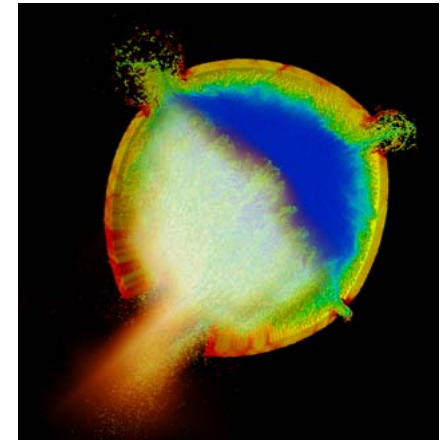


# Software for Adaptive Meshing on Large Scale Fluid-Structure Interaction Problems



**Martin Berzins**

- 1. Driving Problems and Technology**
- 2. Uintah Software and Modeling Details**
- 3. Scalable Adaptive Meshing Algorithms**
- 4. Detonation Results and Future Work**

Thanks to:

**Computer Science**      Justin Luitjens, Qingyu Meng, John Schmidt  
**Chemistry + Mech Eng.** Todd Harman, Joseph Peterson, Chuck Wight  
Steve Parker

DoE for funding the CSAFE project from 1997-2010, DOE NETL, INCITE

NSF for funding via SDCI and PetaApps (Abani Patra)

TACC NICS TRAC ORNL

<http://www.uintah.utah.edu>

# The path to exascale?



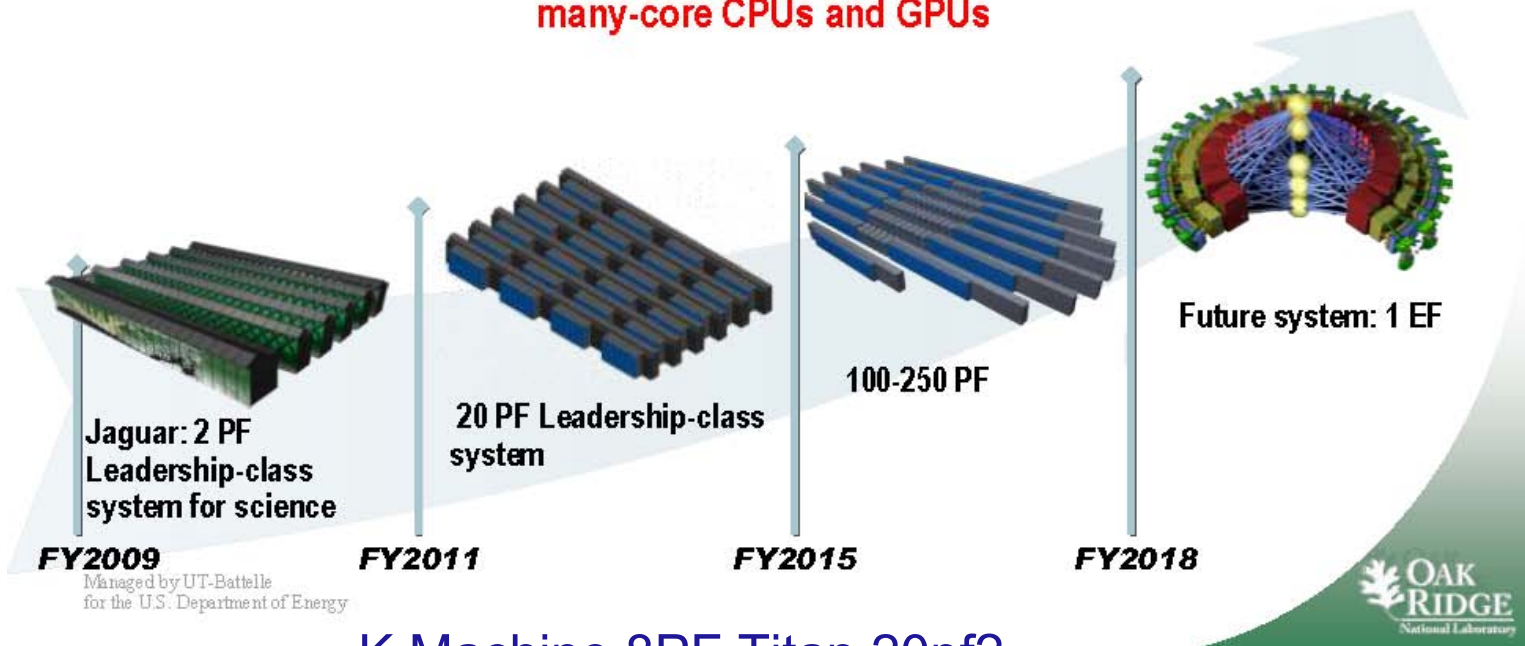
## Exascale Roadmap

**Delivering the next 1000x capability in a decade**

Mission need: Provide the computational resources required to tackle critical national problems

Must also provide the expertise and tools to enable science teams to productively utilize exascale systems

Expectation is that systems will be heterogeneous with nodes composed of many-core CPUs and GPUs



K Machine 8PF Titan 20pf?

Source Al. Geist

# DARPA Exascale Hardware Study

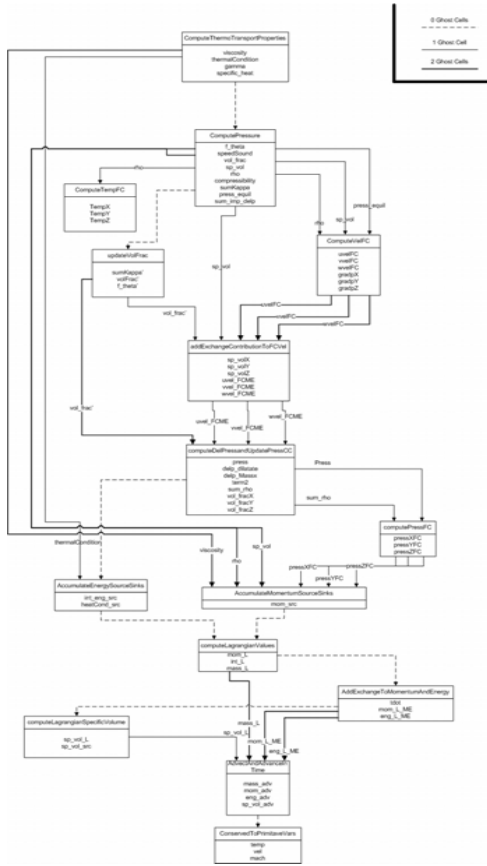
- DARPA public report (Peter Kogge et al.)
- Describes Challenges in going to Exascale at national level and petascale at University level.
- **Exascale machine Aggressive Strawman:**
  - 742 cores per socket, 12 sockets per node, 32 nodes per rack
  - 166,113,024 cores, 223,872 sockets
  - 4 flops per cycle per core @1.5Ghz, 1.029 PFlops
  - Power 67MW! **DoE aims for 25MW**
- **Extraordinary concurrency is the only game in town**
- **Power, fault tolerance, programmability are key**

**IMPLICATION IS PETASCALE AT LOCAL LEVEL – terascale laptops!**

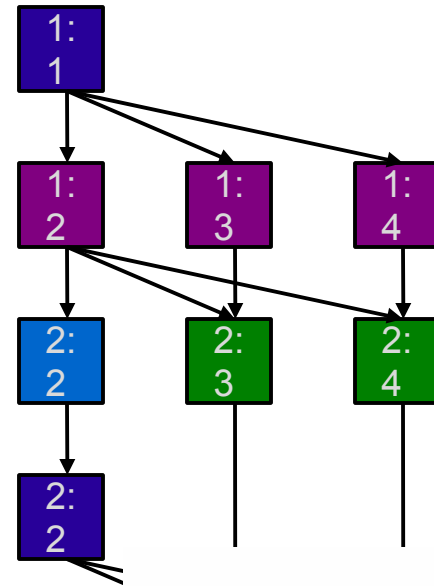
# DARPA Exascale Software Study

- **DARPA public report by (Vivek Sarkar et al.)**
- **(Sterling) Silver model for exascale software:**
  - Have abstraction for high degree of concurrency for directed dynamic graph structured calculations.
  - Enable latency hiding by overlapping computation and communications
  - Minimize synchronization and other overheads
  - Support adaptive resource scheduling
  - Unified approach to heterogeneous processing
- **Silver model is a graph-based asynchronous-task work queue model.**
- **Some instances of this type of approach in use now. CnC, Charm++, Plasma, StarSS, Uintah** **Very disruptive technology - forces us to rethink programming model**
- **DOES IT WORK?**

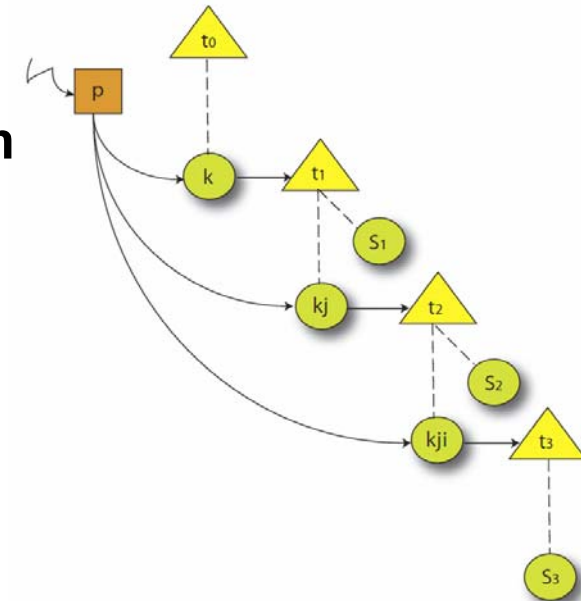
# Graph Based Languages/frameworks



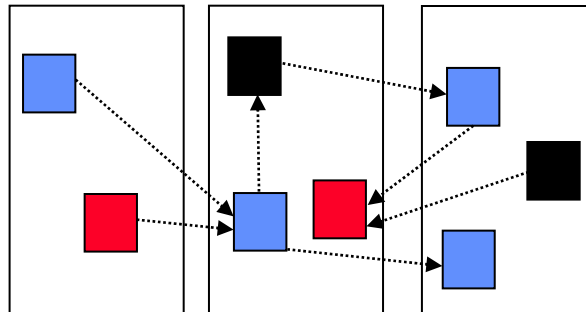
**Plasma (Dongarra):**  
**DAG based**  
**Parallel linear**  
**algebra**  
**software**



**Intel CnC:**  
**new language for**  
**graph based parallelism**



**Uintah Taskgraph**  
**based PDE Solver**



**Charm++: Object-based Virtualization**

# Weak and Strong Scalability

**Strong scalability**  $T(n, p) = \frac{T(n, 1)}{p}$

**Weak Scalability**  $T(np, p) = T(n, 1)$

Constant time for larger problem on more cores

Both weak and strong scalability only if  $T(n, 1) = \alpha n$

E.G. If  $T(n, 1) = \alpha n^k$  and  $k = 2$

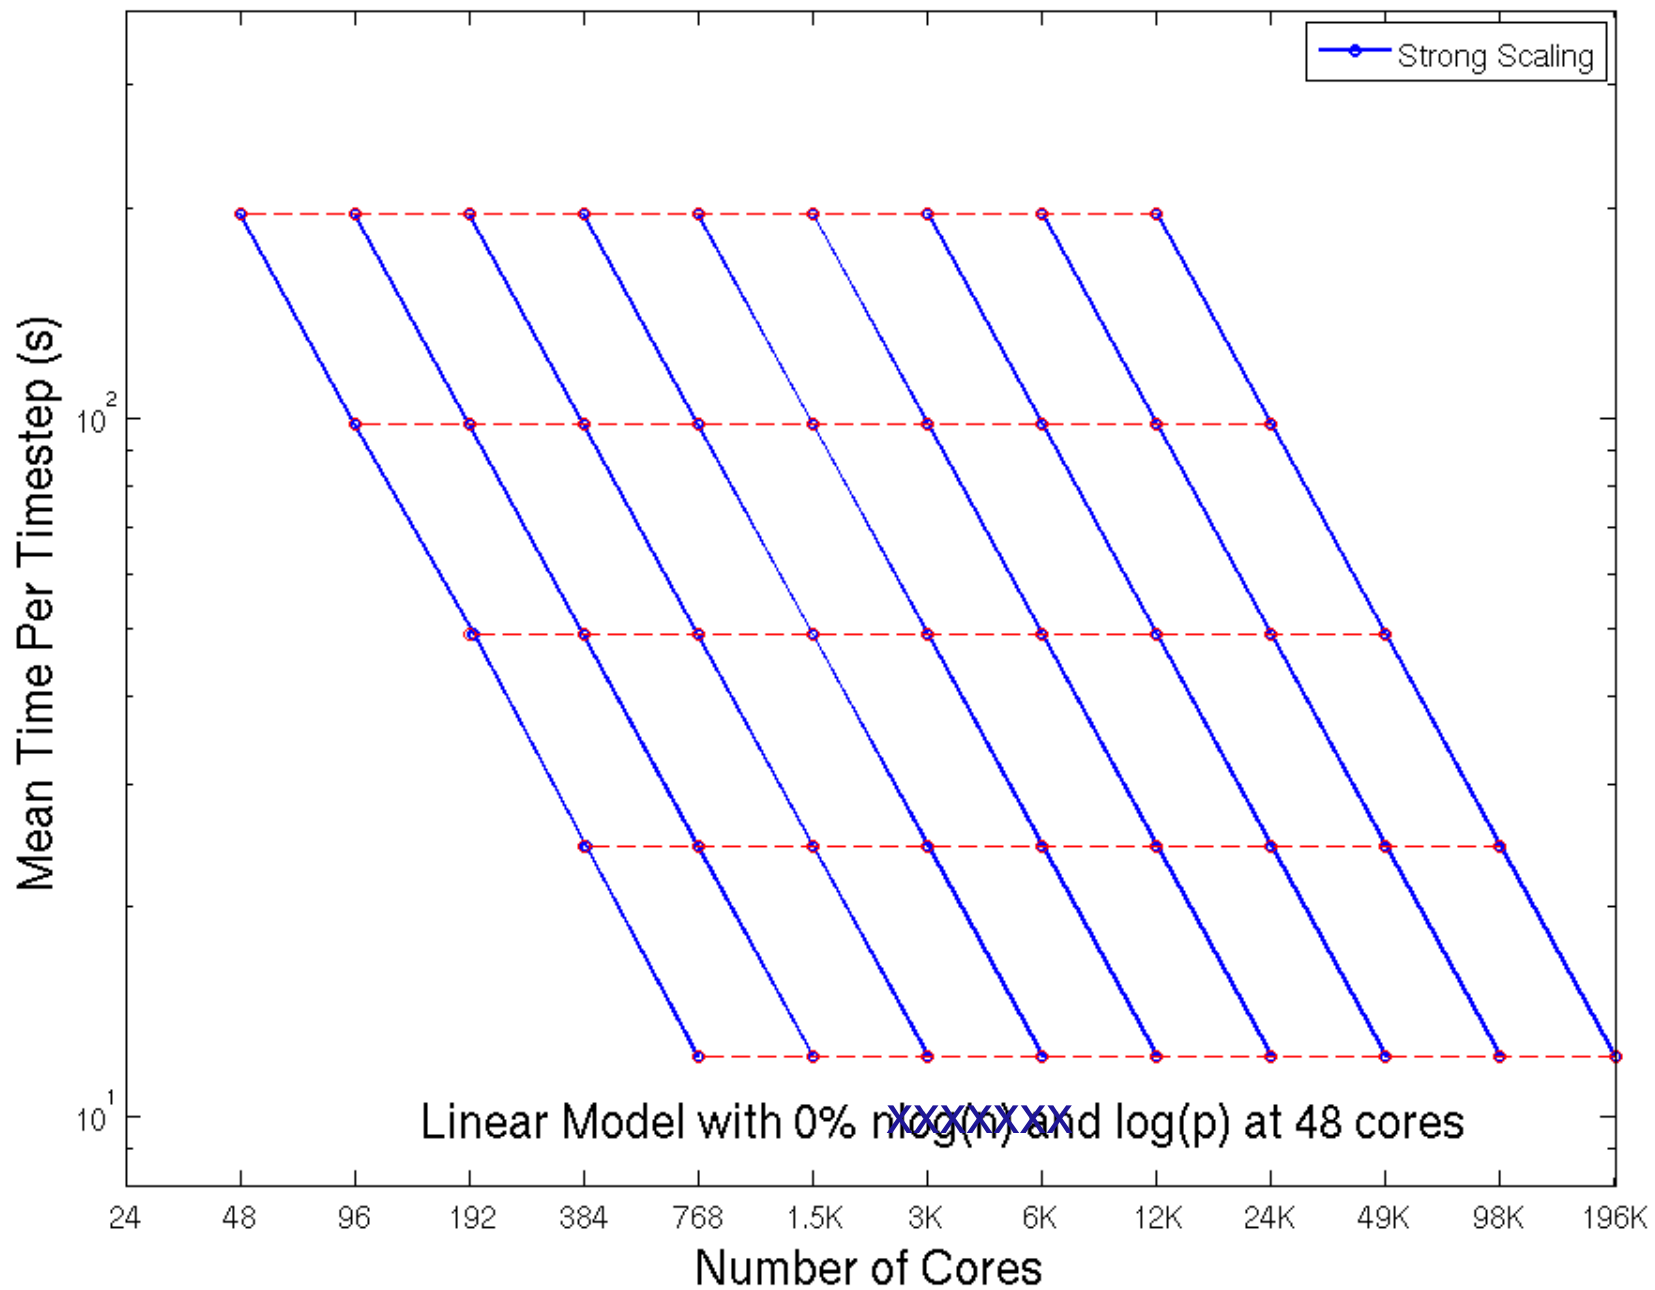
Then doubling the problem size gives four times as much work and hence twice runtime on 2x cores

