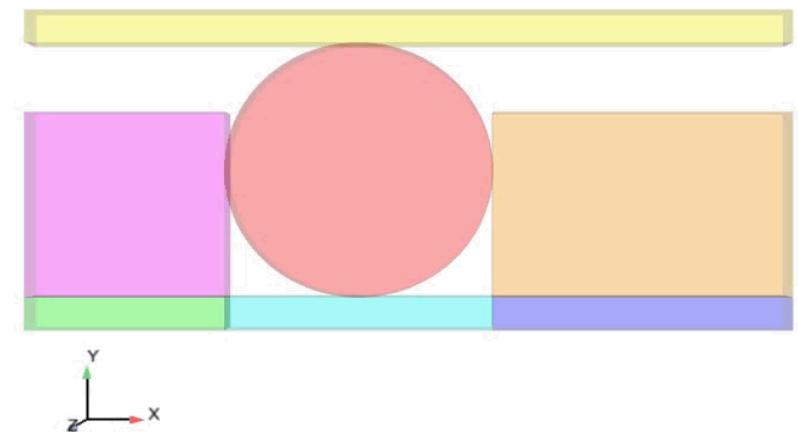
- There are several types of contact



Time = 0.000

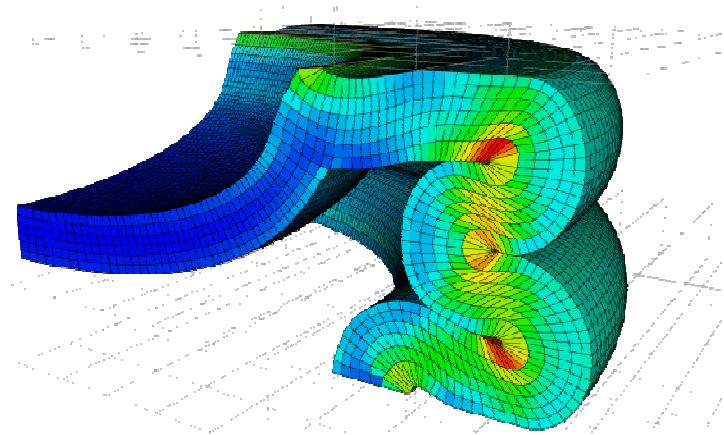
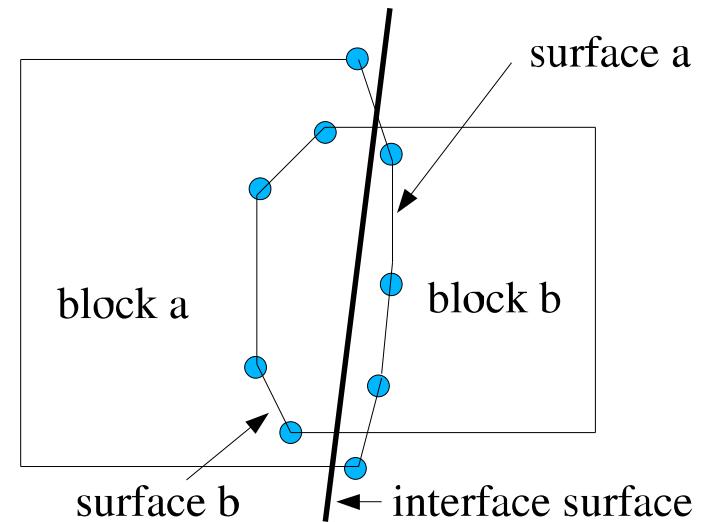


Time = 0.000

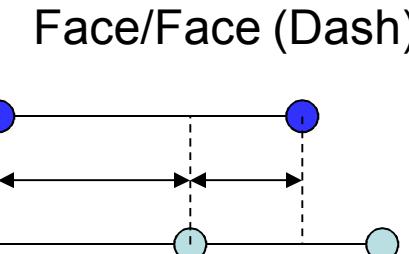
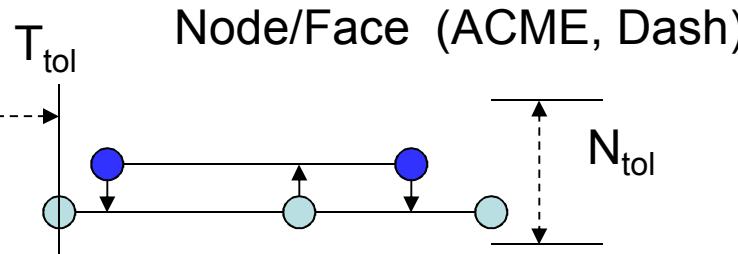


# One challenge is contact

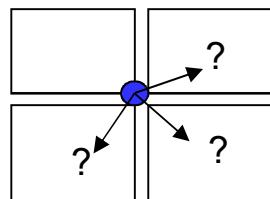
- Contact can be local or non-local
- Contact algorithms have different phases and differ in character from FP calculations
  - Block skinning: identify surfaces
  - Search:  $O(N \log N)$  but not local in memory
  - Enforcement: Iterative, augmented Lagrange algorithm for contact constraints
- Contact presents challenges for UQ and margin calculations
  - Implicit contact provides pre-loads (initial condition for the transient response)
  - Mesh sensitivities and discontinuities



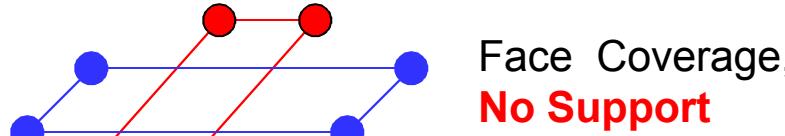
# And there are several algorithmic approaches to contact