## More realistic model

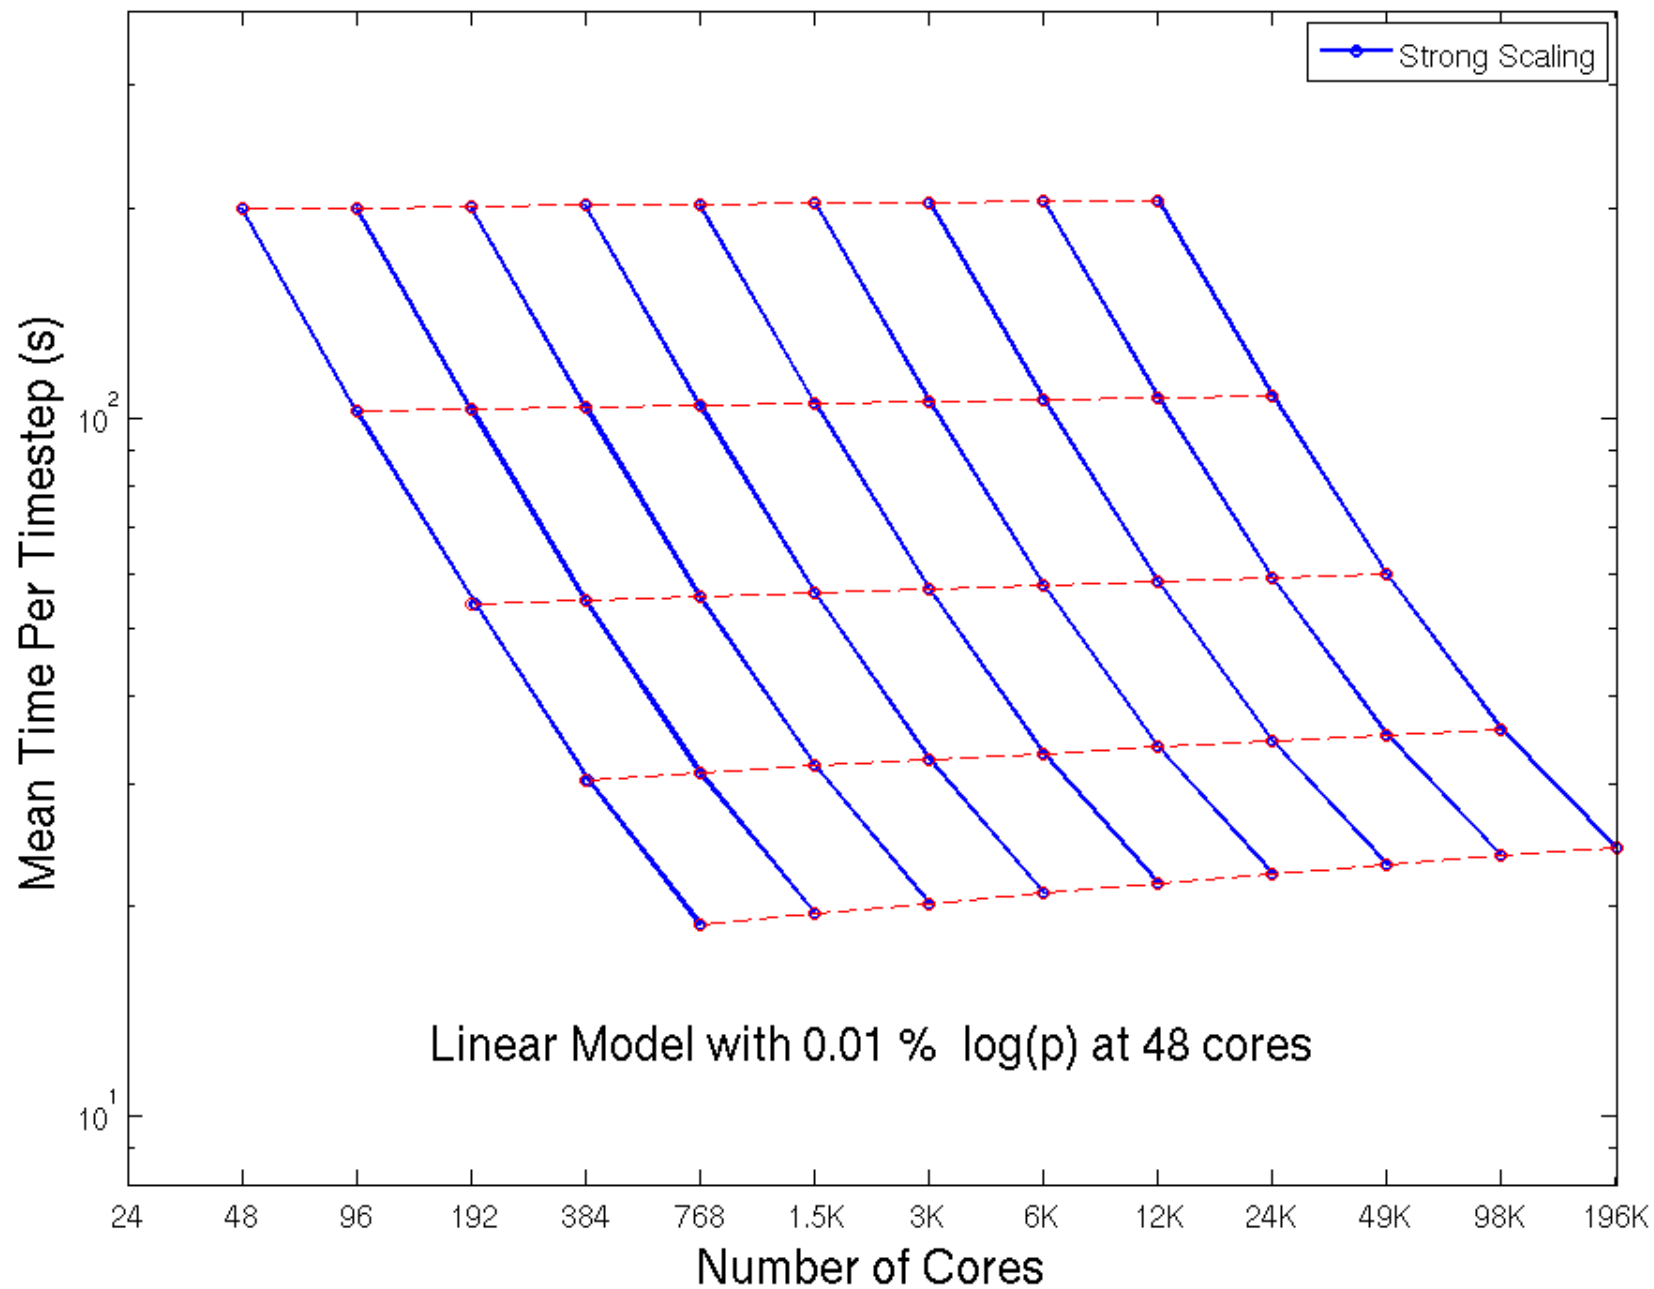
$$T(n, p) = \alpha n + \gamma \log(p)$$

$\log(p_0)\gamma / (\alpha n_0)$  Is fraction of time spent in global collectives at  $n_0 p_0$

# Single Level ICE Model: Scaling

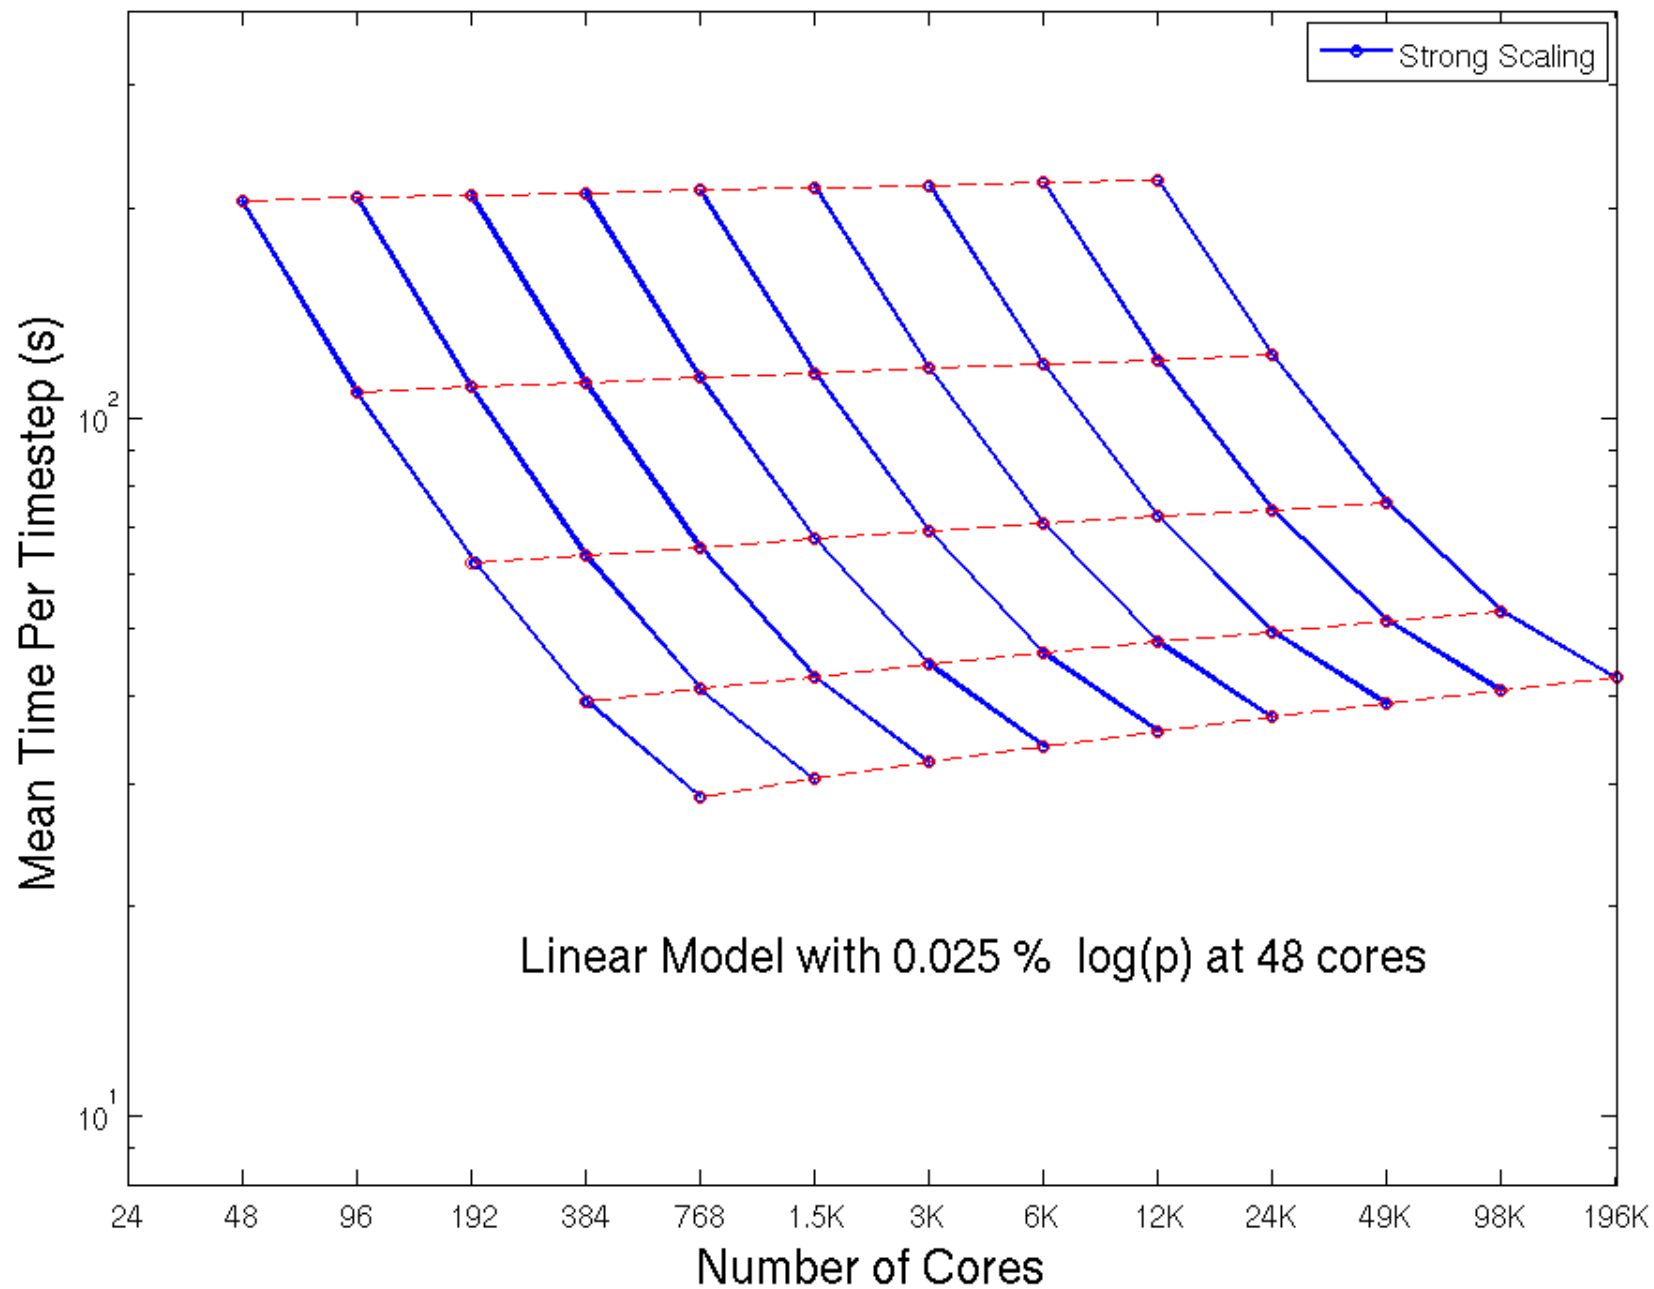


# Single Level ICE Model: Scaling

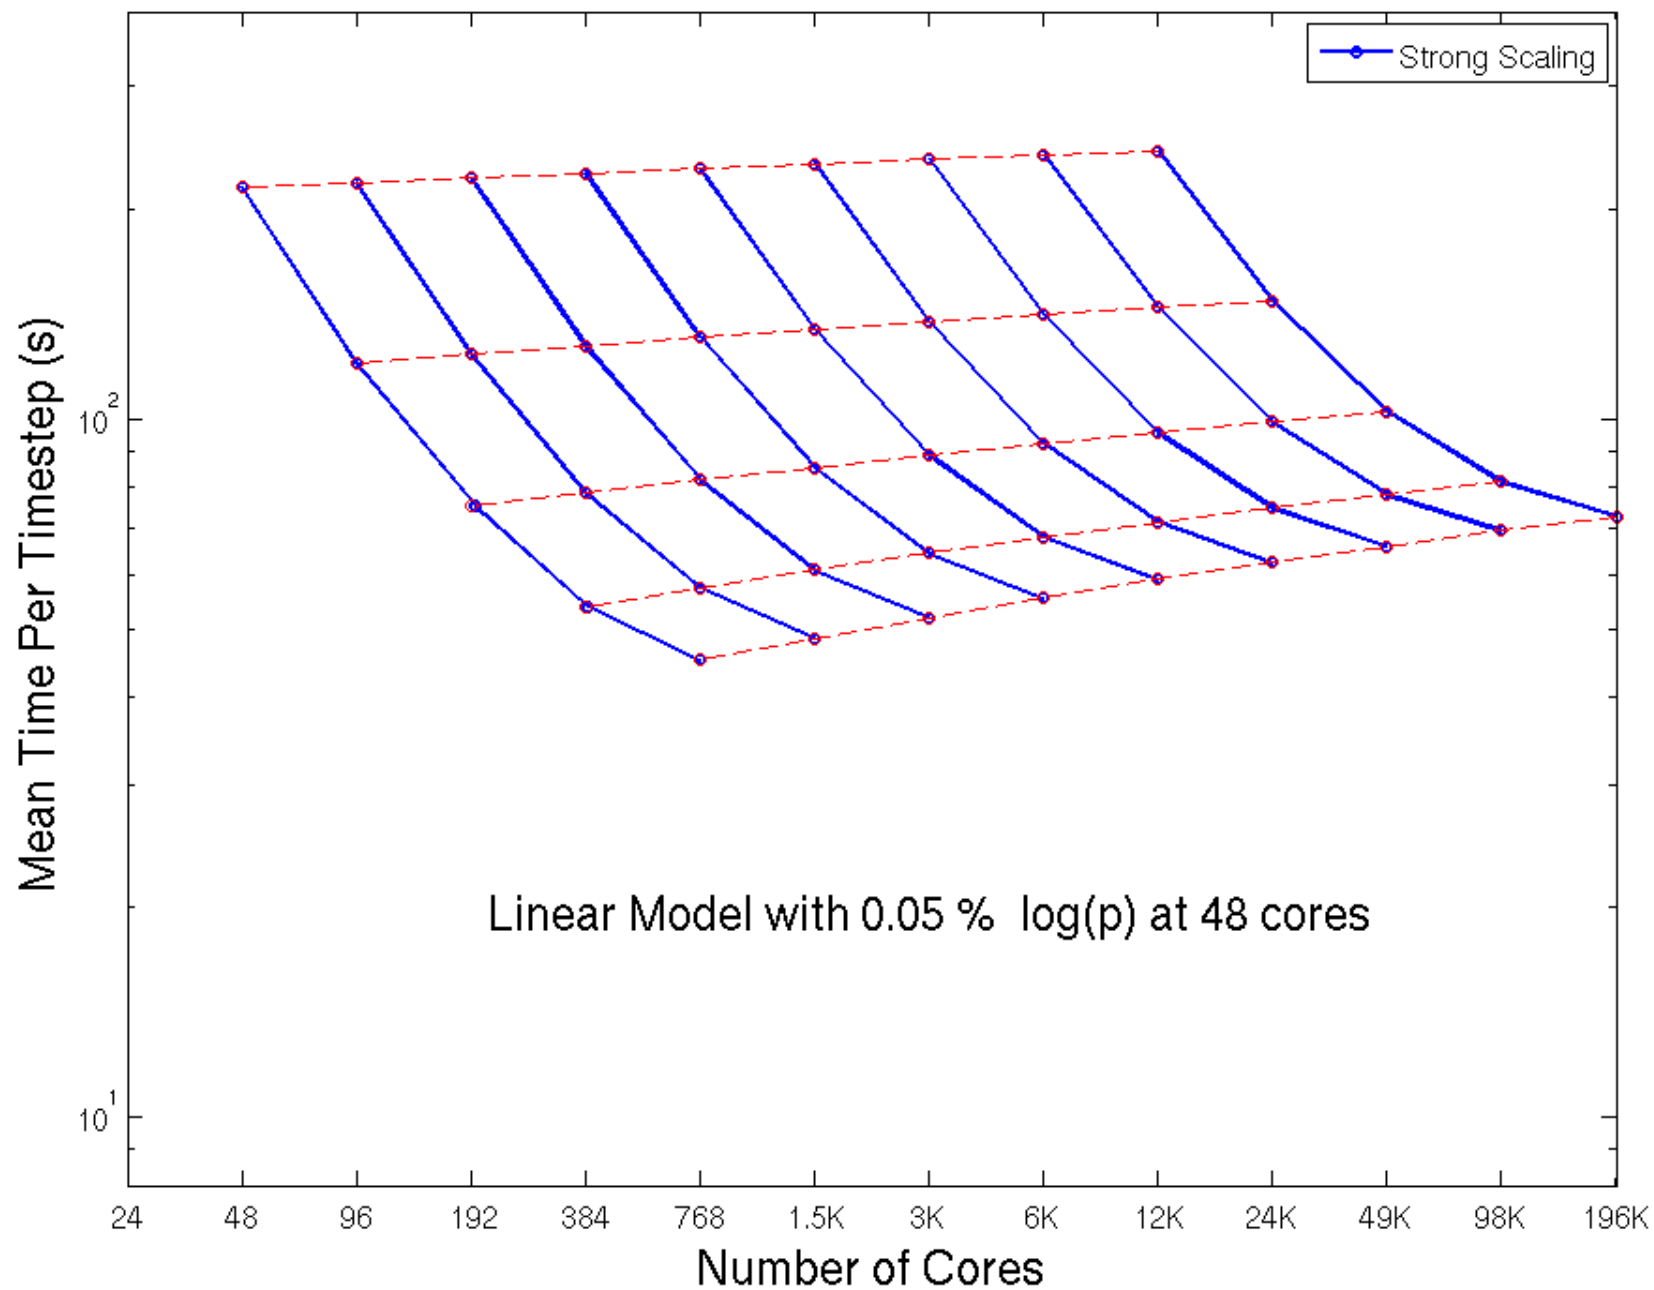




# Single Level ICE Model: Scaling



# Single Level ICE Model: Scaling



# Uintah Parallel Computing Framework

- Uintah (1998-2005) used DOE parallel computers, typical run – 512 to 2K cores far-sighted design by Steve Parker:

Solution of broad class of fluid-structure interaction problems

Patch-based AMR using particles and mesh-based fluid solver

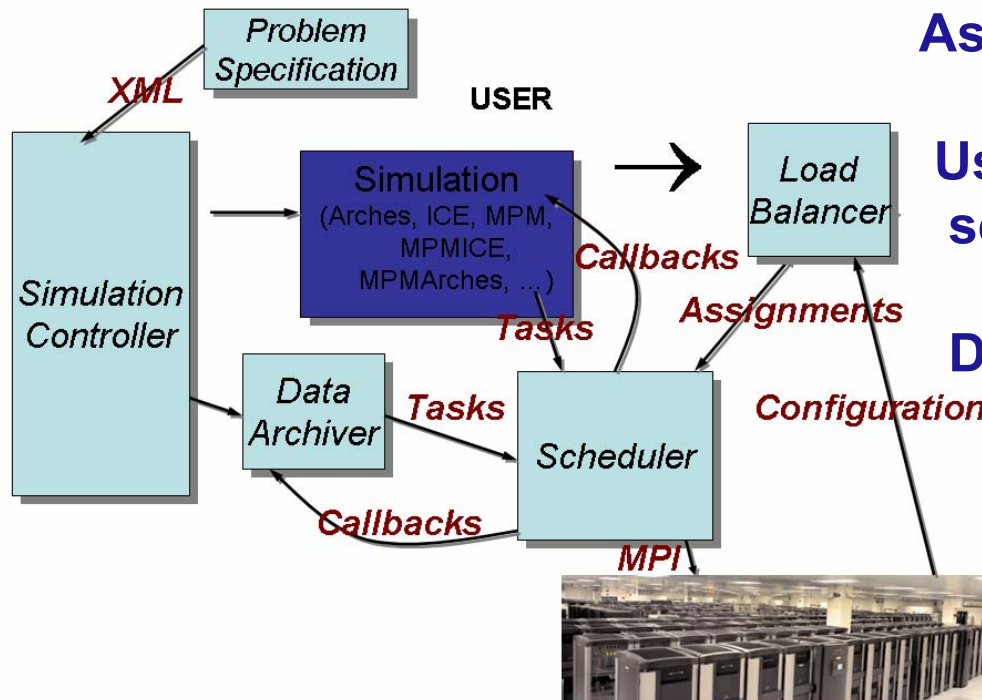
Automated task-graph generation for scheduling parallelism

Automated load balancing

Asynchronous communication

User only writes “serial” code for a serial mesh patch

Data is pulled from a data warehouse and results sent there



# Uintah Parallel Computing Framework

- Uintah had “legacy” code aspects –original design sound
- MUCH OF THE CODE HAS BEEN REWRITTEN
- New scalable AMR algorithm
- New measurement-based load balancer
- Dynamic execution (including out of order) of tasks
- New Hybrid MPI/Pthreads execution model
- Better use of Hypre solvers for time-dependent problems
- Much algorithmic development of discretisation methods
- 
- Uintah now uses NSF ( Ranger Kraken) DOE (Jaguar) computers, typical run – 2K to 196K cores
- How do we apply Uintah to model Developing Detonations?  
How do we start to think about scaling to beyond petascale



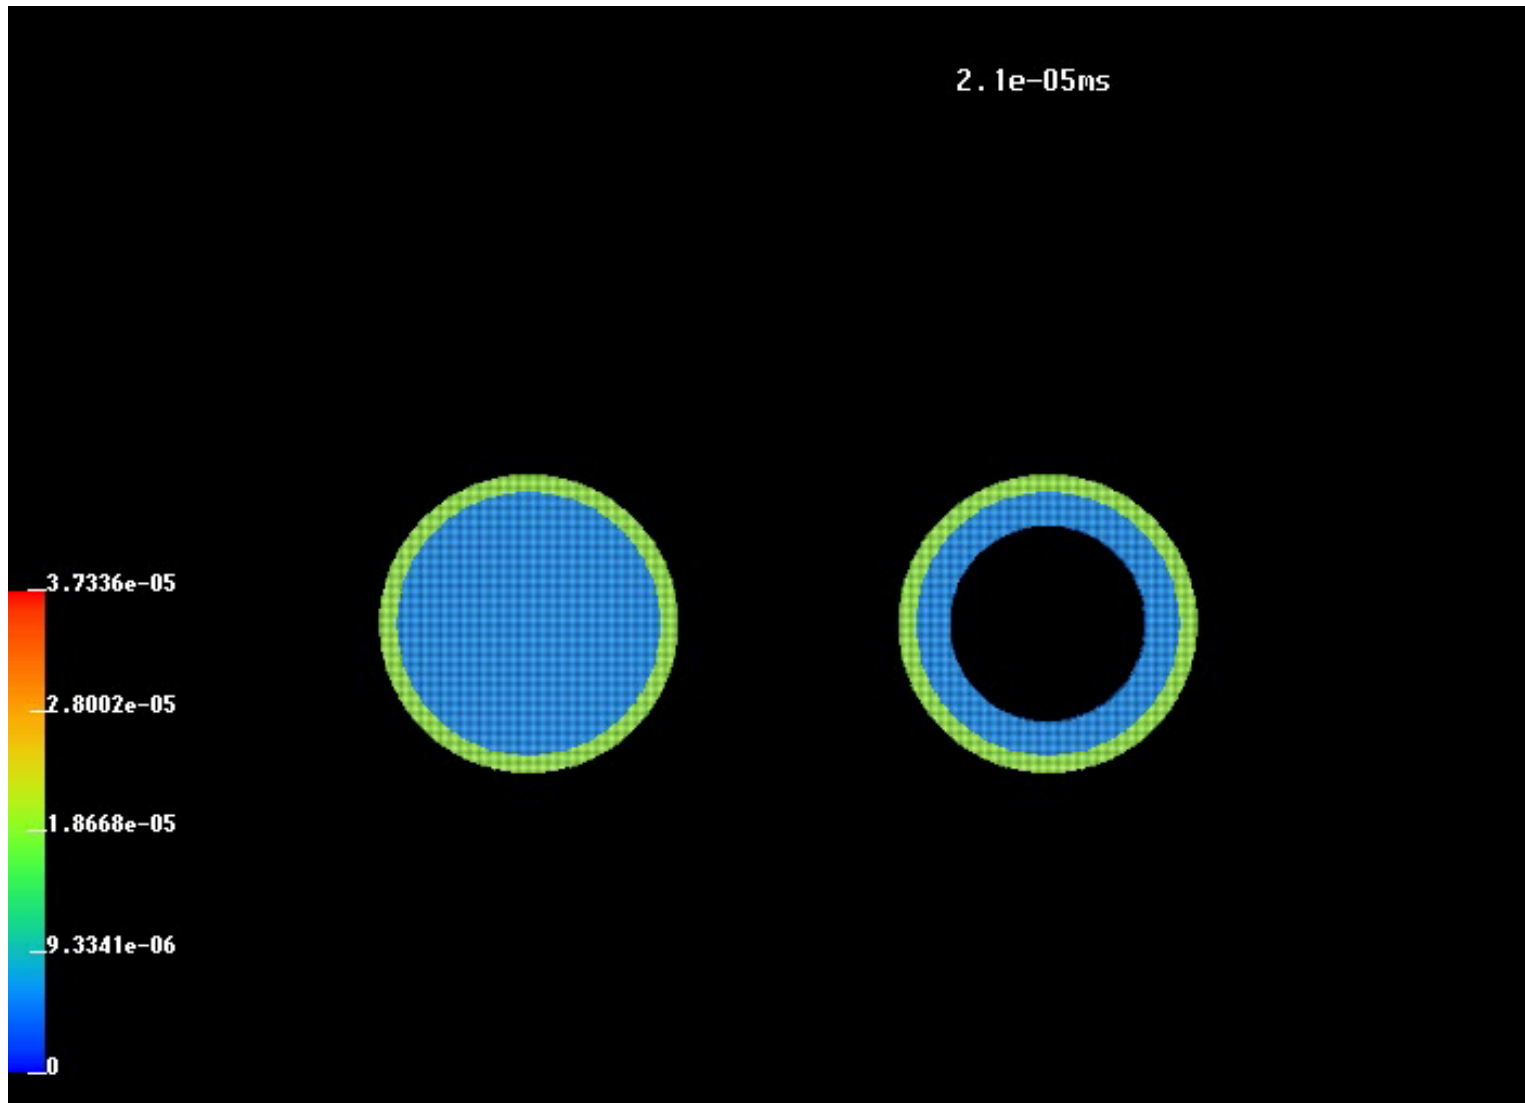
## Spanish Fork Accident 8/10/05

Speeding truck with 8000 explosive boosters each with 2.5-5.5 lbs of explosive  
Experimental evidence suggests that a transition from deflagration to detonation took place. Why?



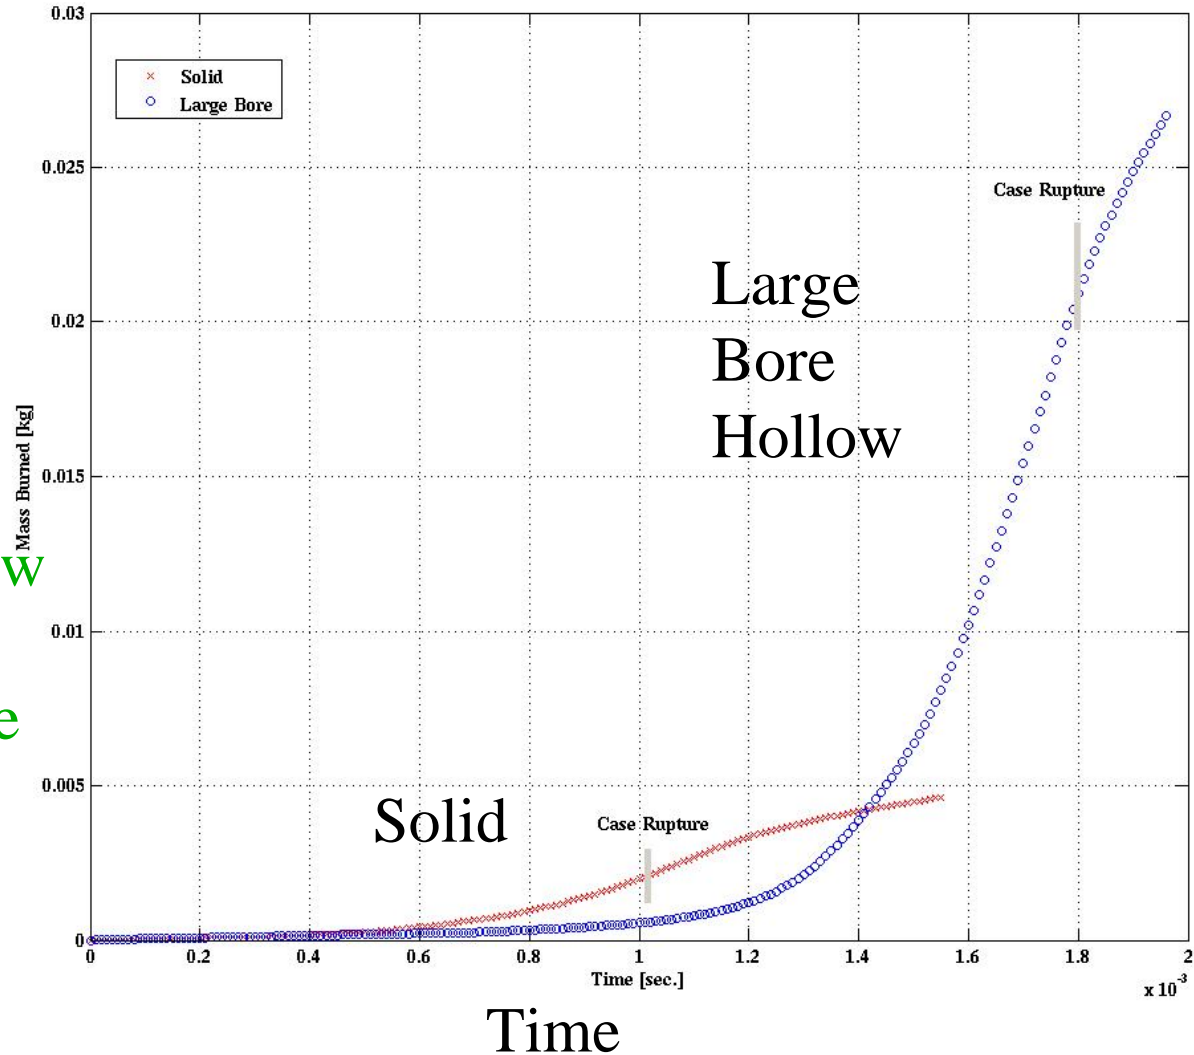
# Counterintuitive Dual Container Experiment

left – solid explosive, right- explosive with air has 4x energy release



# Explosive Mass Burned Comparison

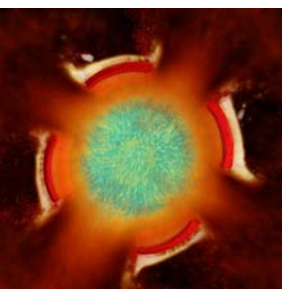
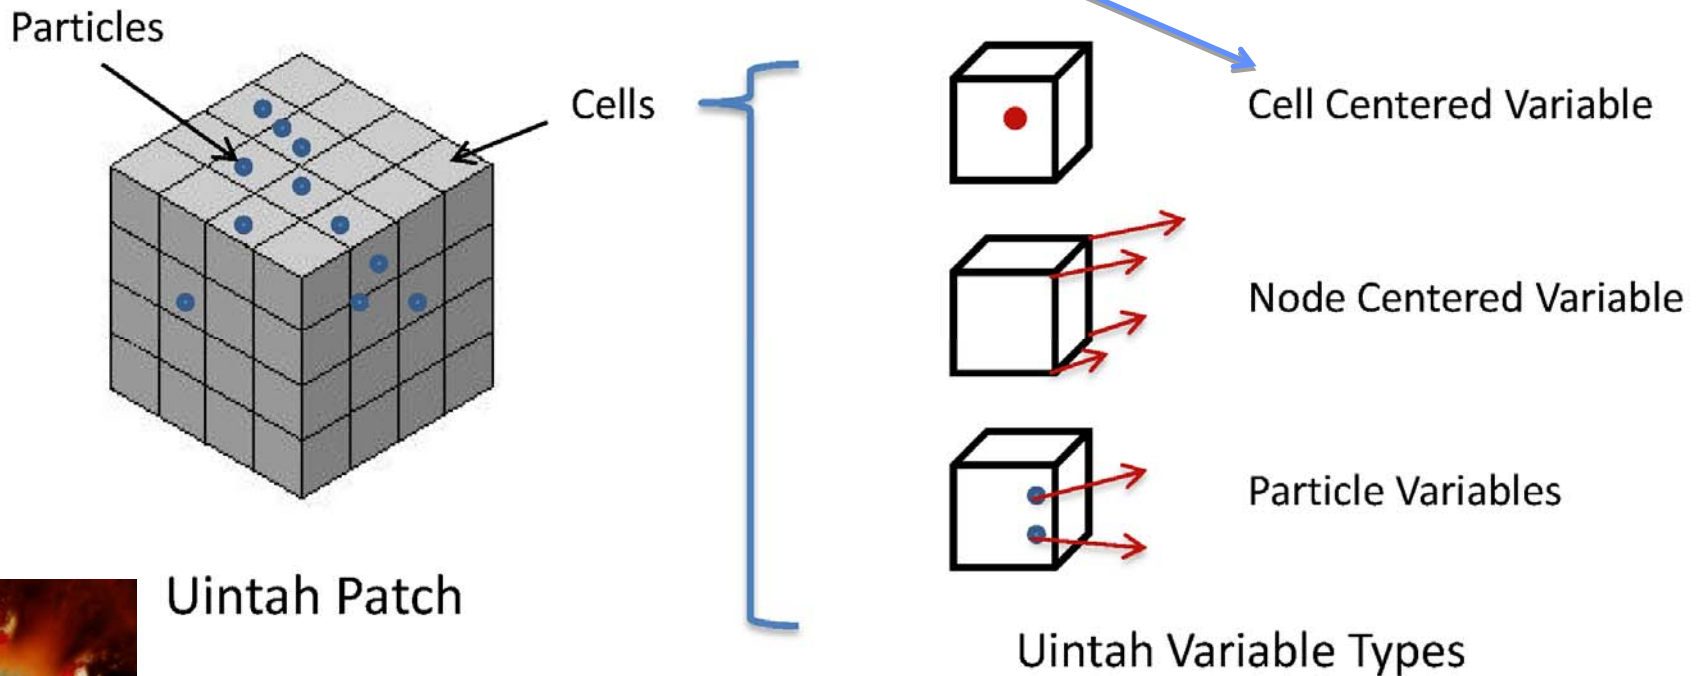
Mass  
Large Bore Hollow  
case burns 4.5 x  
Mass of solid case



# Uintah MPM-ICE-AMR Software

MPM (solids) and ICE (fluids) exchange data several times per timestep (not just boundary condition exchange)

**ICE is a cell-centered finite volume method for Navier Stokes equations**



**MPM is a novel method that uses particles and nodes Cartesian grid used as a common frame of reference**



Consider convection-diffusion type form

# ICE Algorithm

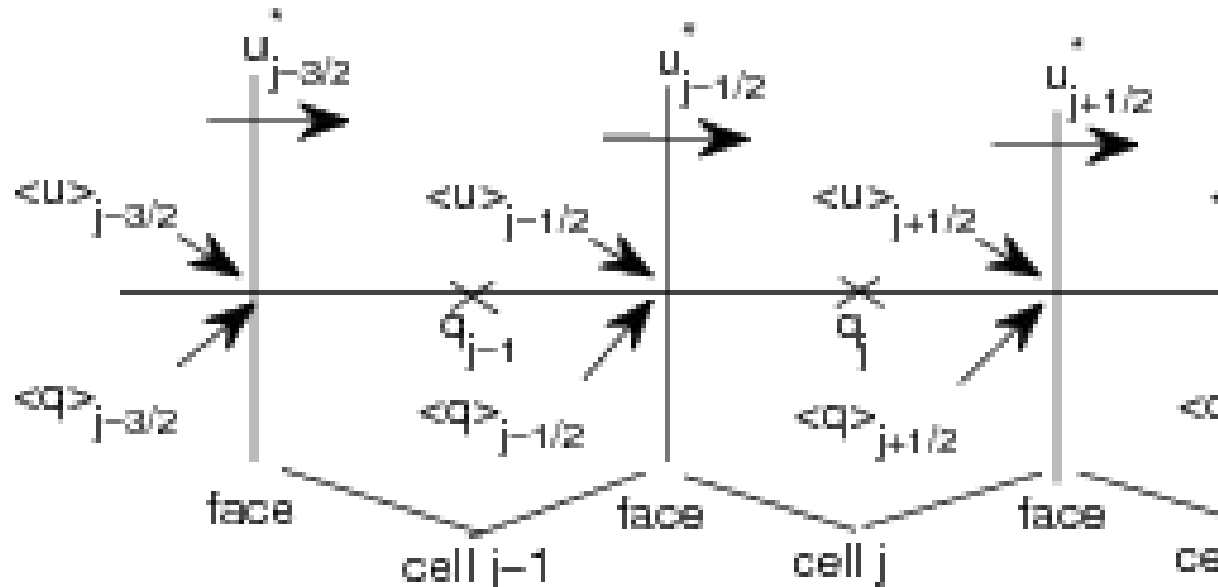
$$\frac{\partial q}{\partial t} + \frac{\partial(uq)}{\partial x} = - \frac{\partial f^p(q)}{\partial x}$$

**Fluxing velocity**  $u_{j+1/2}^*$  **at**  $t_n + \Delta t / 2, x_j + \Delta x / 2$

**Lagrangian**  $q_j^l V_j^l = q_j^n V_j^n - \Delta t (f^p(\langle q_{j+1/2}^{n+1/2} \rangle) - f^p(\langle q_{j-1/2}^{n+1/2} \rangle))$

**Eulerian**  $q_j^{n+1} V_j^{n+1} = (q_j^l V_j^l) - \Delta t (\langle q \rangle_{j+1/2}^n u_{j+1/2}^* - \langle q \rangle_{j-1/2}^n u_{j-1/2}^*)$

**Apply limiters at a number of stages To get a positivity preserving algorithm**



Original Algorithm for face value at  $t_n$

# ICE Algorithm

$$\langle u \rangle_{j+1/2} = \frac{\rho_j u_j + \rho_{j+1} u_{j+1}}{\rho_j + \rho_{j+1}}$$

Left  
Value

$$u_{j+1/2}^L = u_j + \frac{1}{2} \Phi(r_j)(u_j - u_{j-1}), r_j = \frac{u_{j+1} - u_j}{u_j - u_{j-1}}$$

Right  
Value

$$u_{j+1/2}^R = u_{j+1} - \frac{1}{2} \Phi(r_{j+1})(u_{j+2} - u_{j+1}), r_{j+1} = \frac{u_{j+1} - u_j}{u_{j+2} - u_{j+1}}$$

If  $u_{j+1/2}^L \rho_{j+1/2}^l + u_{j+1/2}^R \rho_{j+1/2}^R > 0$

then

$$\langle u \rangle_{j+1/2} = u_{j+1/2}^L$$

else

$$\langle u \rangle_{j+1/2} = u_{j+1/2}^R$$

$\Phi(r)$  is a standard limiter as used for hyperbolic eqns

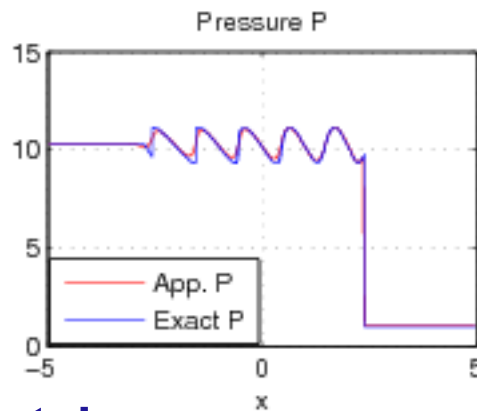
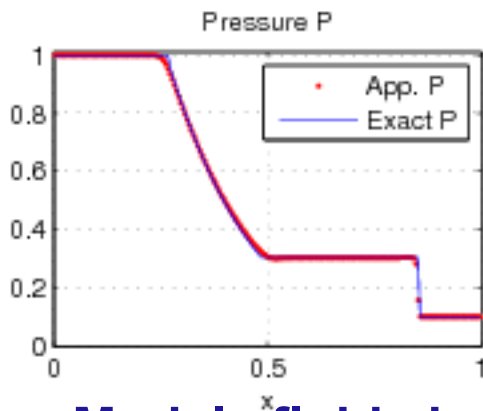
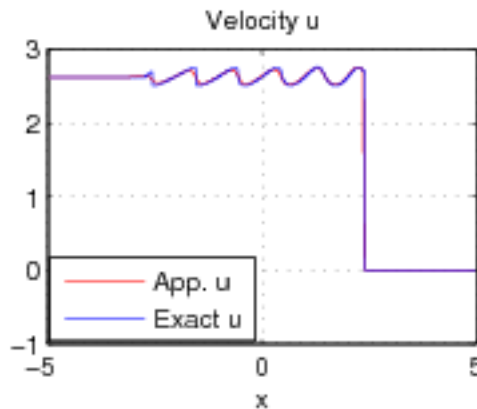
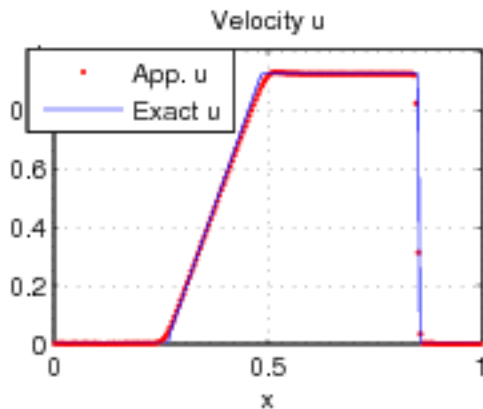
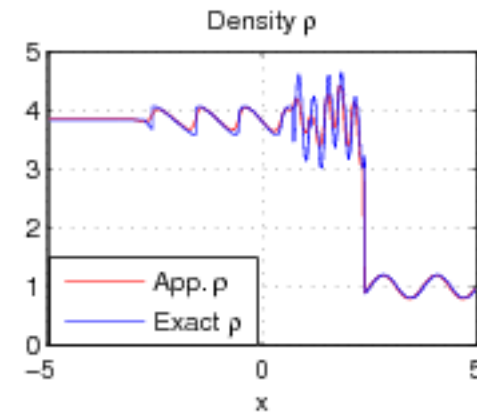
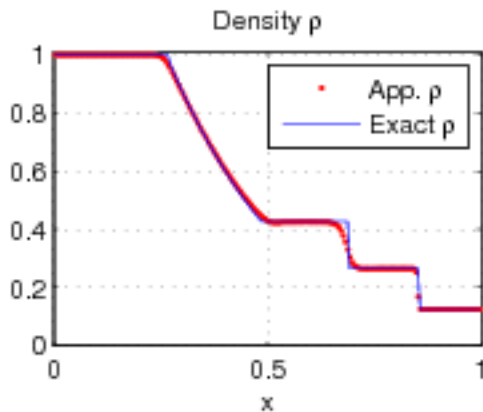
# ICE Navier Stokes Algorithm of Kashwa et al. Improved 08/09 for High Speed Flow Tran and Berzins