## Ambiguous Cases:

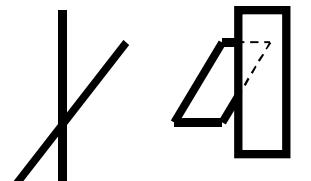


Corner Contacts, **Poor**



Face Coverage,  
**No Support**

Sharp Point  
Contact, **Good**

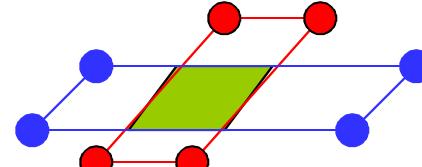


Beam and shell  
Edges, **No support**

## Ambiguous Cases:

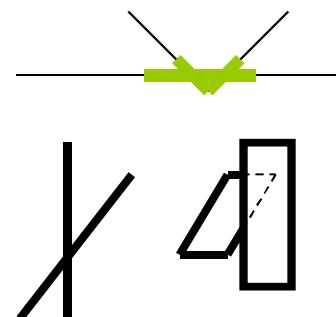


Corner Contacts,  
**Good**



Face  
Coverage, **Good**

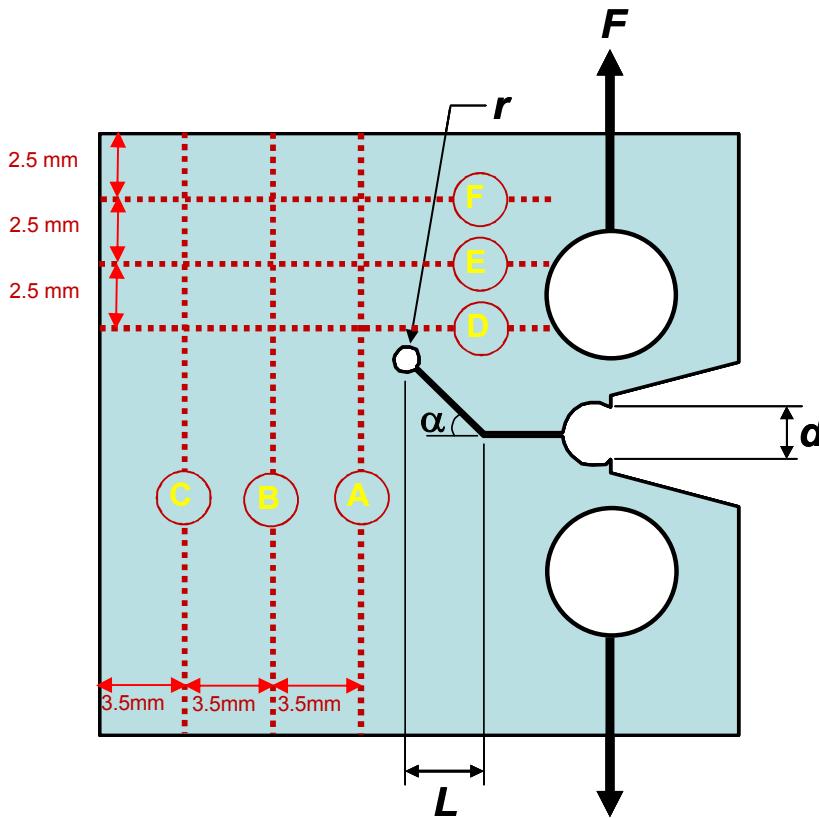
Sharp Point  
Contact, **Poor**



Beam and shell  
Edges, **OK**

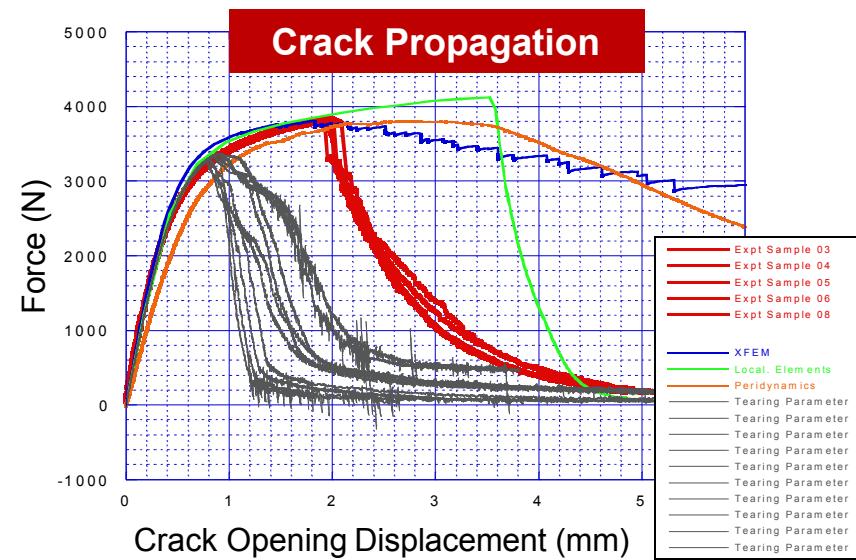
# The ability to model fracture and failure remains a research challenge

- The Sandia fracture challenge shows both how far we have come and how far we have to go

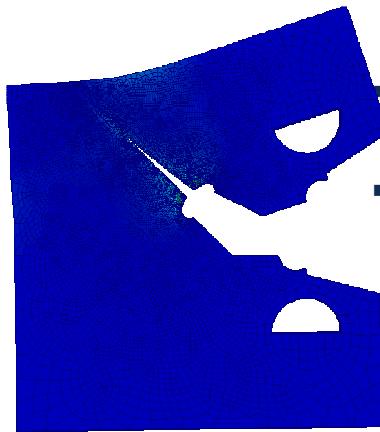


For the specimen shown below

- What is the *load-line displacement*  $\Delta d$  and the *peak force* prior to crack initiation?
- What is the *order of crack propagation* (e.g. A-B-D-C, etc.)?
- What is the *force* and *displacement* at which the crack reaches the *1<sup>st</sup> line*?
- What is the *force* (kN) and load-line *displacement* (mm) at which the crack reaches *line E*?