**Euler Equations  
Elimination of oscillations  
and second order version**

**Examples:**

**Left : Sod Shock Tube with  
200 points**

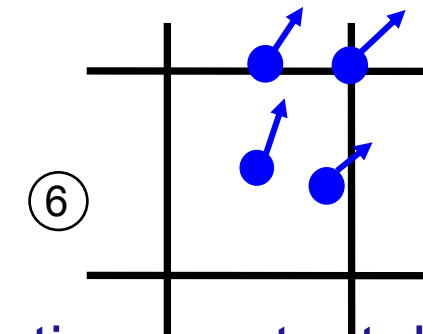
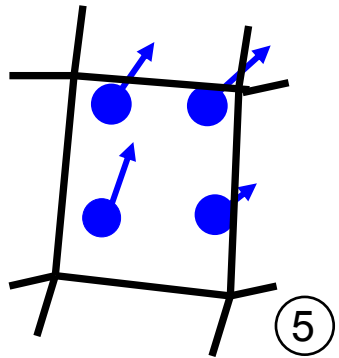
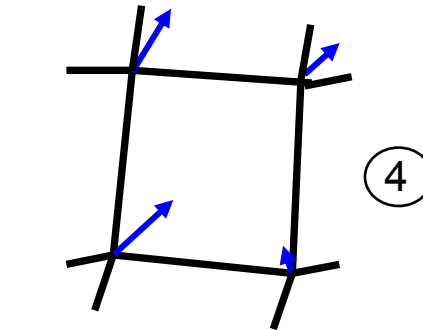
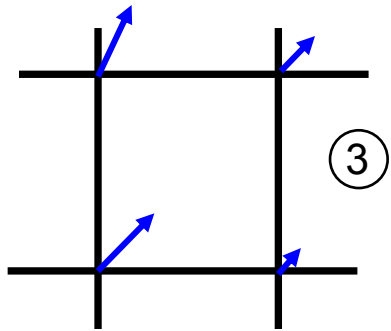
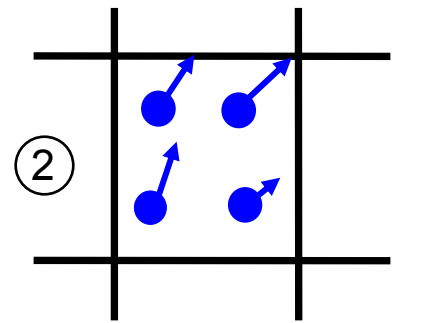
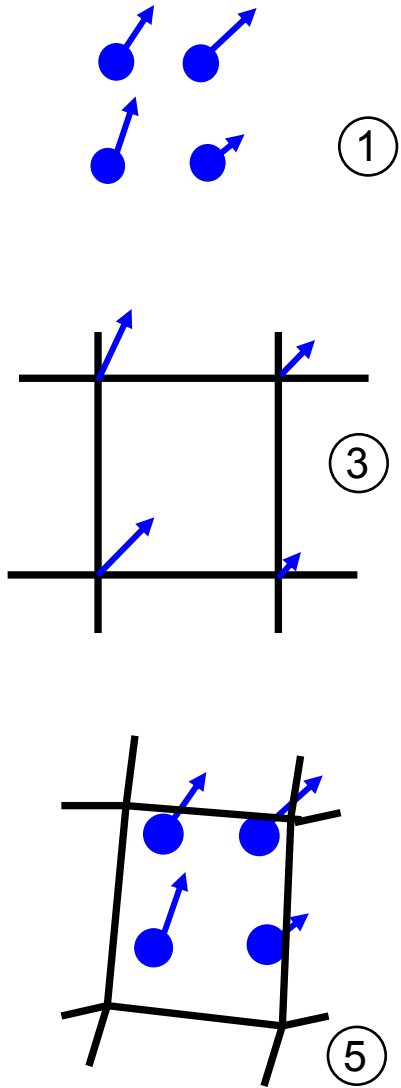
**Right: Shu-Osher Problem  
With 800 points**



**Models fluids in Uintah**

# The Material Point Method (MPM)

Sulsky Guilkey Bardenhagen et al.



Particles with properties  
(velocity, mass etc)  
defined on a mesh

Particle properties mapped  
onto mesh points

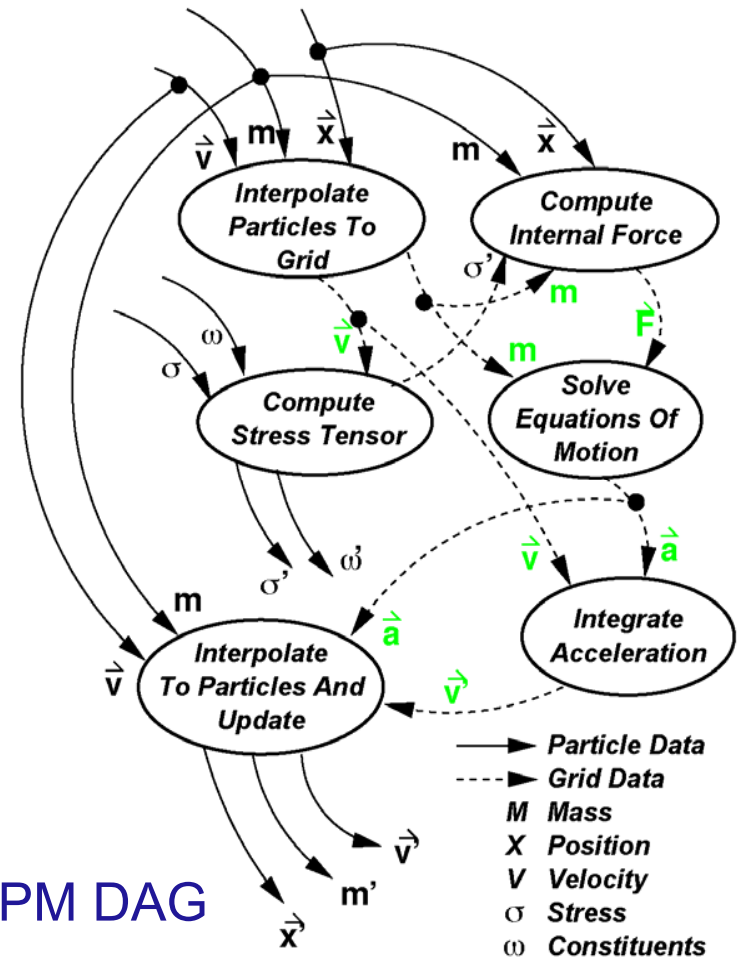
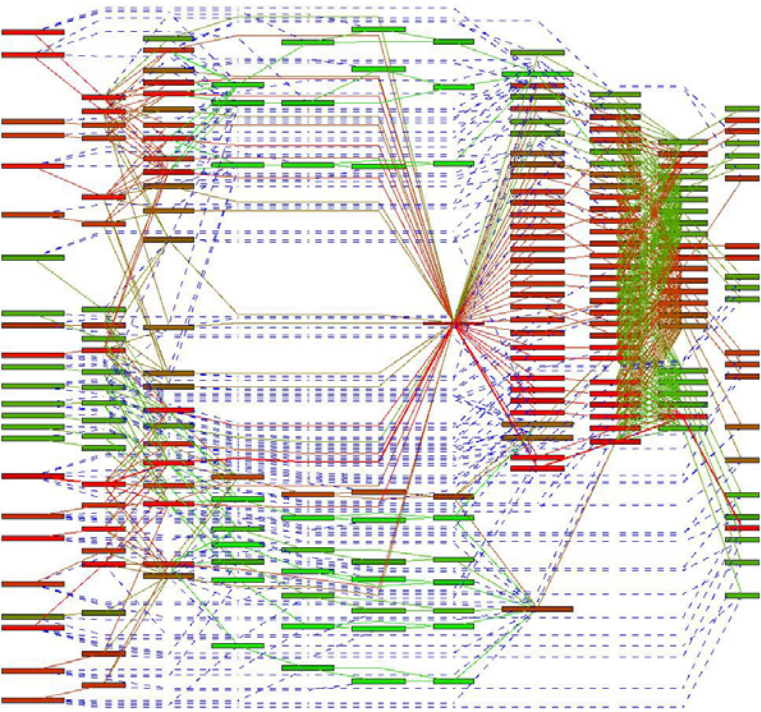
Forces, accelerations, velocities  
calculated on mesh points

Mesh point motion calculated  
but only the particles moved  
by mapping velocities back to  
particles

Handles deformation, contact, high strain, fragmentation  
models solids in Uintah

- Each task defines its computation with required inputs and outputs
- Uintah uses this information to create a task graph of computation (nodes) + communication (along edges)
- Similar to Charm++ TBlas, CnC DAG approach increasingly popular for efficient parallelism with irregular communications

## Directed Acyclic Graphs (DAGs)



# Example Uintah Task from the ICE Algorithm

Compute face-centered Velocities:

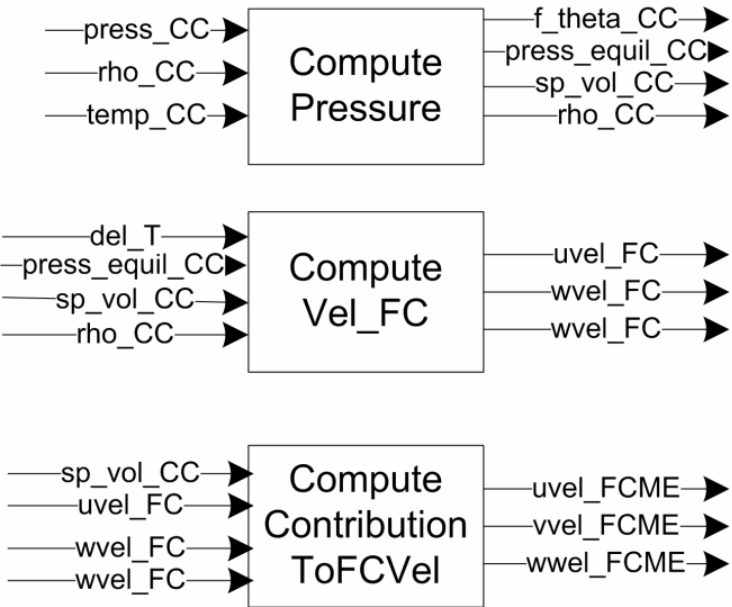
$$\vec{U}^{*f} = f(\Delta t, P_{eq}, \vec{g}, \rho, \vec{U})$$



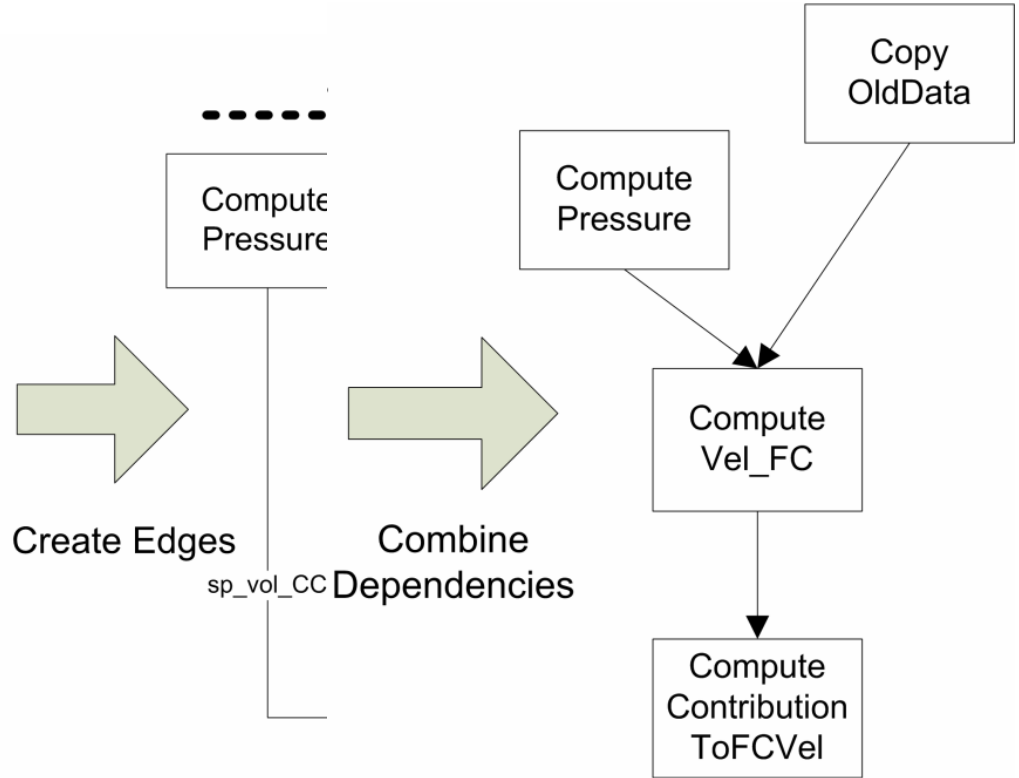
**Input variables**  
(include boundary conditions)

**Output variables**

# Task Graph Compiling



Tasks

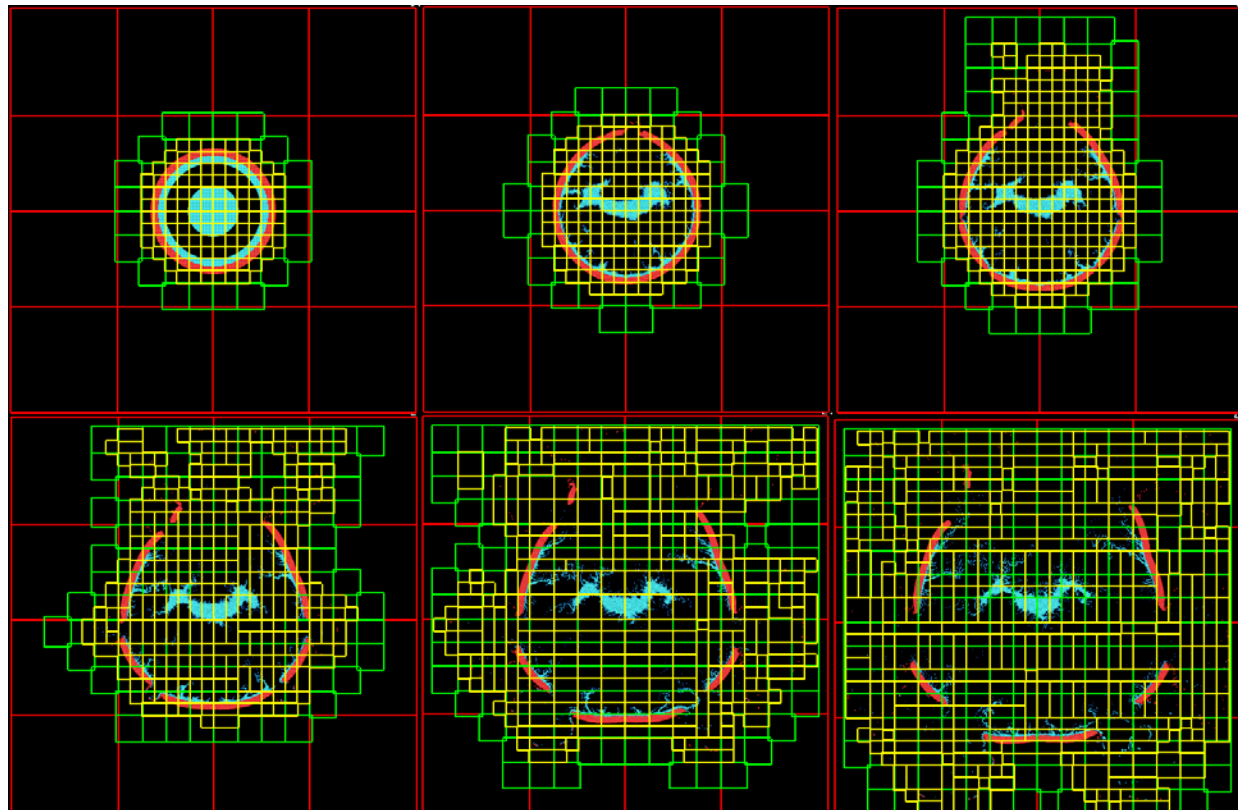


TaskGraph

# AMR for Multiple Space/Time Scales

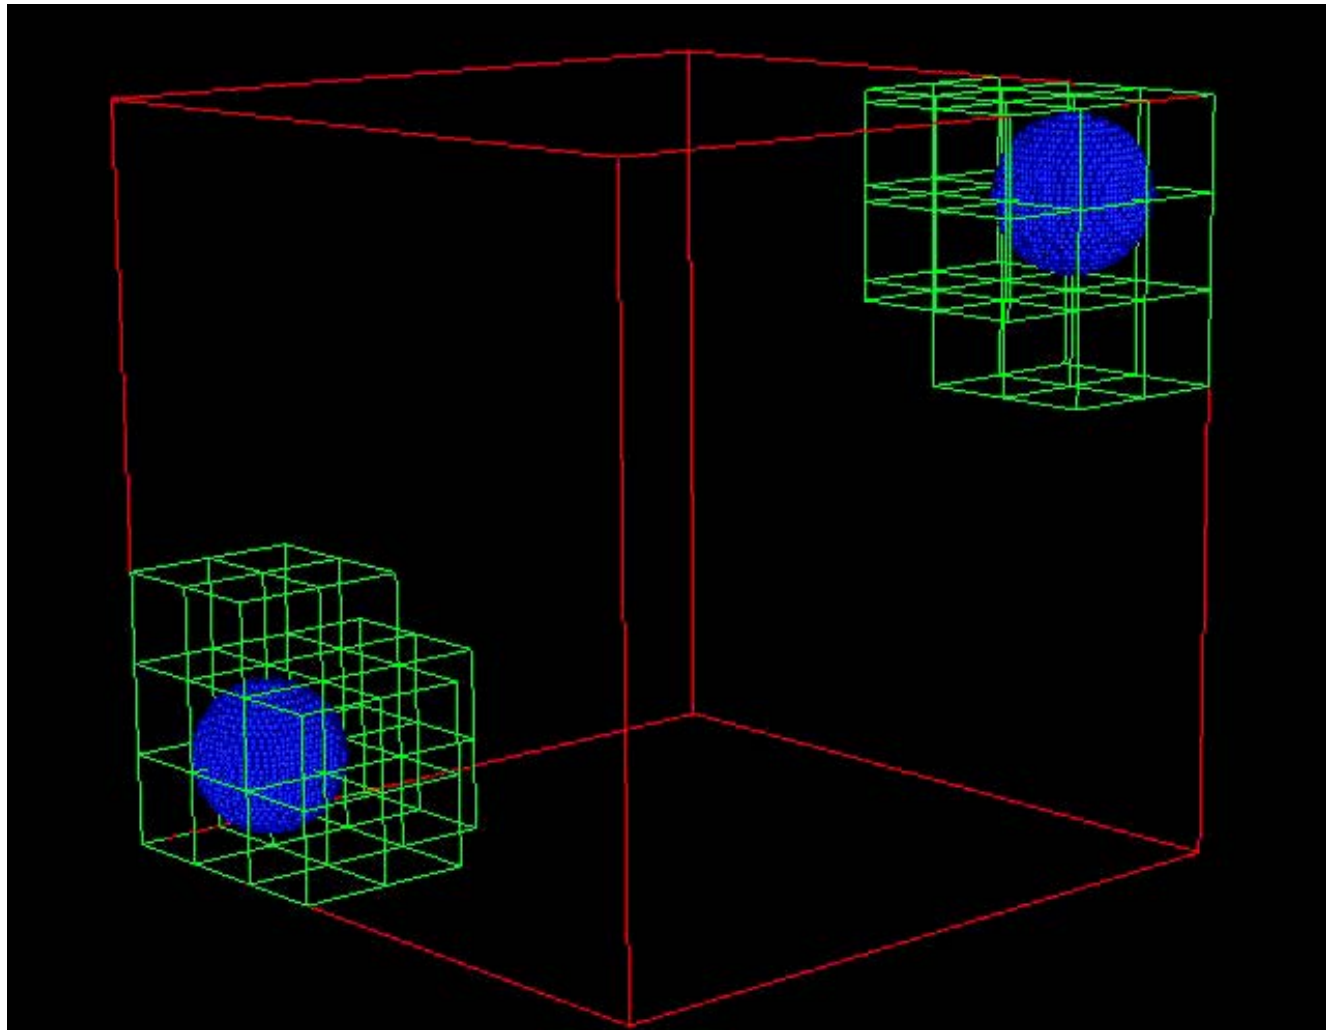
- Solvers that are designed to work together
- One mesh for all material phases

**Refine/coarsen  
spatial mesh where  
gradients/ errors  
are large/small**





# EXAMPLE OF UNTAX MOVING MESHES WITH PARTICLES





**End to end simulation of container**



# Load Balancing Weight Estimation

- **Algorithmic Cost Models** based on discretization method and machine, requires accurate information from the user
- **Time Series Analysis** used to forecast time for execution on each patch - automatically adjusts according to simulation and architecture with no user interaction

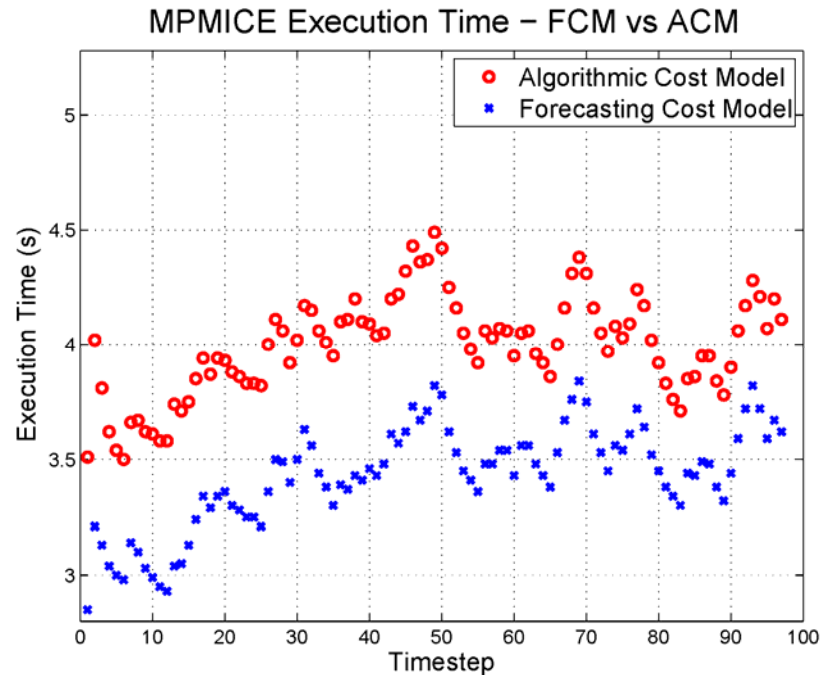
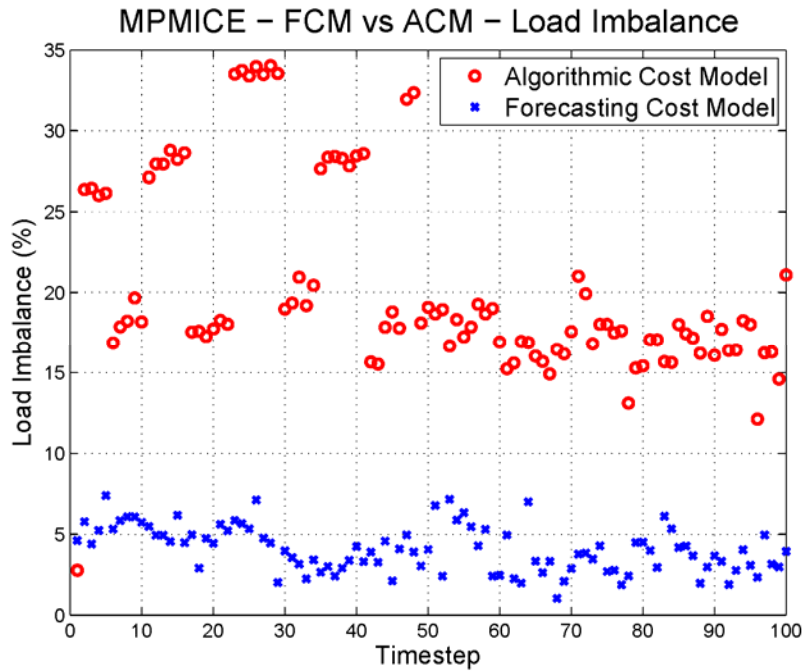
## Simple Exponential Smoothing:

**$E_{r,t}$ : Estimated Time**     **$O_{r,t}$ : Observed Time**     **$\alpha$ : Decay Rate**

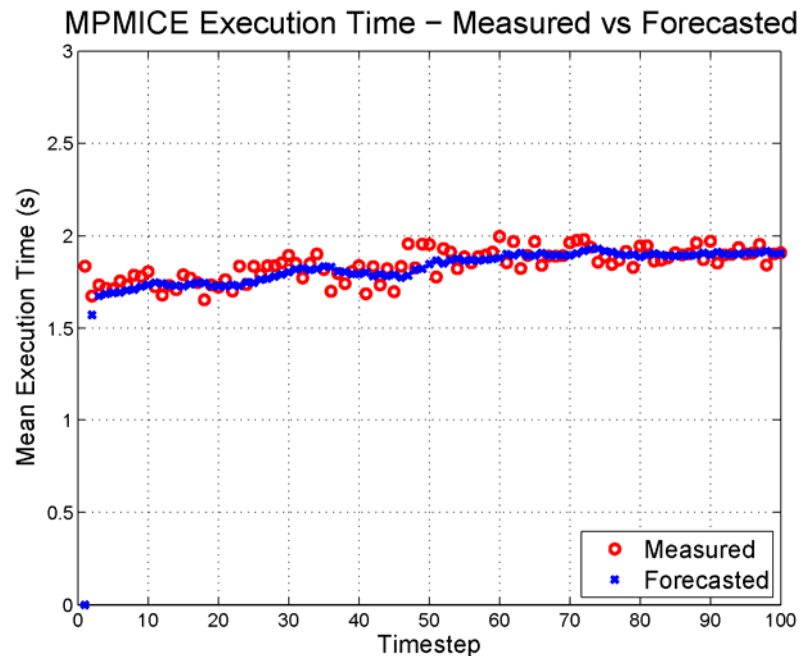
$$\begin{aligned} E_{r,t+1} &= \alpha O_{r,t} + (1 - \alpha) E_{r,t} \\ &= \alpha \underbrace{(O_{r,t} - E_{r,t})}_{\text{Error in last prediction}} + E_{r,t} \end{aligned}$$

Error in last prediction

[IPDPS10 paper], Charm++ uses a similar idea without feedback

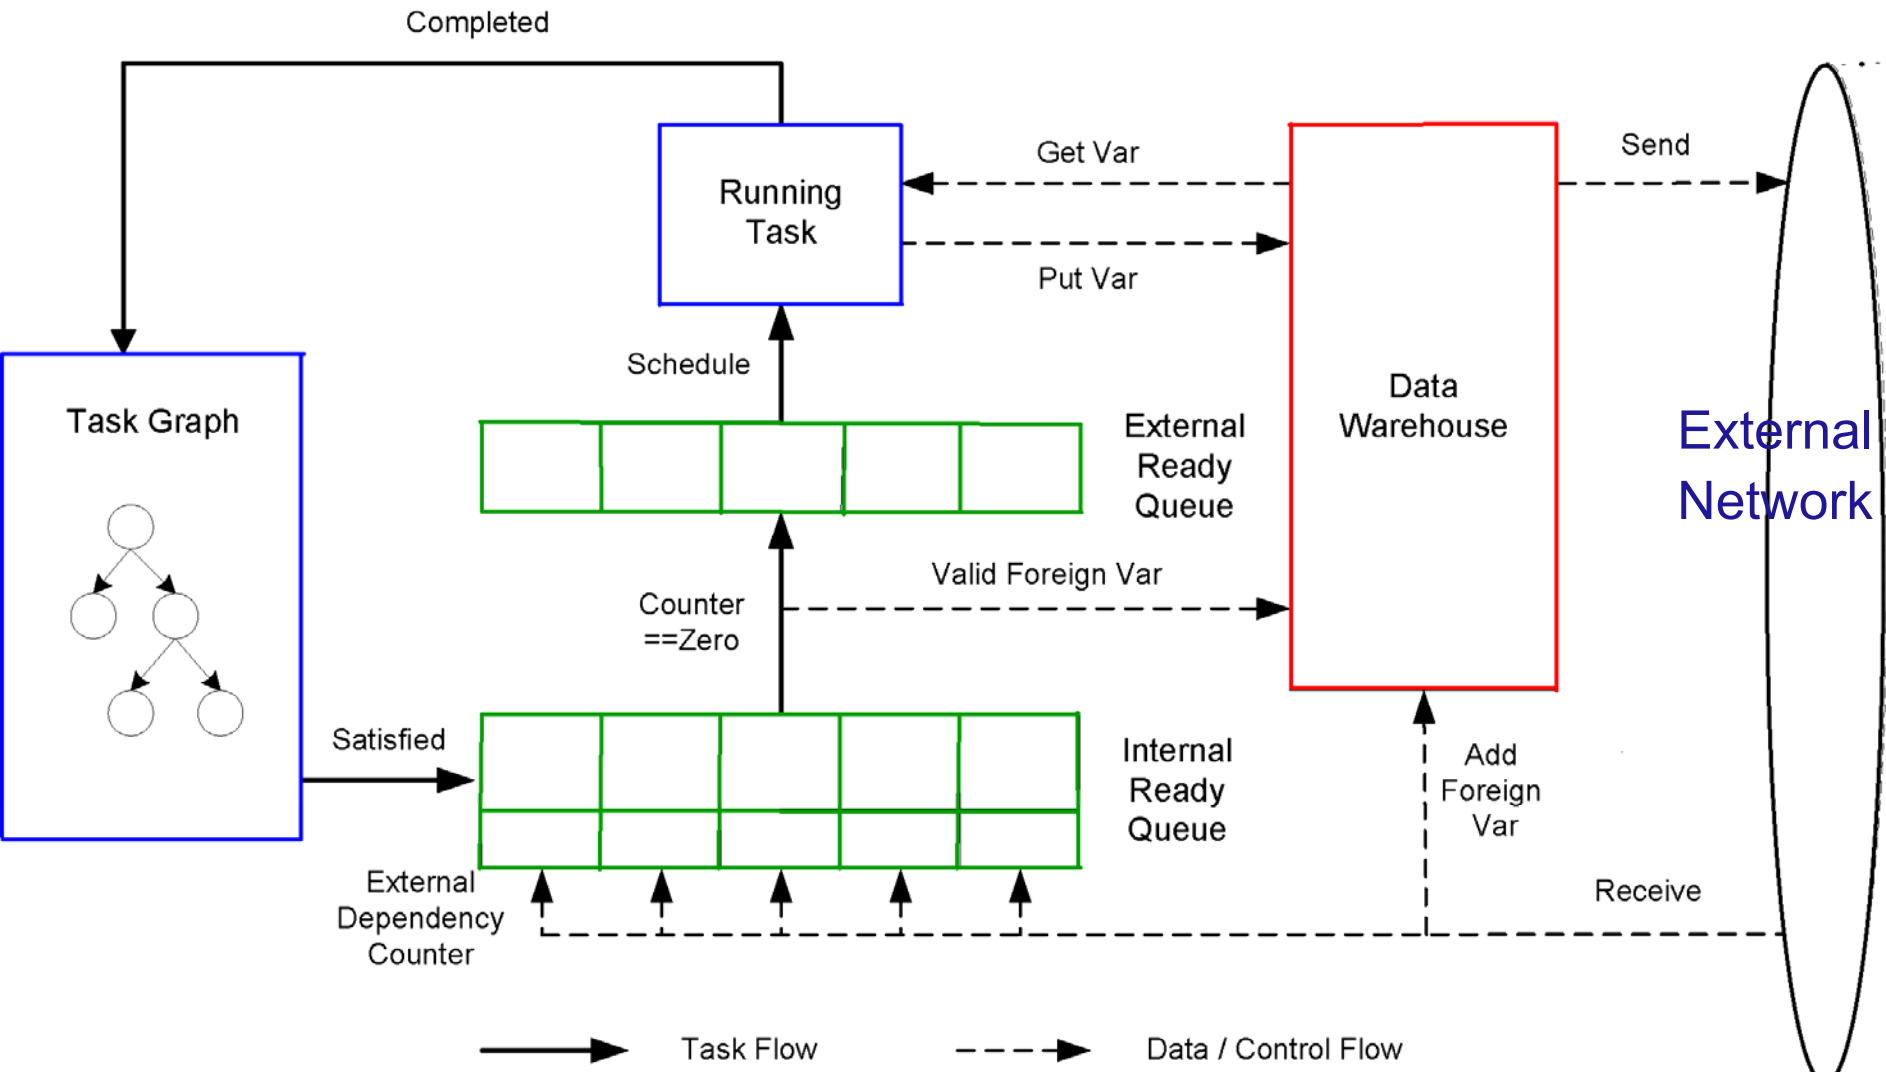


**Comparison between  
Forecast Cost Model FCM  
& Algorithmic Cost Model  
Particles + Fluid code  
FULL SIMULATION**

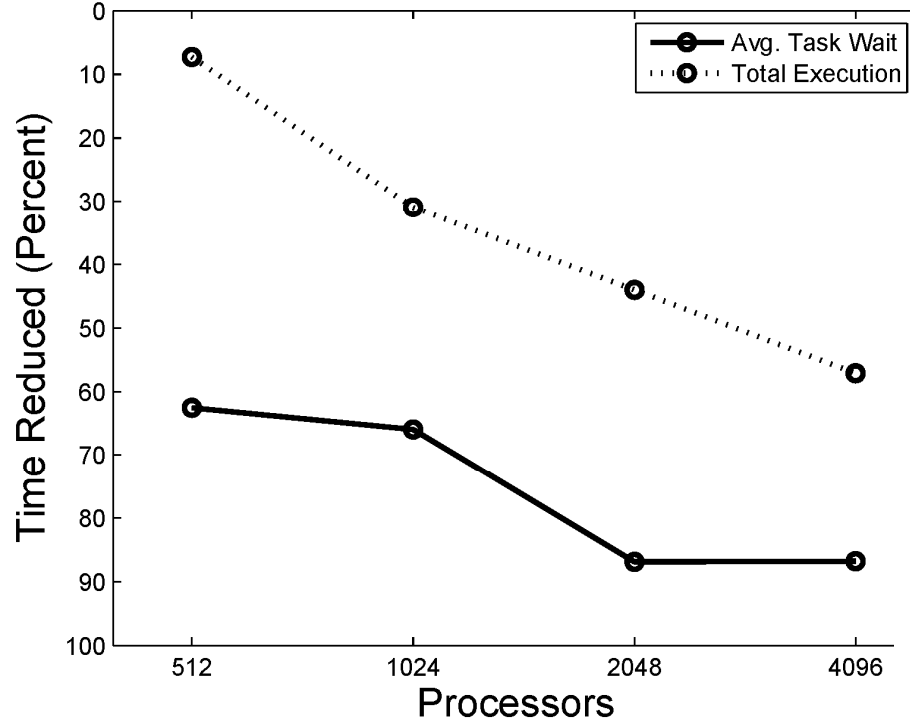


# DYNAMIC TASK EXECUTION

When a task's external dependencies are satisfied it can execute



ICE Dynamic vs Static Scheduling (TACC Ranger) 62K AMD cores Sun infiniband

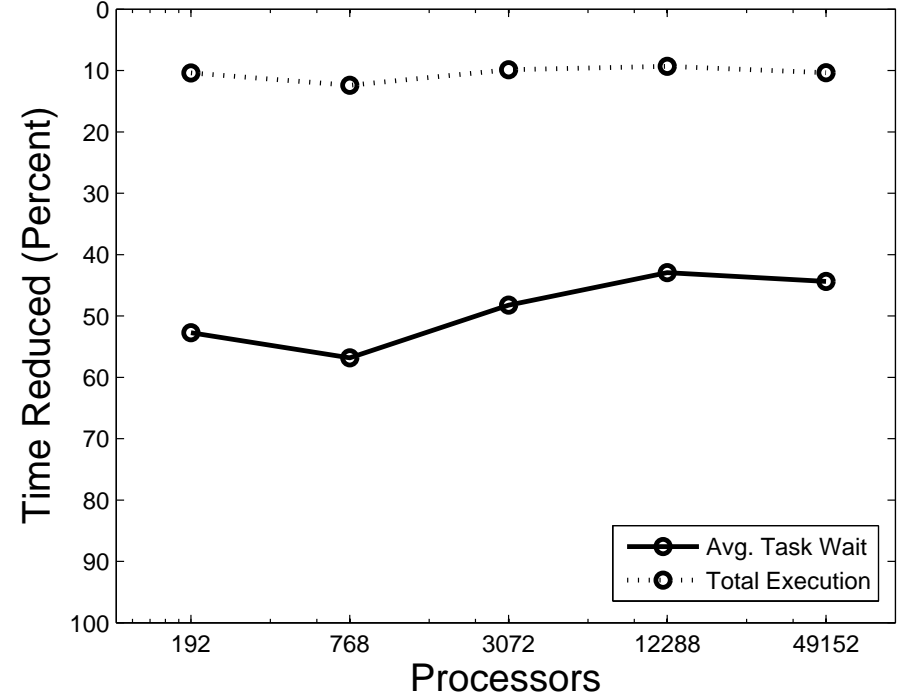


# Static vs Dynamic Scheduling Improvements

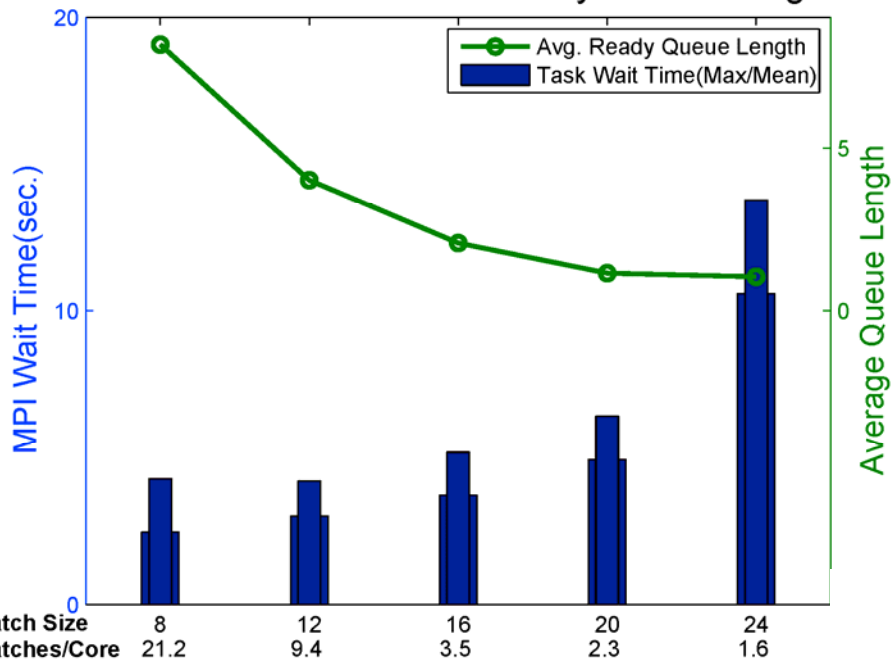
**Ranger has slower communications than Kraken and so we see less improvement in overall time through overlapping and out-of-order execution**

98K AMD 6 cores CRAY Seastar torus

ICE Dynamic vs Static Scheduling (NICS Kraken)