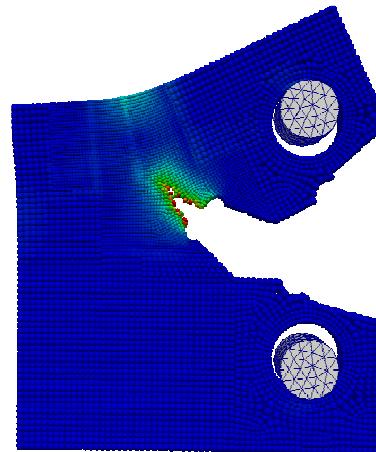


# Four mathematical approaches to fracture



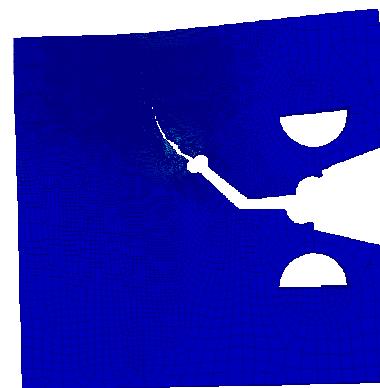
## Localization Elements

- Interface elements placed in mesh *a priori*
- 3 potential crack paths competed against each other to determine most likely path



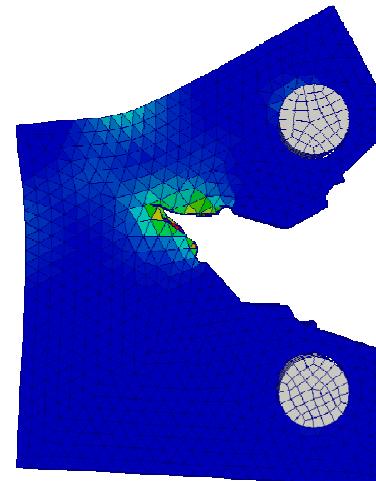
## Peridynamics

- Elastic-plastic continuum material used
- Bonds broken based on critical stretch
- Pins modeled with finite elements, interact with peridynamics via contact



## Element Death

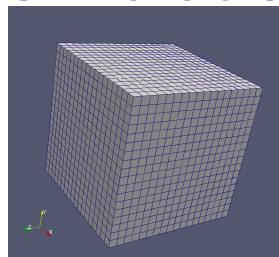
- Elastic-plastic material with triaxial-based tearing parameter to determine failure
- Elements killed when failure criterion reached
- Used implicit solver with limited number of failed elements per iteration



## XFEM

- XFEM in Sierra is still in development stage. This is one of first applications
- Elastic-plastic continuum material used, no cohesive zone law
- Coarse elements used for time step and because of crack tangling

# And of course, the traditional heart of any simulation is the solver.

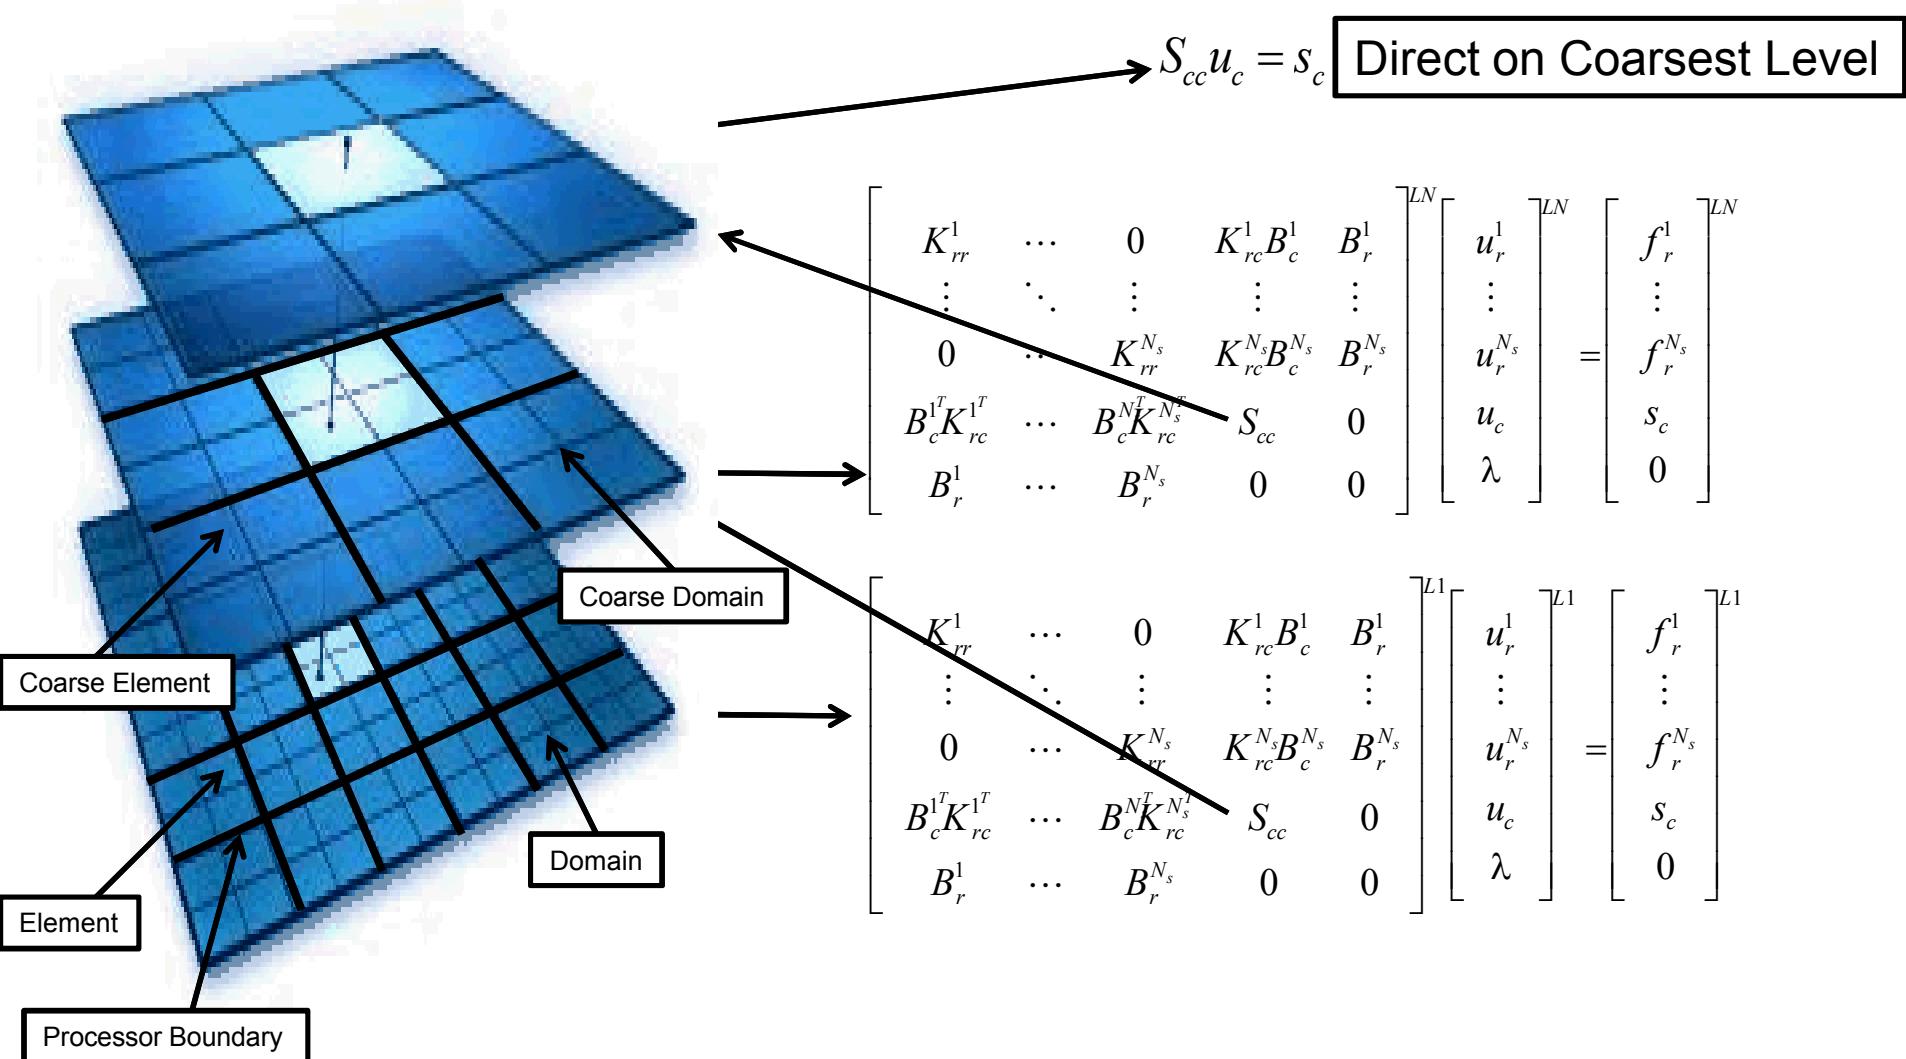


## Time to solve one linear system of equations: $Ax = b$

Cube Mesh	Dense	Skyline	Sparse	FETI
10x10x10	5 s	1 s	1 s	1 s
15x15x15	100 s	6 s	3 s	2 s
20x20x20	1300 s	40 s	10 s	6 s
25x25x25	DNF	180 s	30 s	12 s
30x30x30	-	-	93 s	25 s
<b>1000x1000x1000</b>	-	-	-	?

Linear solvers dominate nonlinear implicit solution time.  
We must have a scaleable technology to be able to take advantage of current and next generation machines

# Multilevel solvers represent a path to scalability

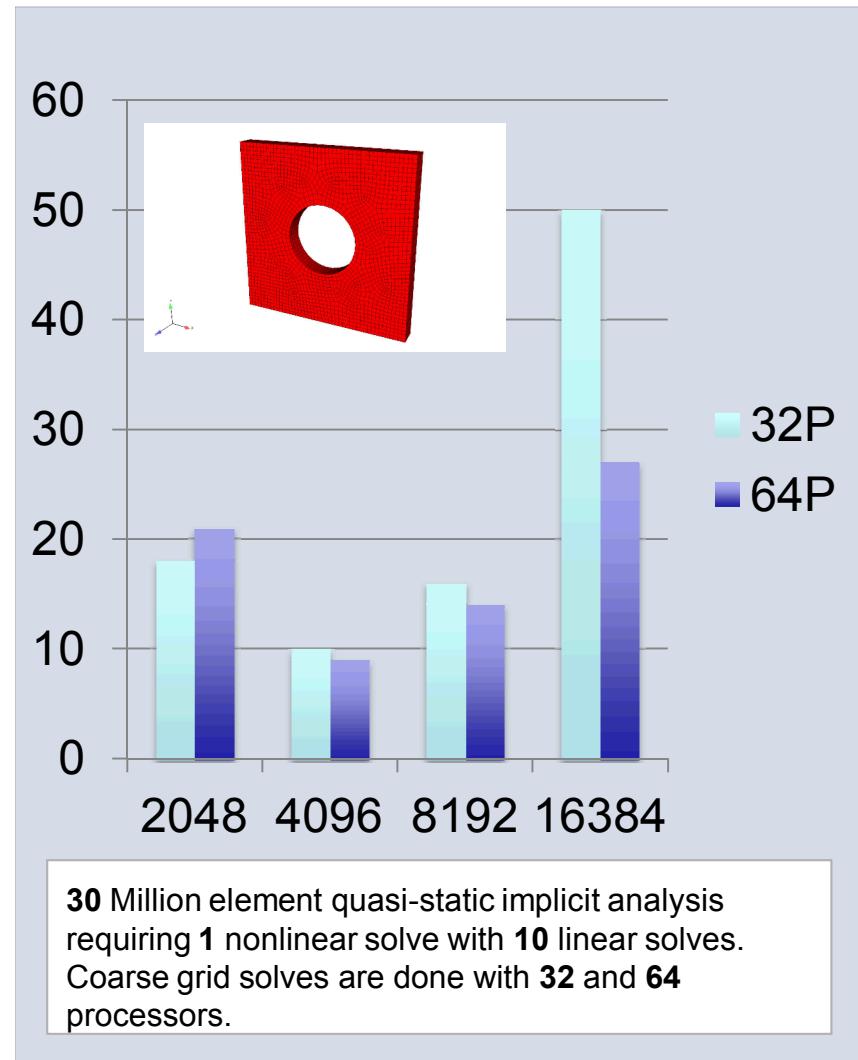


# The computing itself also remains a challenge

- The state of the art in HPC computing is currently explicitly-programmed domain-decomposition-based parallelism. This is about to change.
- Until now, mathematicians and engineers needed little more than a passing knowledge of computer architecture, including for example,
  - What is MPI and message passing?
  - Which algorithms scale logarithmically?
  - What data structures and algorithms facilitate vectorization?
  - What is a cache and what data structures and algorithms are cache friendly?
- Soon we will need to know
  - Hierarchical load balancing for heterogeneous computers
  - Resilient algorithms
  - Power management
  - More complex programming models
  - Asynchronous task-based parallelism
  - Etc.