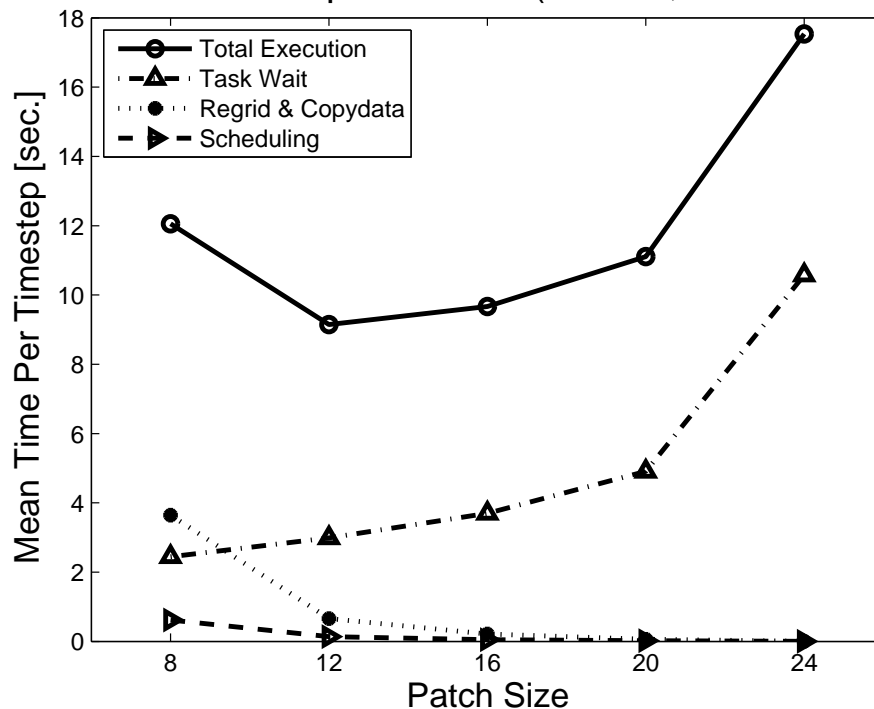
ICE Wait Time and Task Ready Queue Length



**Task queue length drops and wait time increases as patch sizes grow for a fixed mesh**

**Execution time is a trade-off  
Between granularity improvements  
and overhead due to more patches**

ICE with different patch sizes (Kraken, with 24K cores)

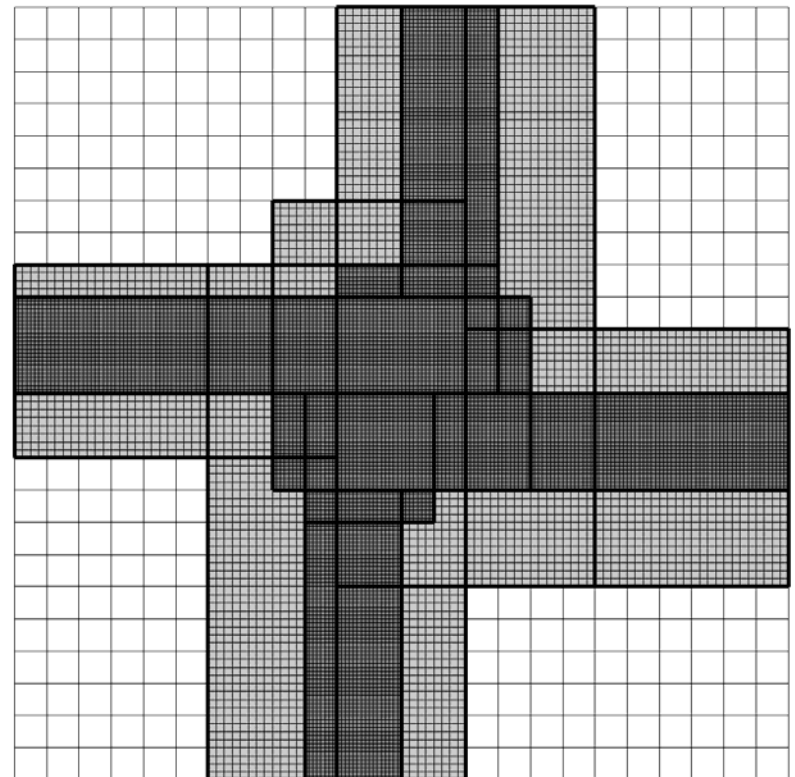
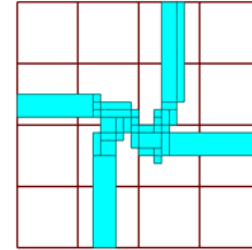
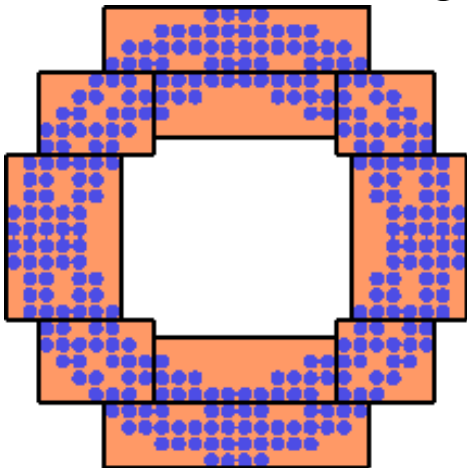


**GRANULARITY EFFECTS**



# Scalable AMR Regridding Algorithm

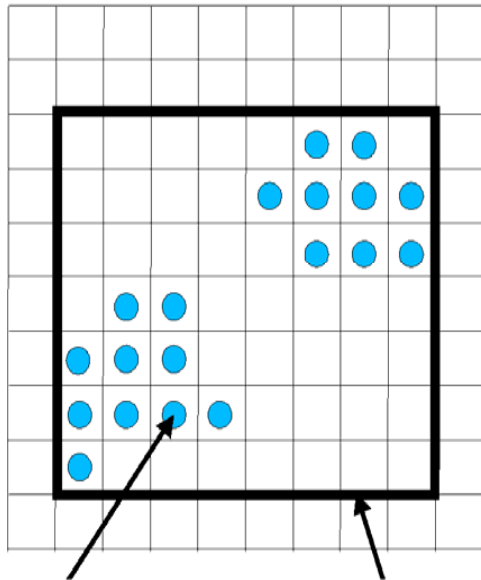
- Berger-Rigoutsos
  - Recursively split patches based on histograms
  - Histogram creation requires all-gathers  $O(\text{Processors})$
  - Complex and does not parallelize well
  - Irregular patch sets
  - Two versions – version 2 uses less cores in forming histogram



# AMR regridding algorithm (Berger-Rigoutsos)

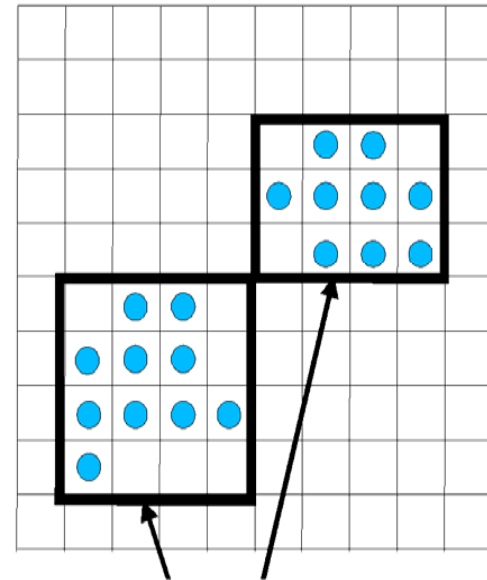
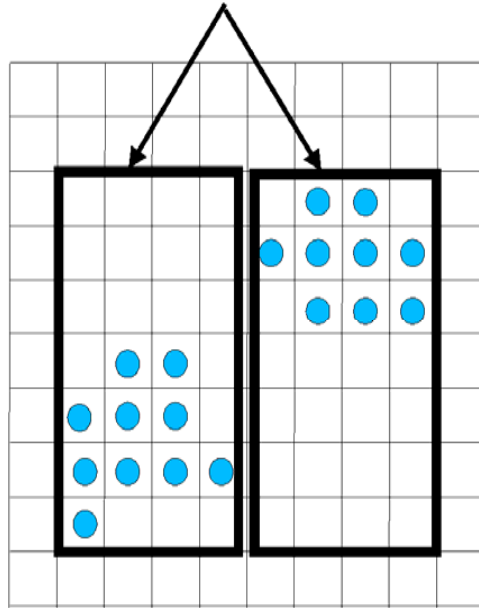
Number of points  
in each column

3 3 3 1 1 3 3 2



Tagged points Initial box

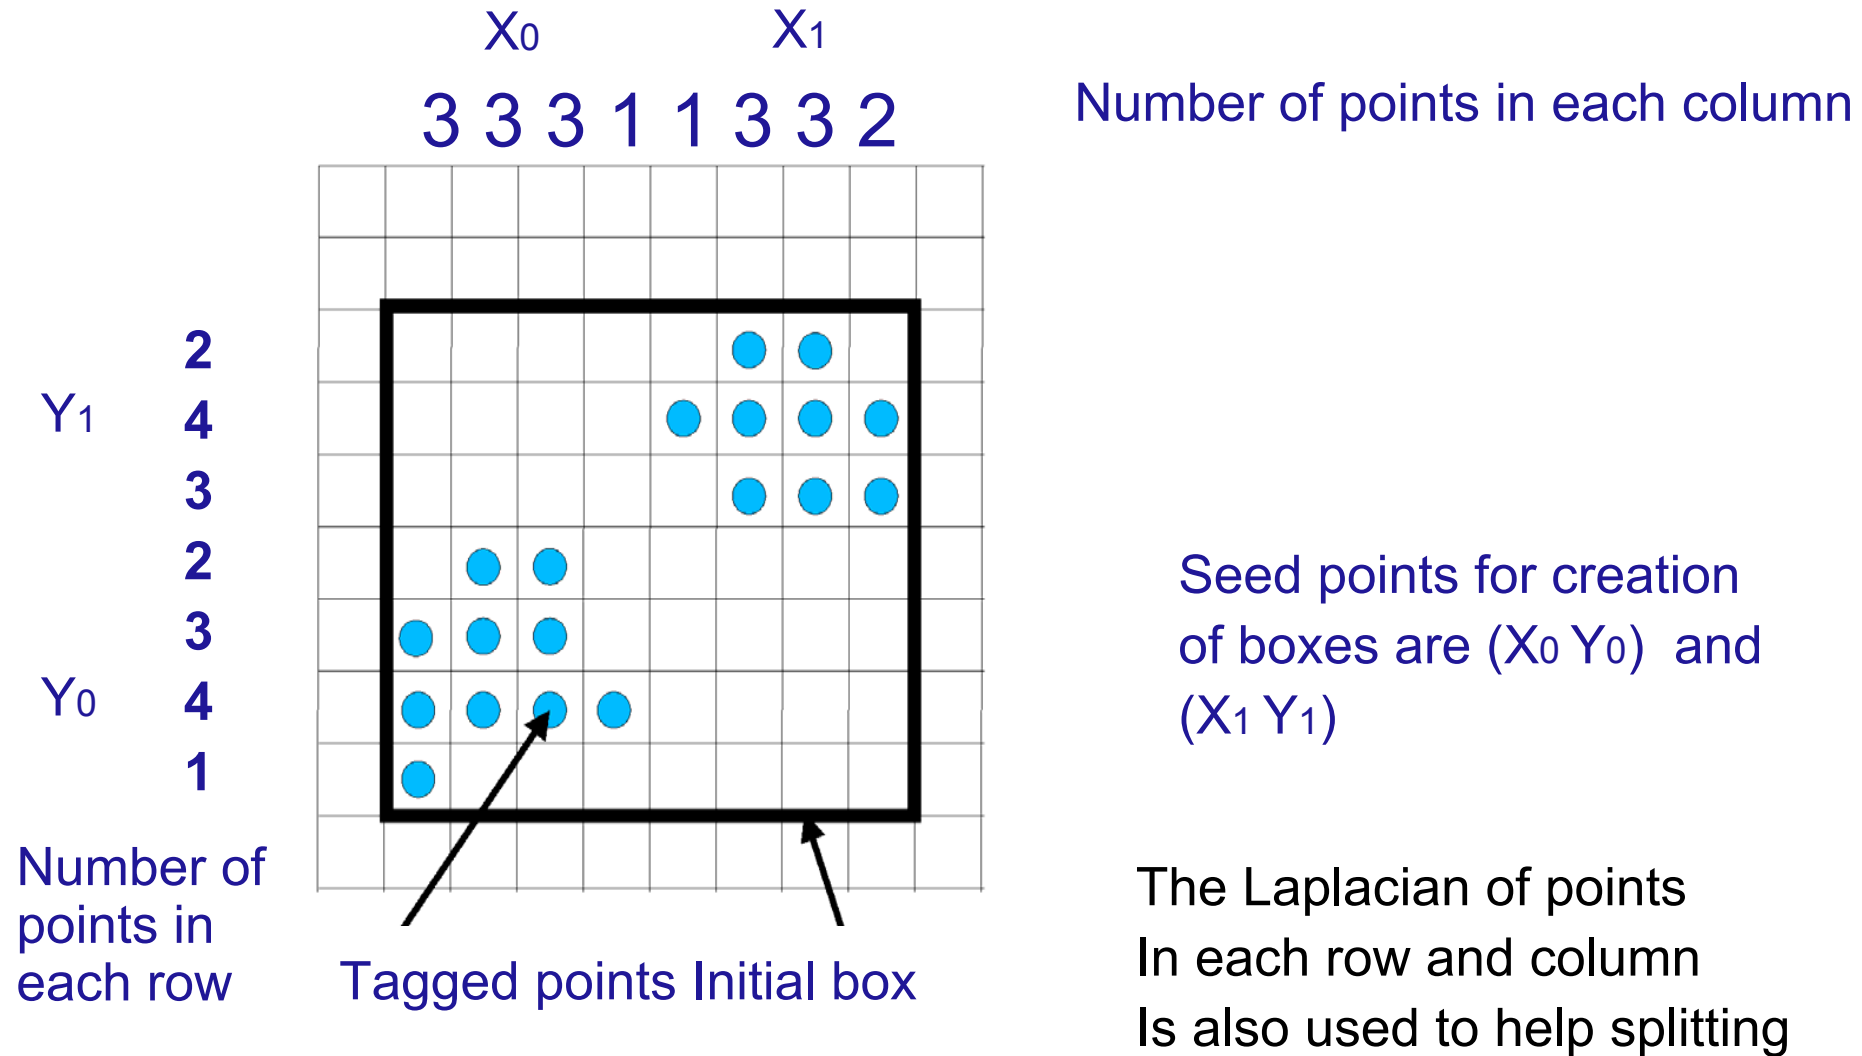
Box is split into 2



Process is repeated on the  
two new boxes

- (1) tag cells where refinement is needed
- (2) create a box to enclose tagged cells
- (3) split the box along its long direction based on a histogram of tagged cells
- (4) fit new boxes to each split box and repeat the steps as needed.

# AMR regridding algorithm (Berger-Rigoutsos)



**This step is repeated at least  $\log(B)$  times where  $B$  is the number of patches**

# MPI\_ALLGATHER NEED

- Every processor has to pass its part of the histogram to every other so that the partitioning decision can be made
- The cost of an allgather is  $\text{Log}(P)$

Version 1:  $(2B-1)\log(P)$  messages

Version 2: Only certain cores take part  
[Gunney et al.]

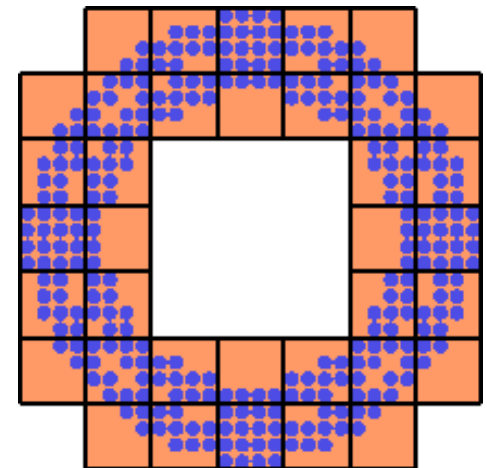
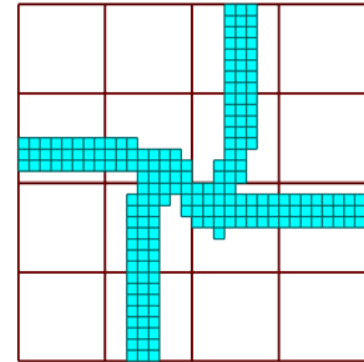
$2P - \log(P) - 2$  messages

**Alternative: use Berger  
Rigoutsos locally on each  
processor**



# Scalable AMR Regridding Algorithm

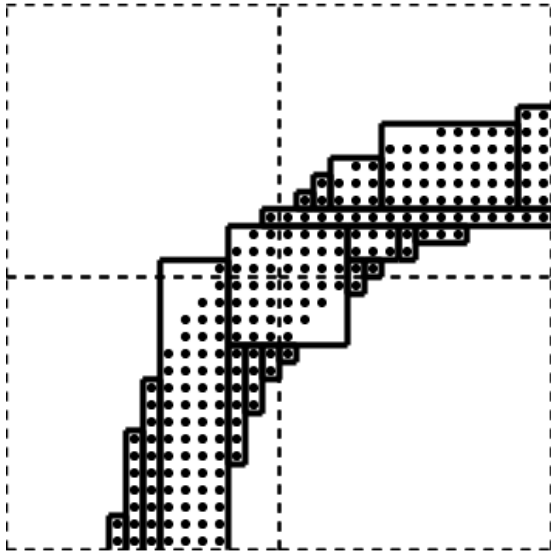
- Tiled Algorithm
  - Tiles that contain flags become patches
  - Simple and easy to parallelize
  - Semi-regular patch sets that can be exploited
    - **Example: Neighbor finding**
  - Each core processes subset of refinement flags in parallel-helps produce global patch set



[Analysis to appear in Concurrency]