# We have to consider a solver's algorithmic and implementation scalability. What size problems are we talking about?

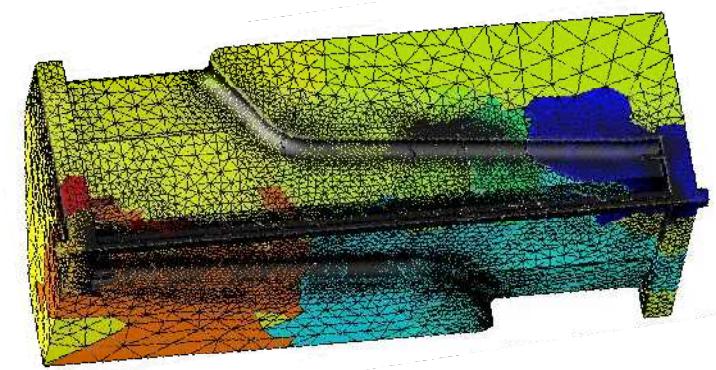
- Implicit Quasi-Statics Scalability
- Successfully ran on 16,384 processors with Implicit Nonlinear and Linear solvers, solving
- ~100 million equations
- 1 Nonlinear Solve, 10 Iterations.
- ~1 minute per linear solve.
- Scalability study identified areas of improvement in 3-Level FETI parallel iterative linear solver.
- Specifically need an automatic algorithm for coarse grid solves.



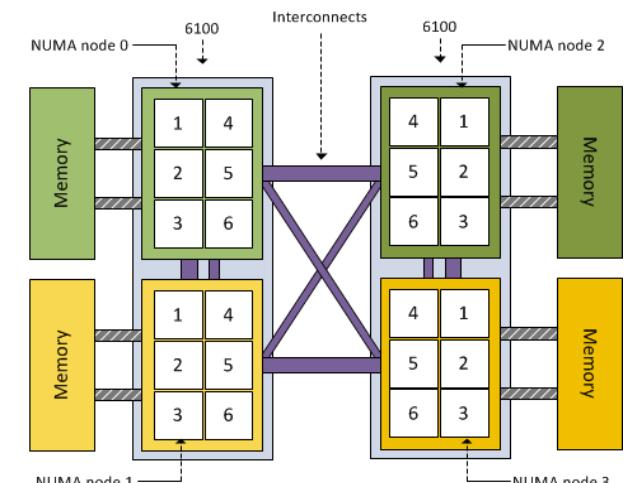
# Dynamic and hierarchical load balancing

- We need a much more sophisticated approach to workload and memory management. Example drivers for load balancing

- Hierarchical architectures and memories
- Adaptivity
- Localized multiscale methods
- Contact
- Power management



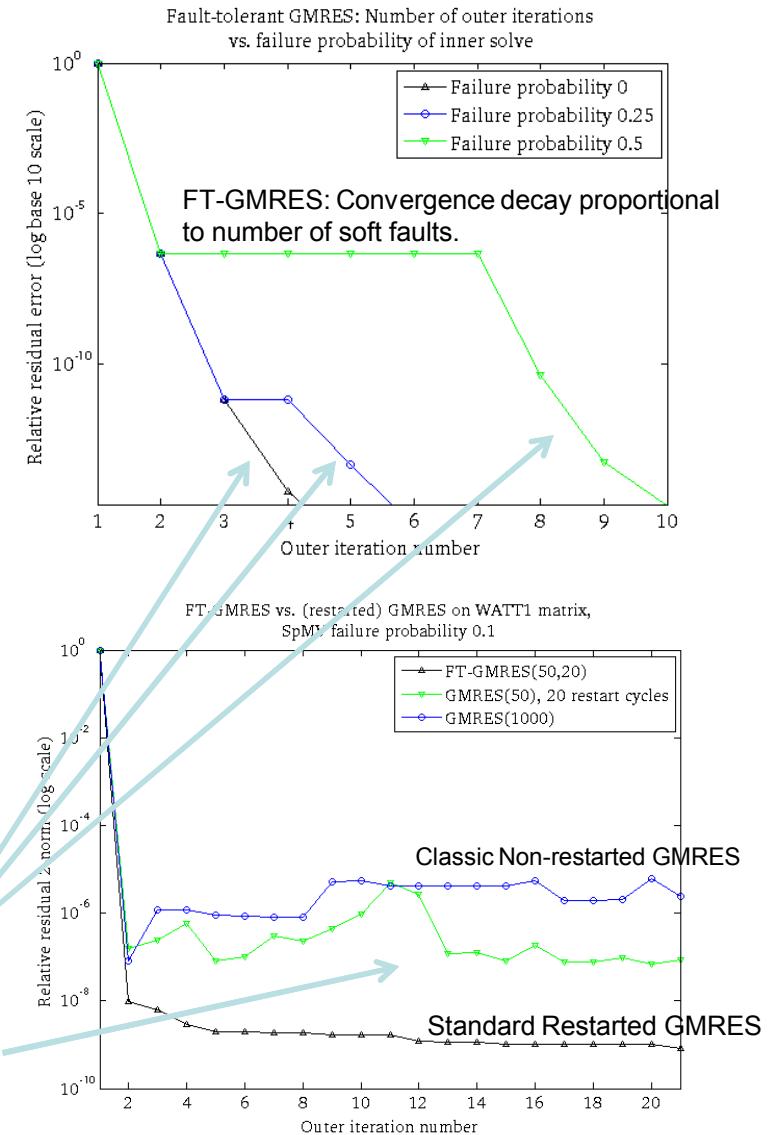
- We have implemented a “topology” parameter
  - TOPOLOGY=96 (96 cores)
  - TOPOLOGY=4,2,2,6 (4 nodes, 2 sockets, 2 dies, 6 cores)
- Topology-aware load balancing has resulted in a 60% speedup in one (simple) example



<http://frankdenneman.nl/2011/01/amd-magny-cours-and-esx/>

# Resilience and Fault Tolerance

- Exascale systems will be less reliable, and codes will be expected to take more responsibility
  - Exception handling/detect and fail
  - Checkpoint/restart
  - Redundant calculations
  - Asynchrony
  - Fault-tolerant algorithms
- Selective reliability enables new solvers
  - System exposes reliability tradeoffs
  - Algorithm identifies what *must* be reliable
- Fault-tolerant GMRES
  - Inner solver “preconditions” outer solver
  - Inner solver runs unreliably, outer solver reliably
  - Reuse any existing solver stack as “inner solver”
  - Most time spent in cheap unreliable mode
  - Faults only delay convergence; don’t prevent it
  - Standard restarting GMRES (simpler approach) is not sufficient.



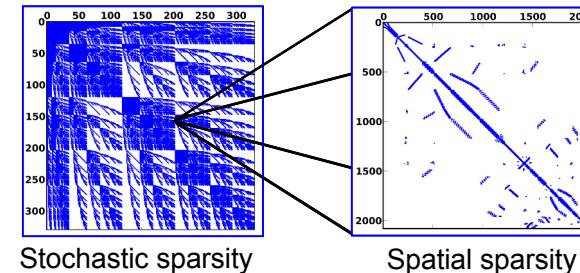
FT-GMRES

# Embedded UQ

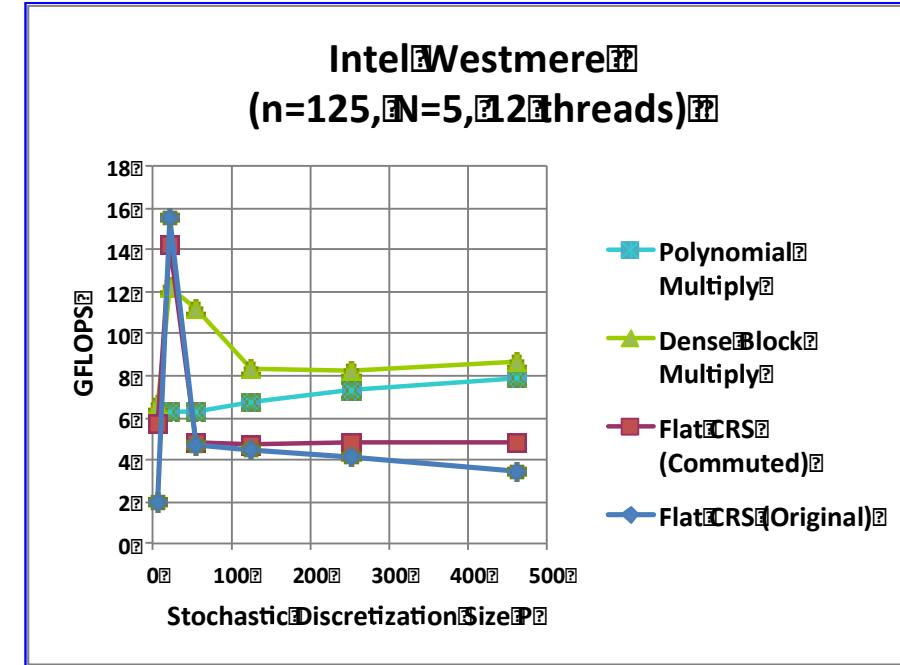
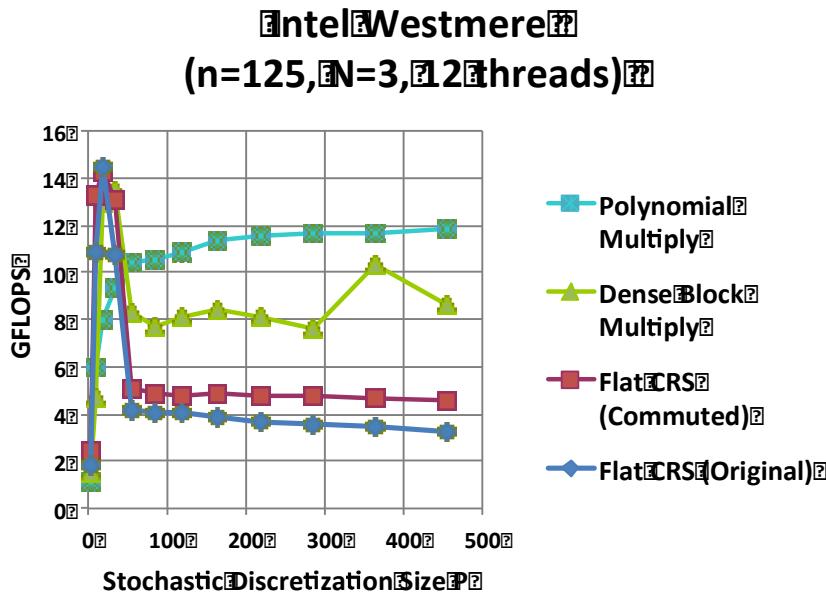
- Current methods are sample based and use capacity computing
- Embedding leads to
  - Better accuracy
  - Overall more efficient than ensembles
  - Reordering allows more parallelism and use of threads
- Can be coupled with traditional sample-based methods
- Steady-state stochastic problem (for simplicity):  
Find  $u(\xi)$  such that  $f(u, \xi) = 0$ ,  $\xi : \Omega \rightarrow \Gamma \subset R^M$ , density  $\rho$
- Stochastic Galerkin method (Ghanem and many, many others...):  

$$\hat{u}(\xi) = \sum_{i=0}^P u_i \psi_i(\xi) \rightarrow F_i(u_0, \dots, u_P) = \frac{1}{\langle \psi_i^2 \rangle} \int_{\Gamma} f(\hat{u}(y), y) \psi_i(y) \rho(y) dy = 0, \quad i = 0, \dots, P$$
- Method generates new coupled spatial-stochastic nonlinear problem (intrusive)

$$0 = F(\mathbf{U}) = \begin{bmatrix} F_0 \\ F_1 \\ \vdots \\ F_P \end{bmatrix}, \quad \mathbf{U} = \begin{bmatrix} u_0 \\ u_1 \\ \vdots \\ u_P \end{bmatrix} \quad \frac{\partial F}{\partial \mathbf{U}} :$$