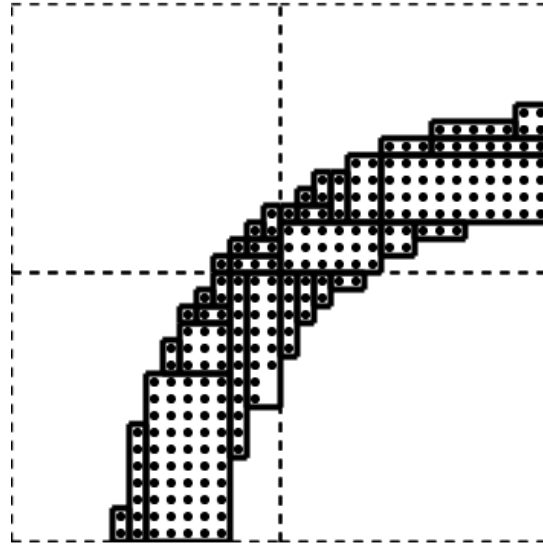
# EXAMPLE – MESH REFINEMENT AROUND A CIRCULAR FRONT

BNR Patches



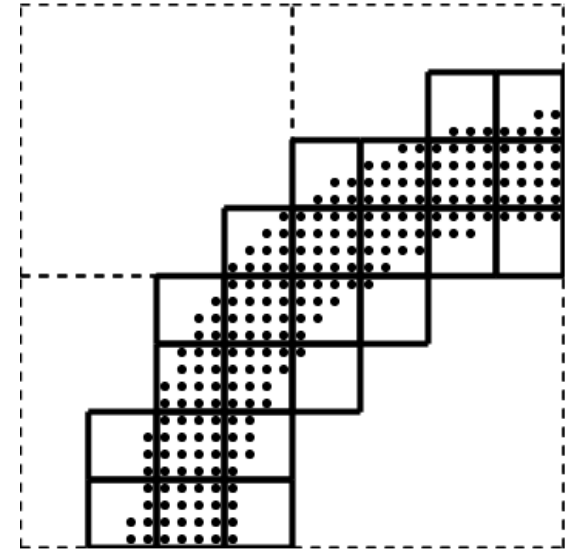
**Global Berger-Rigoutsos**

LBR Patches



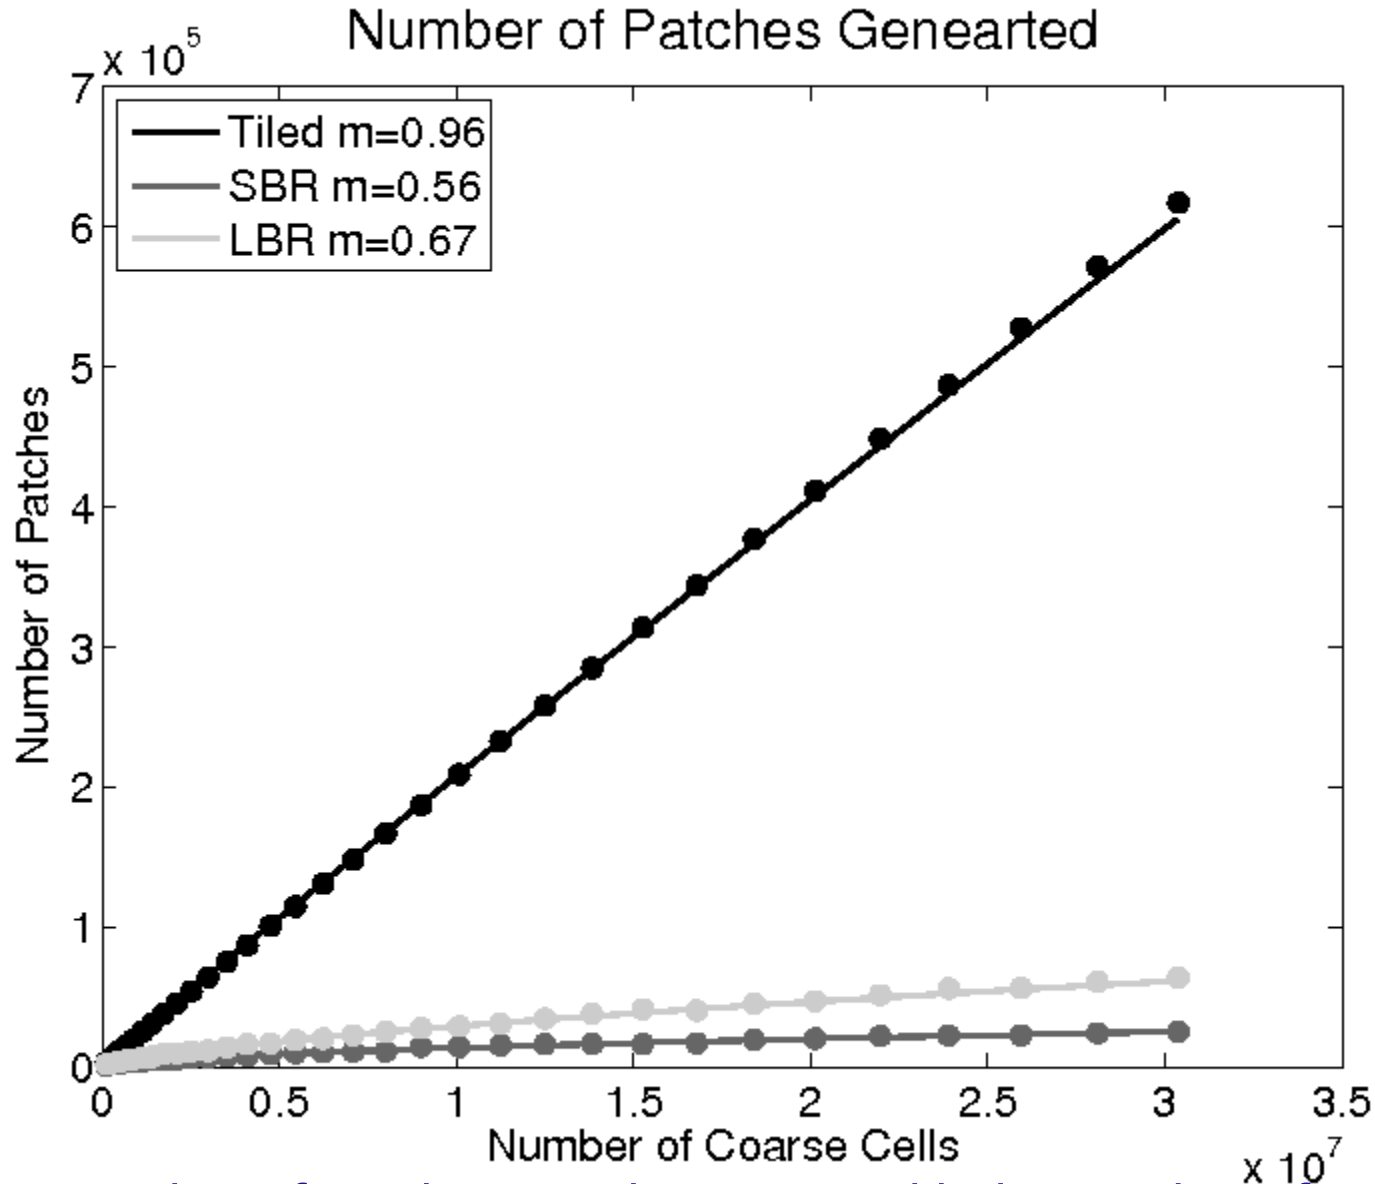
**Local Berger-Rigoutsos**

Tiled Patches



**Tiled Algorithm**

# NUMBER OF PATCHES GENERATED



The number of patches needs to grow with the number of cells on a parallel machine – otherwise we need to partition patches

# Theoretical Models of Refinement Algorithms

$C$  = number of mesh cells in domain

$F$  = number of refinement flags in domain

$B$  = number of mesh patches in the domain

$B_c$  = number of coarse mesh patches in the domain

$P$  = number of processing cores

$M$  = number of messages

$$T: \mathbf{GBRv1} = c_1 (F/P) \log(B) + c_2 M(\mathbf{GBRv1})$$

$$T: \mathbf{GBRv2} = c_1 (F/P) \log(B) + c_3 M(\mathbf{GBRv2})$$

$$T: \mathbf{LBR} = c_4 (F/P) \log(B/B_c)$$

$$T: \mathbf{Tiled} = c_5 C/P$$

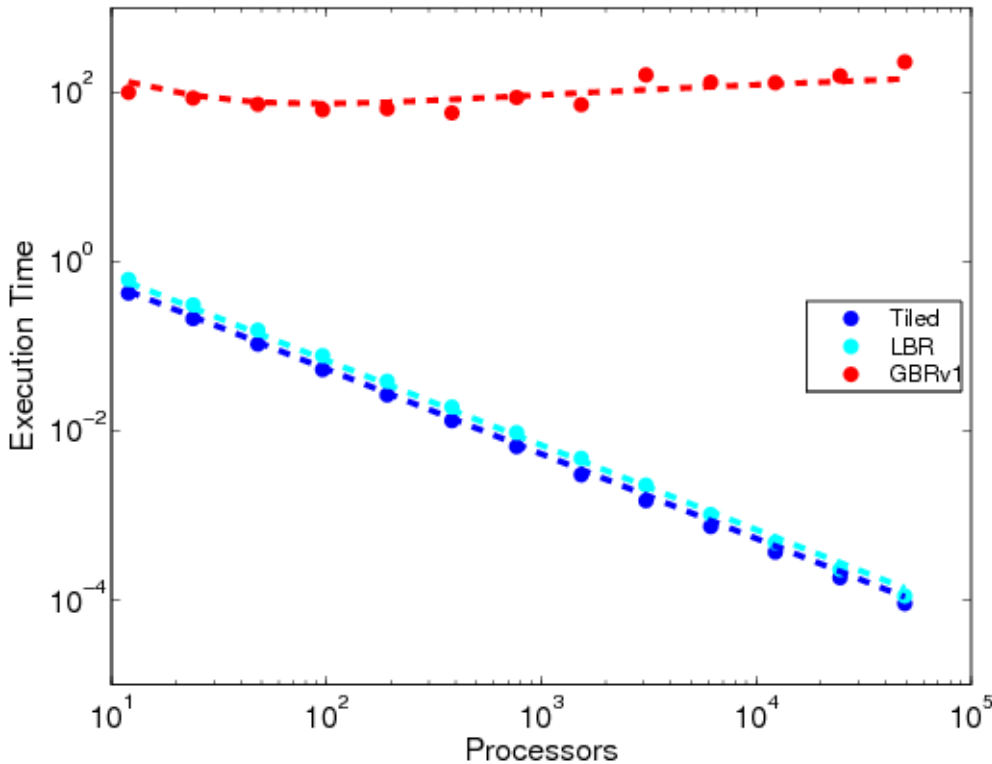
$$T: \mathbf{Split} = c_6 B/P + c_7 \log(P) - \text{ is the time to split large patches}$$



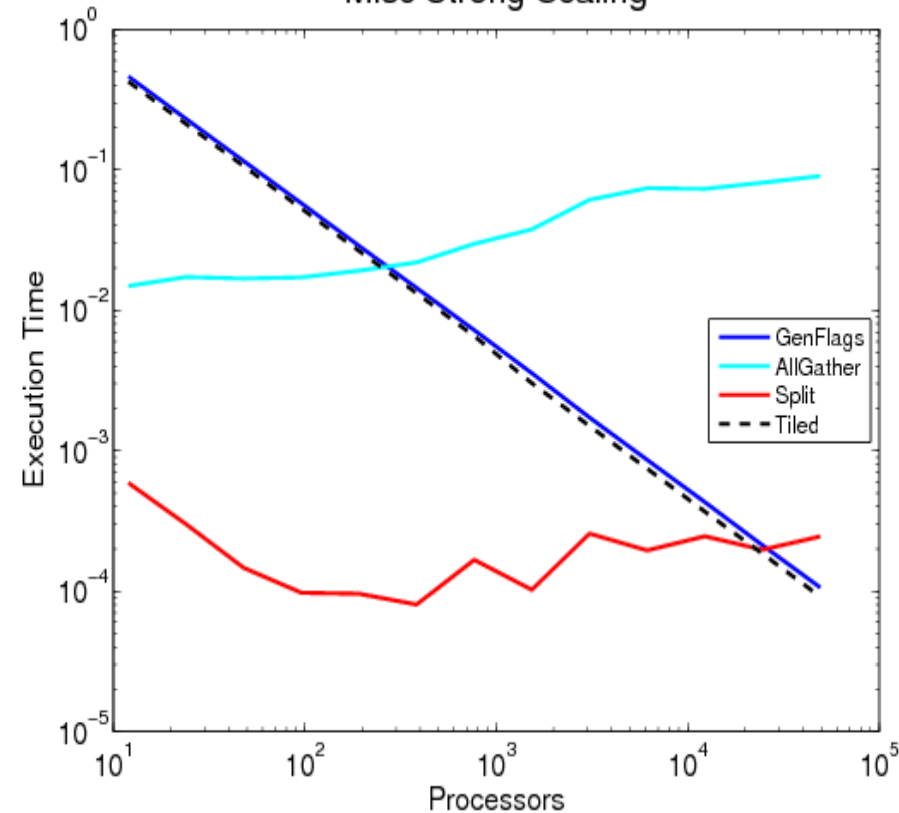
# Strong Scaling of Algorithms and their Components

Dots are data lines are models

### Regridder Strong Scaling

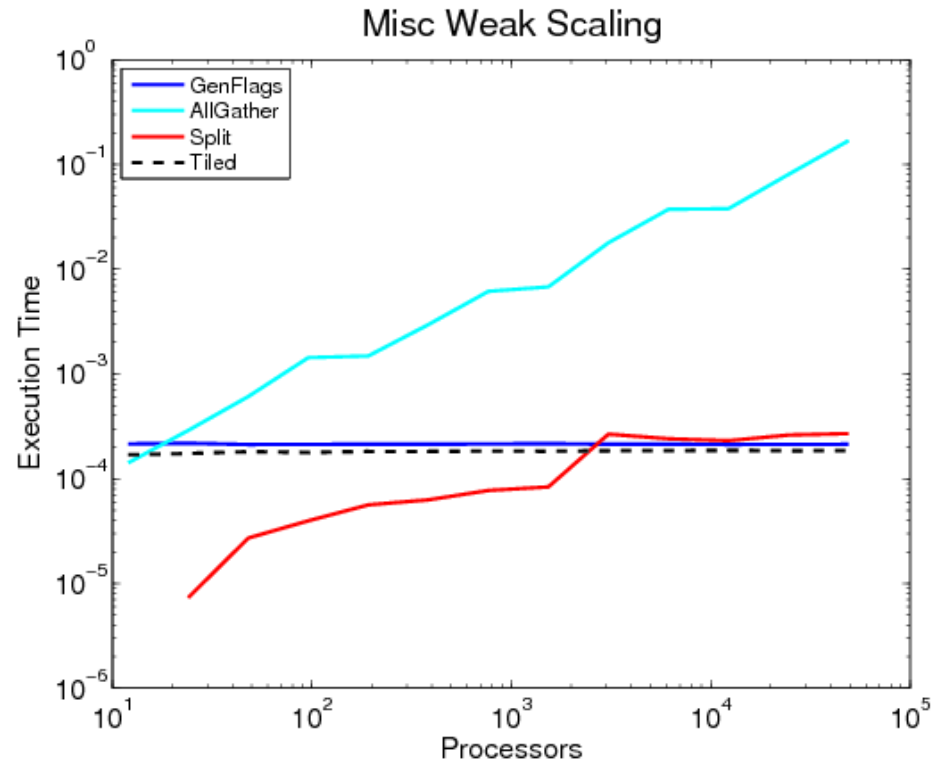
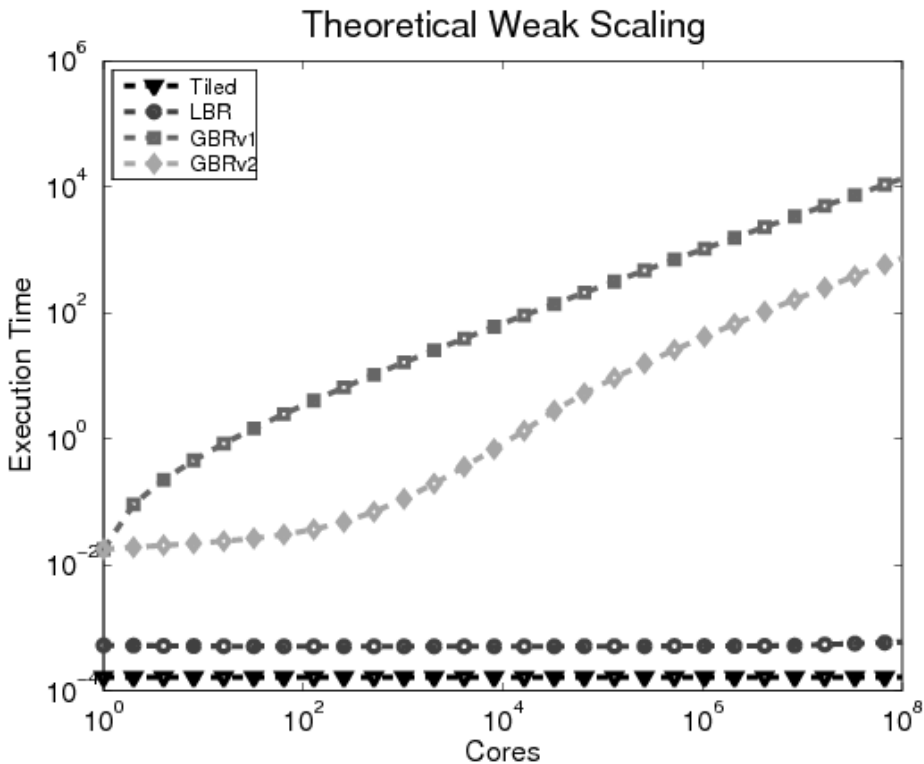


### Misc Strong Scaling



**Strong Scaling** – problem size fixed as number of cores increases should lead to decreasing execution time.

# Weak Scaling of Algorithms and their Components



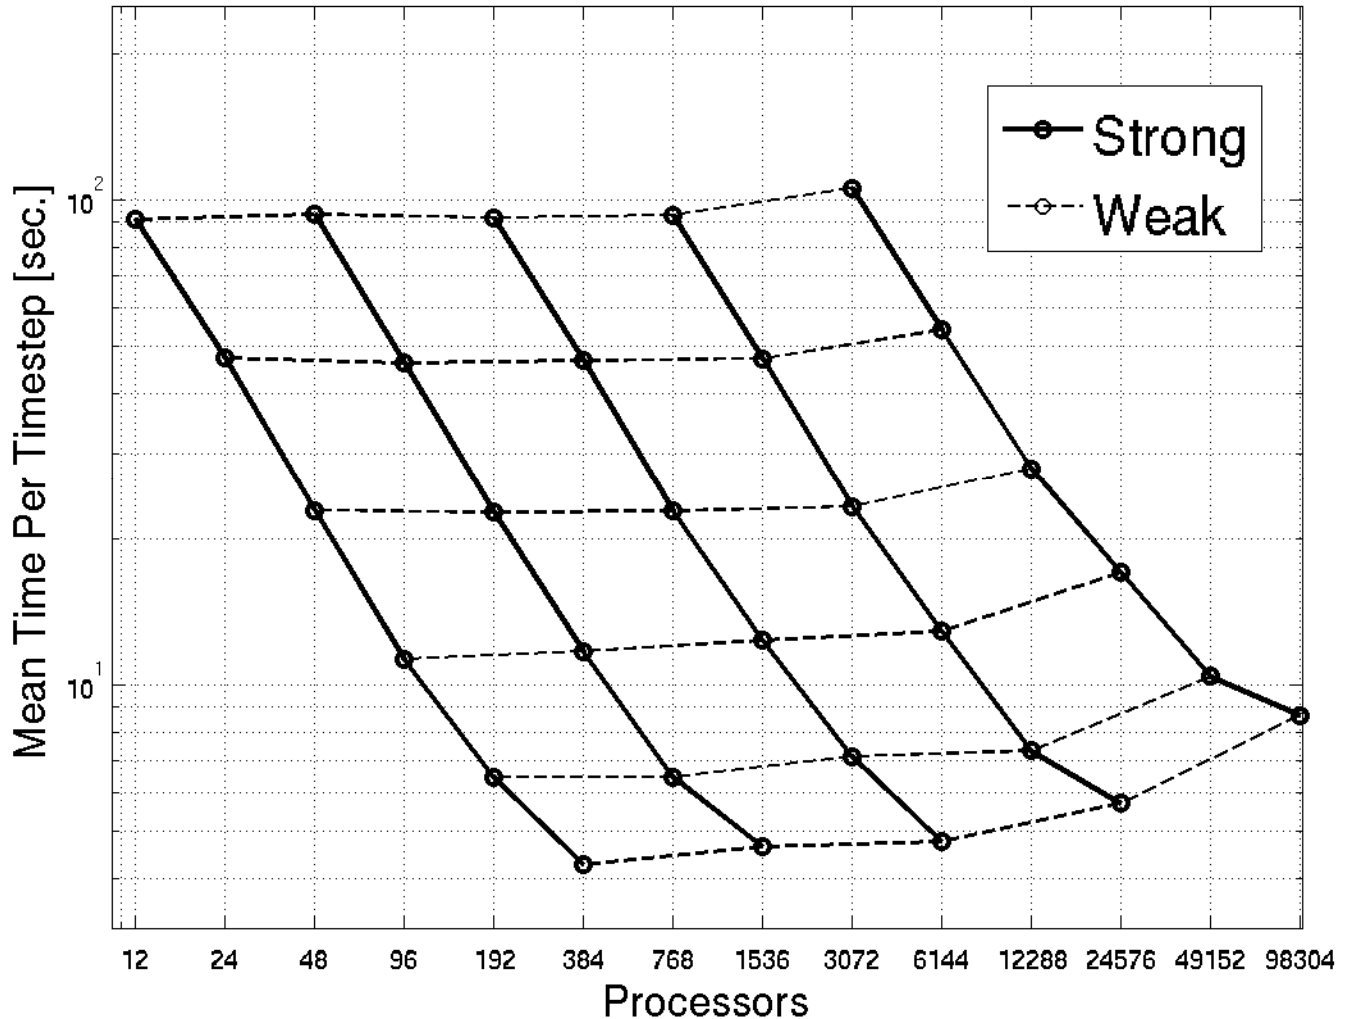
**Weak Scaling** – problem size per core fixed as number of cores increases should lead to constant execution time.