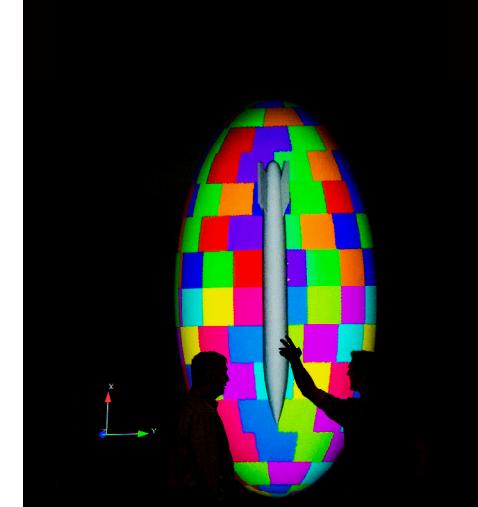
# The computational advantage of moving UQ calculation to the inner loop (GFLOPS – Intel Westmere)



- Standard 3-D first-order FEM grid (5x5x5)
  - Small FEM size due to limited GPU memory, large usage by block and CRS approaches
- N = polynomial order (larger N, denser blocks)

# And finally, there are a few potentially revolutionary advances

- Optimization and design
  - Asking a different set of questions
- Predictive simulations
  - Coupled scales
  - reduced-order models and surrogates
- Full life cycle (cradle-to-grave) engineering
  - Environmental specification for design
  - Design
  - Qualification and testing
  - Manufacturing and infrastructure
  - Surveillance and maintenance
  - Decommissioning



Long pause ----- deep breath

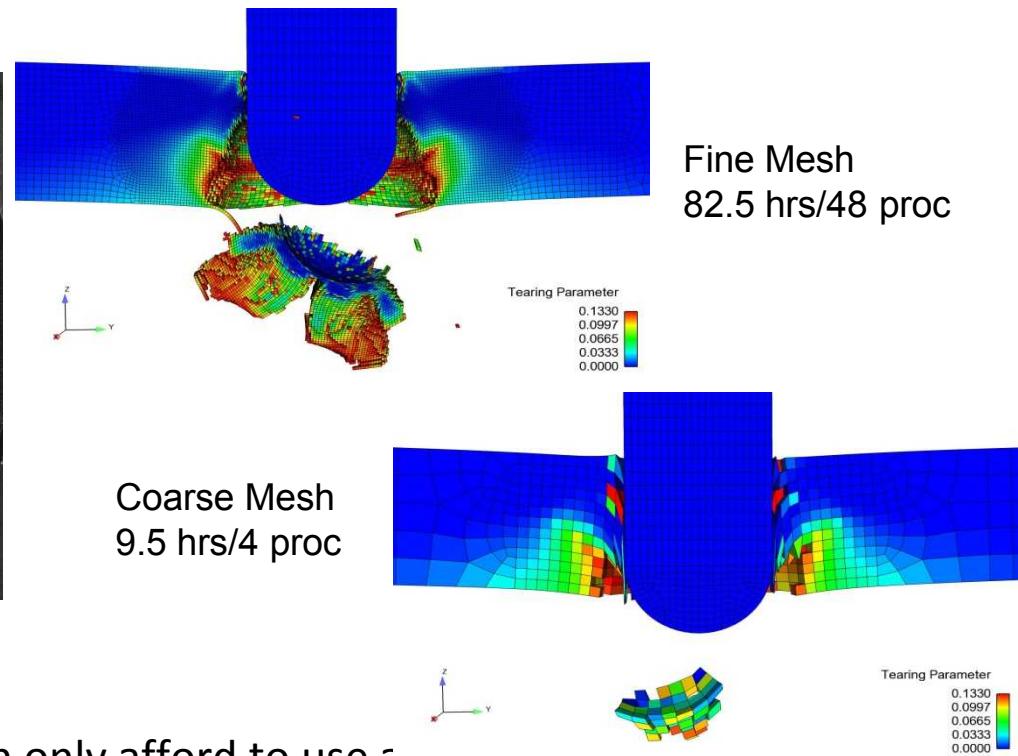
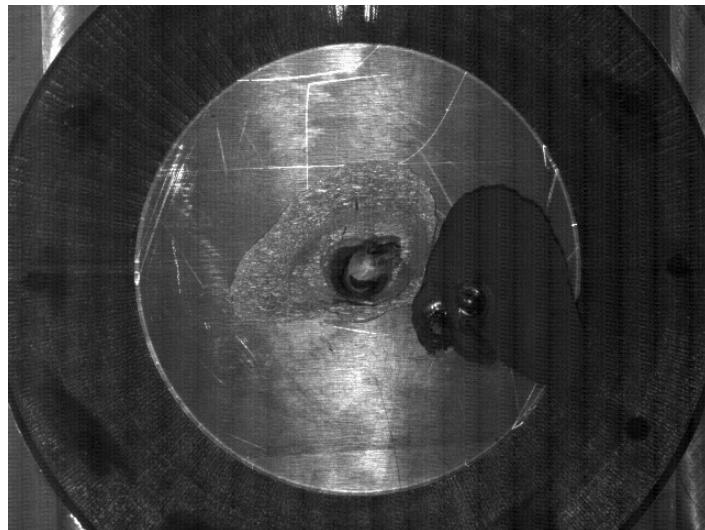
We know the physics and the algorithms. And we have a machine to run on. But is it correct?

# To answer that question is a long process

- Software verification and testing
- Solution verification
- Uncertainty quantification, sensitivities, error estimations and quantified margins
- Validation, including experiments and simulations

# V&V process and methods were tested in a recent L2 milestone.

- **Objective:** Assess predictions of the minimum penetration velocity using Sierra/SM w/tearing parameter as the failure criterion by comparing to test data.



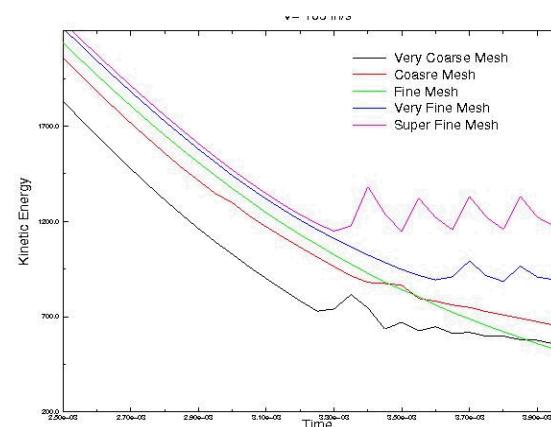
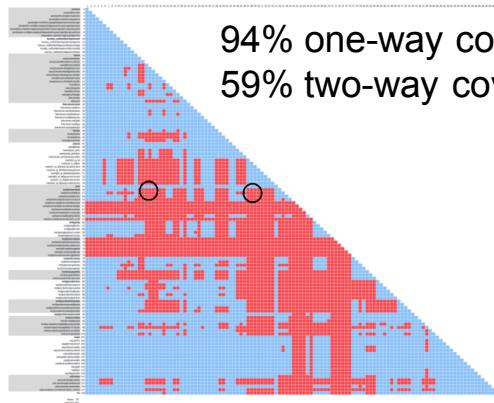
- **More questions:** Typically, we can only afford to use a ~~coarse mesh when many~~ function evaluations are involved! Are these models still valuable? Why do we believe in them? What error do they carry with them?

# The Milestone included verification activities.

## ■ Credibility Assessments:

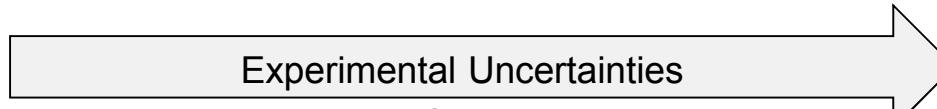
Phenomena	Consensus	Adequacy		
	Importance	Math Model	Sierra/S M Code	Validation
Large elastic-plastic deformation of metals	H	H	M	M
Ductile material failure	H	M	M	L
Contact	H	H	M	M
Friction between punch and test item	M	M	M	L
Enforcement of boundary conditions	L	H	H	L
Inertial loads	H	H	H	M

## ■ Code/Solution Verification:



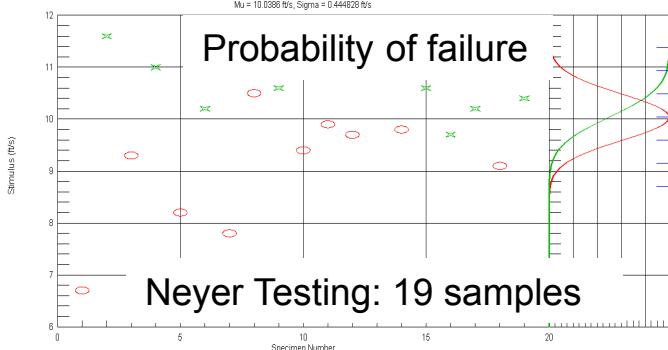
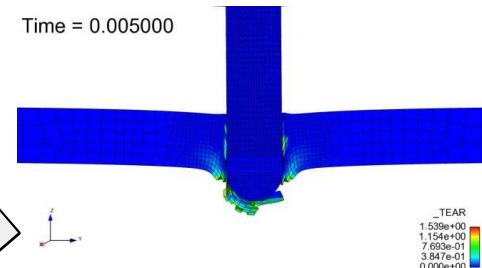
Evaluate mesh convergence rates before convergent behavior degrades

It also included validation activities.



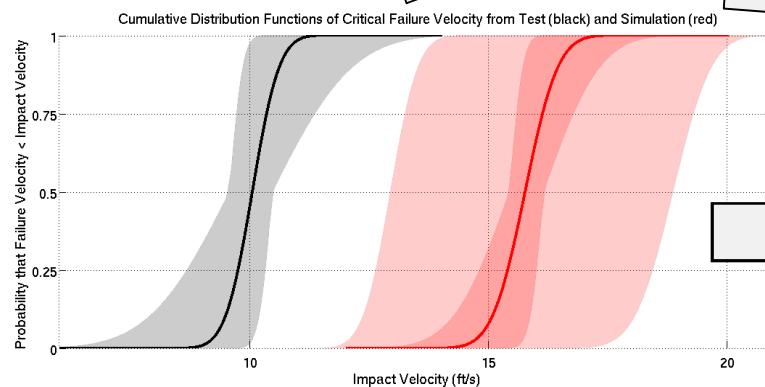
- Bungee force
- Friction (punch and plate/tube)
- Velocity measurement
- Material variability-> characterization process

Sensitivity analysis helped identify important factors → reduce scope



Uncertainties

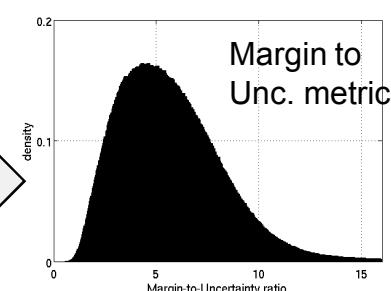
Uncertainties



Validation metrics

Numerical Uncertainties

- Discretization Error
- Algorithmic parameters



# Summary

- Computational simulation may not replace experiment, but it has fundamentally changed the way we approach engineering.
- It has the potential to revolutionize engineering
- Mathematics and computer science (whether taught in traditional Mathematics and C.S. departments or taught in engineering programs) are critical to the success of computational simulation.
- We have moved well beyond the physical regimes in which engineering intuition is sufficient.
- The complexity of the problem (physics, engineered systems, algorithms, code, architectures) required an integrated multidisciplinary approach