# UINTAH SCALABILITY

At 98k Proc  
1 16x16x16  
patch per  
Core and so  
Scalability fades

Problem is  
essentially  
an advected  
blob, but  
formulated  
as compress.  
Navier Stokes  
In 3D

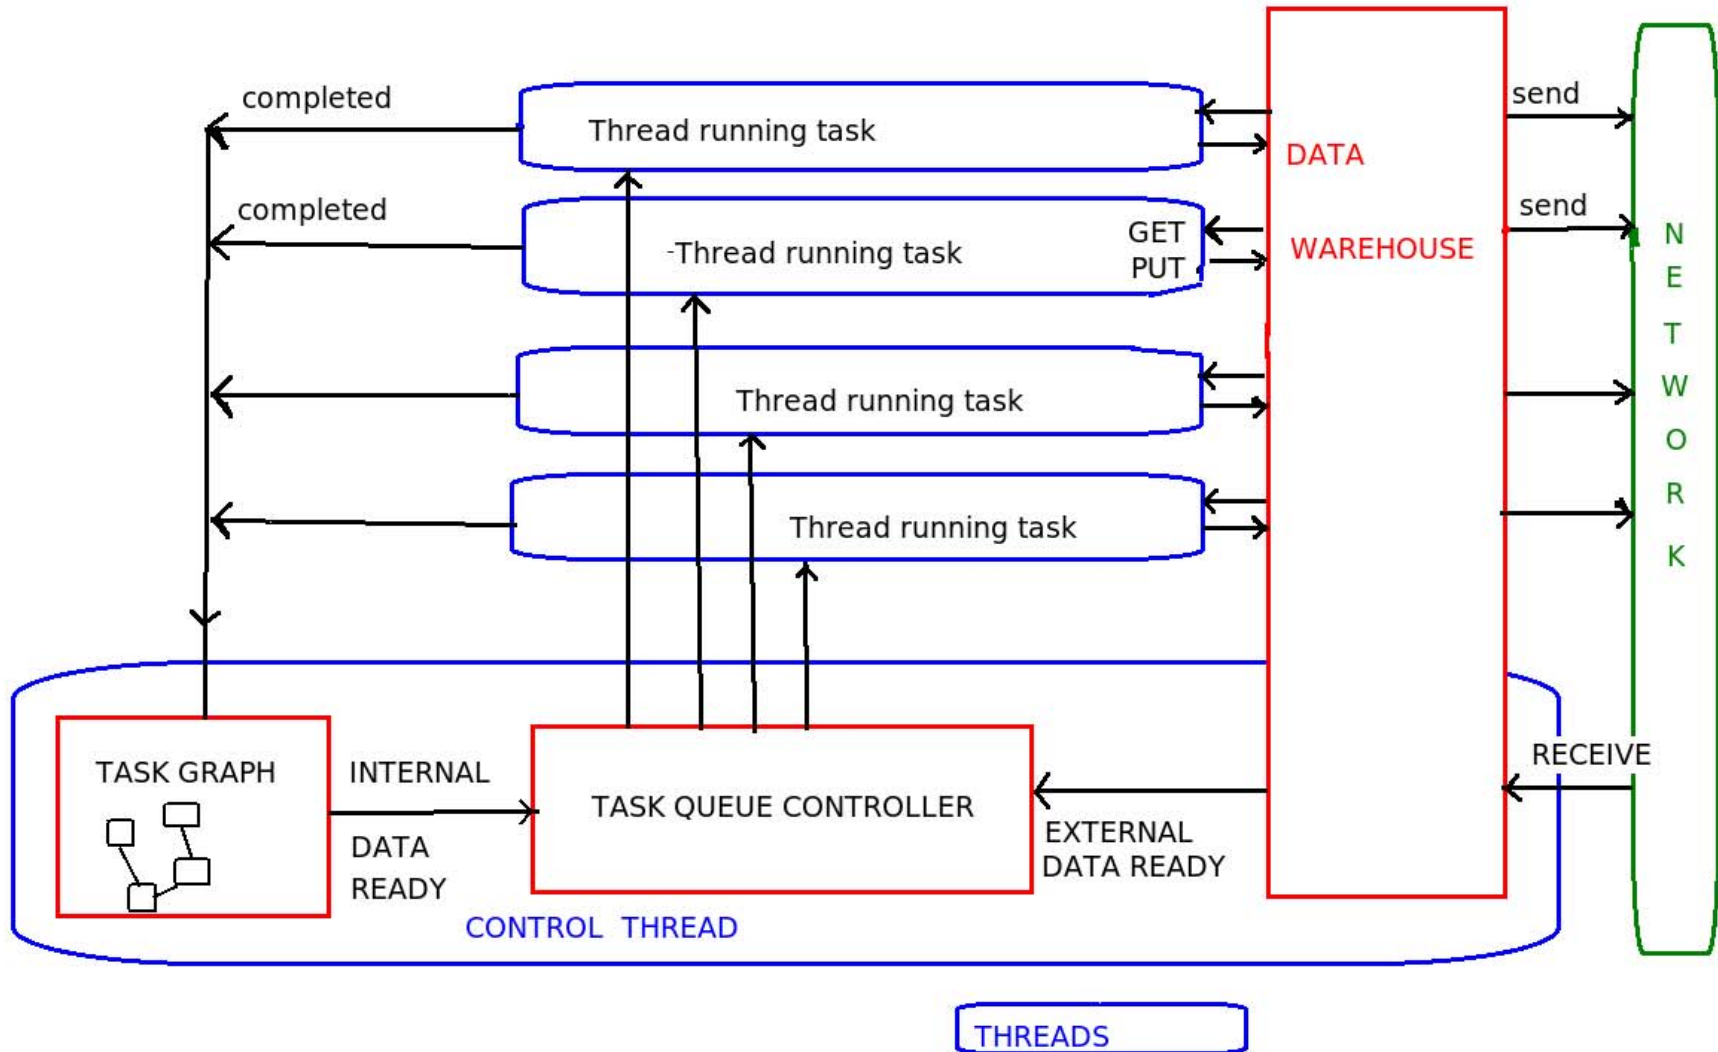
## AMR-ICE Scaling



NSF NICS Kraken 6-core AMD based machine

# Uintah Hybrid MPI/Pthreads [TG11]

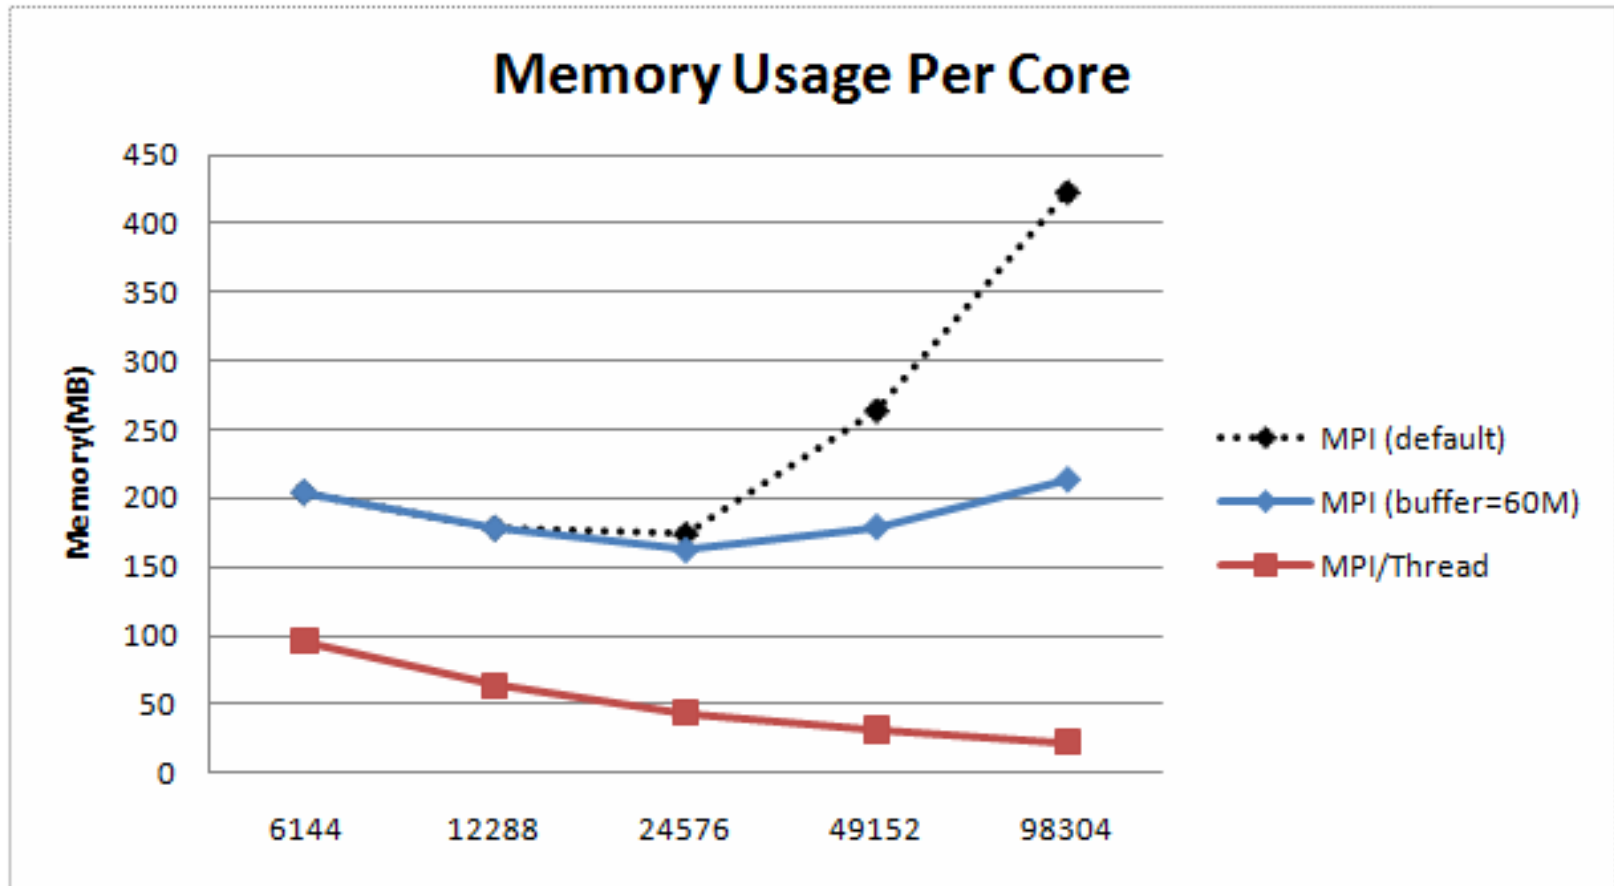
use only one copy of data warehouse per node  
as opposed to per MPI process



# Uintah Memory Used Strong Scaling

Can prove that global memory drops to  $(\text{global memory})/(\text{ncores})$

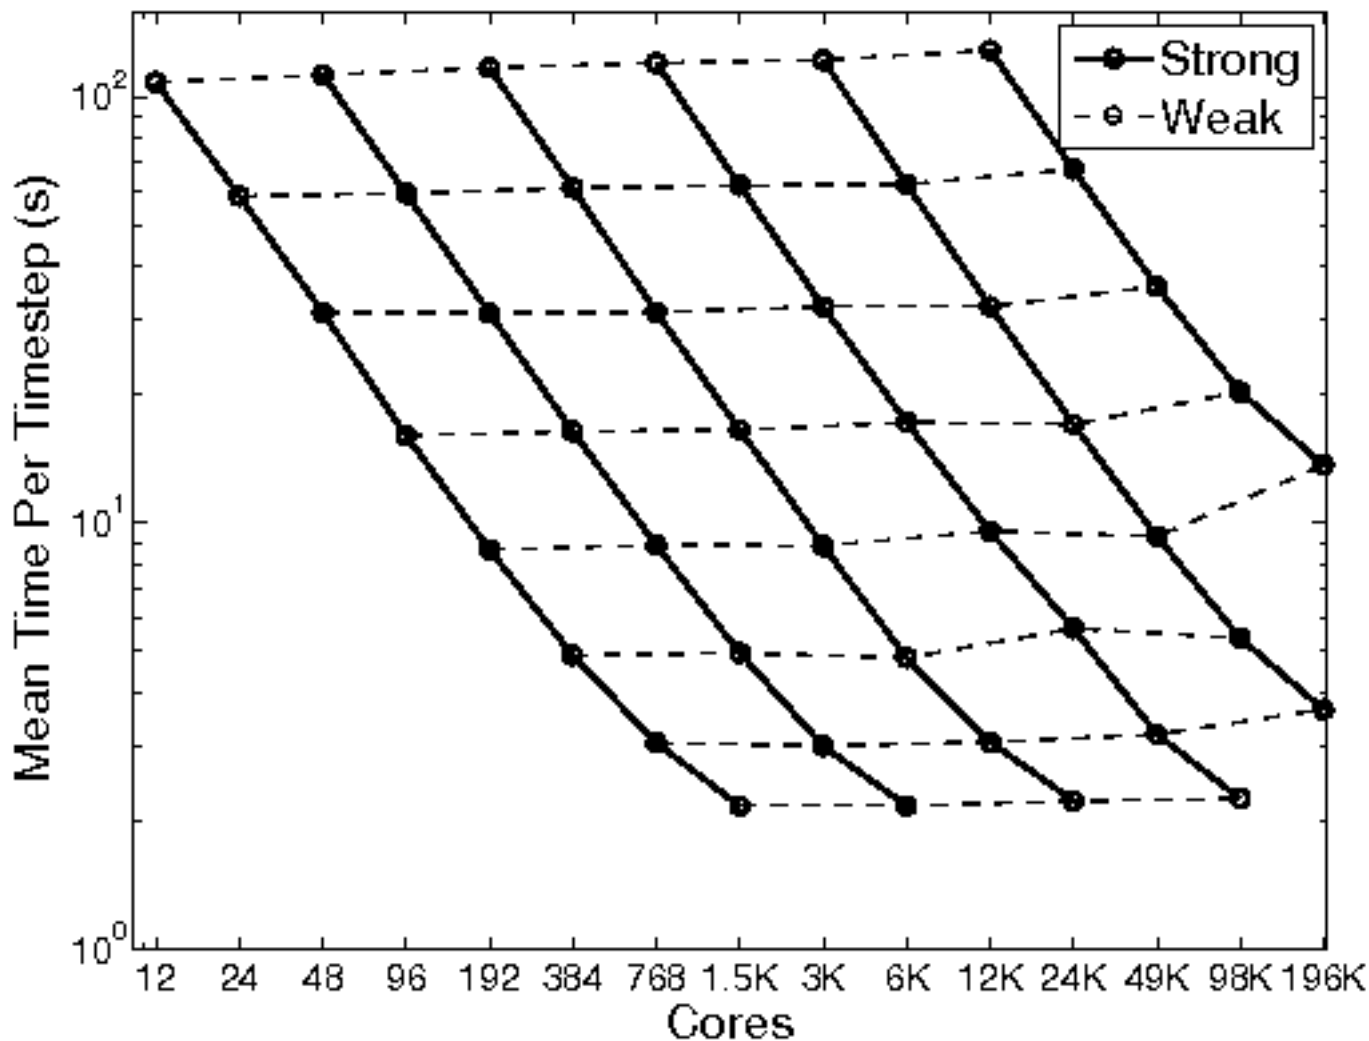
Drop is more if Cray MPI buffers badly set....



Uintah now runs on 200K cores with only 10% of previous global memory per node and similar ~5% cpu times [TG11].

# UINTAH SCALABILITY

## AMR ICE: Scaling



Problem is essentially an advected blob, but formulated as compressible Navier Stokes Equations in 3D

DOE Jaguar 6-core AMD based machine

## DEFLAGRATION

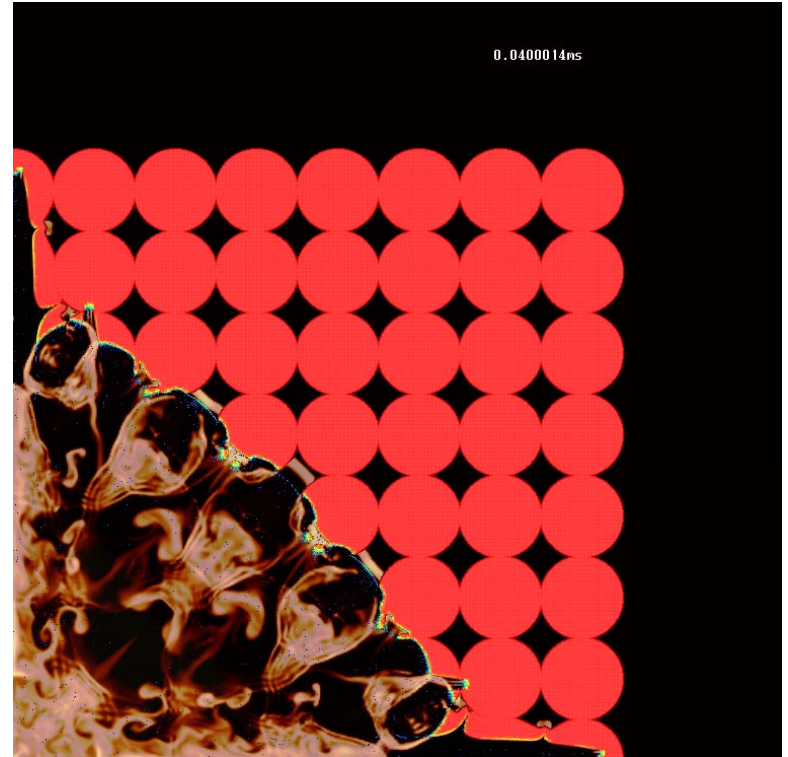
Wave moves at ~400m/s



540K value of T in corner

## DETONATION

Wave moves 8500m/s



5Gpa P in corner

How do we go from this to this?

# DDT: 3 Phase model

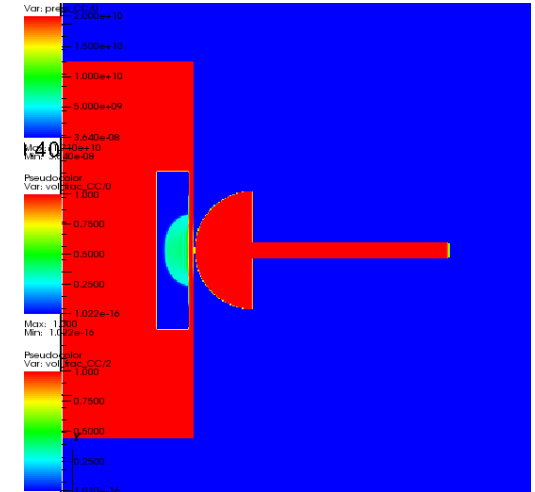
Stevens Test; data 72m/s  
Simulation 65m/s

## 1 WSB two phase burning model

Two-phase thermal decomposition- reactants--  
MPM materials, products--ICE materials  
verified [Atwood]

## 2 Fast convective burning model

Impactor causes initial shock wave and hot spot/core collapse causing second shock which acts as a virtual piston which causes pressurization up to a critical threshold 5.0 GPa [Souers] activated by cracking [Berghout et al.] Pop plot shows correlation between LLNL experiments and computation.

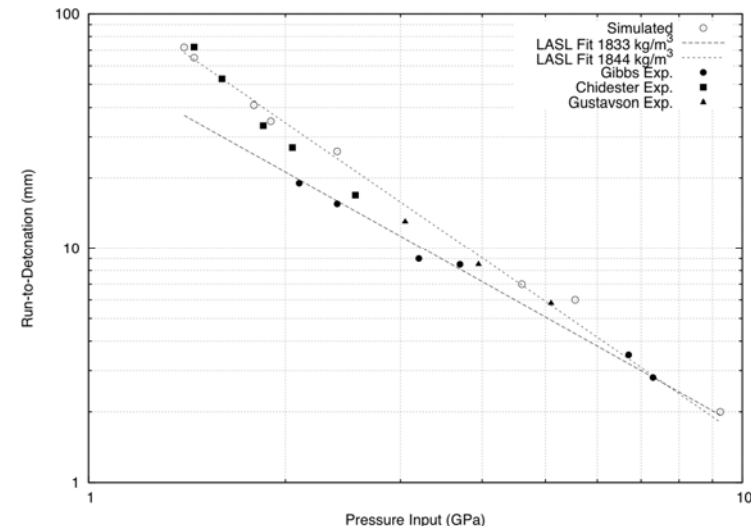


Pop Plot

## 3 JWLL++ detonation model

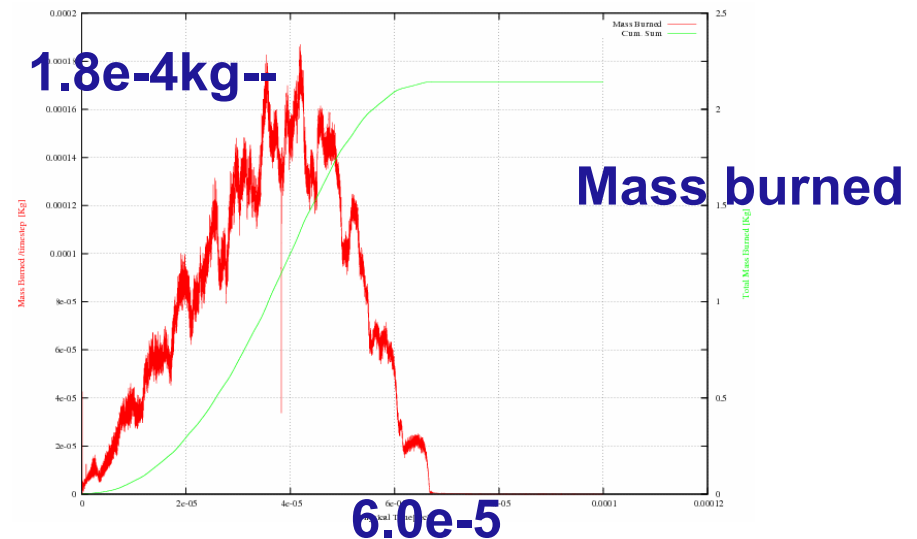
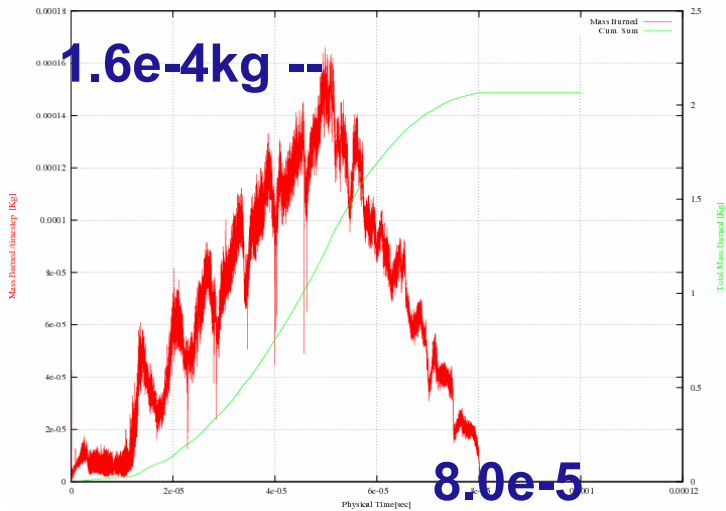
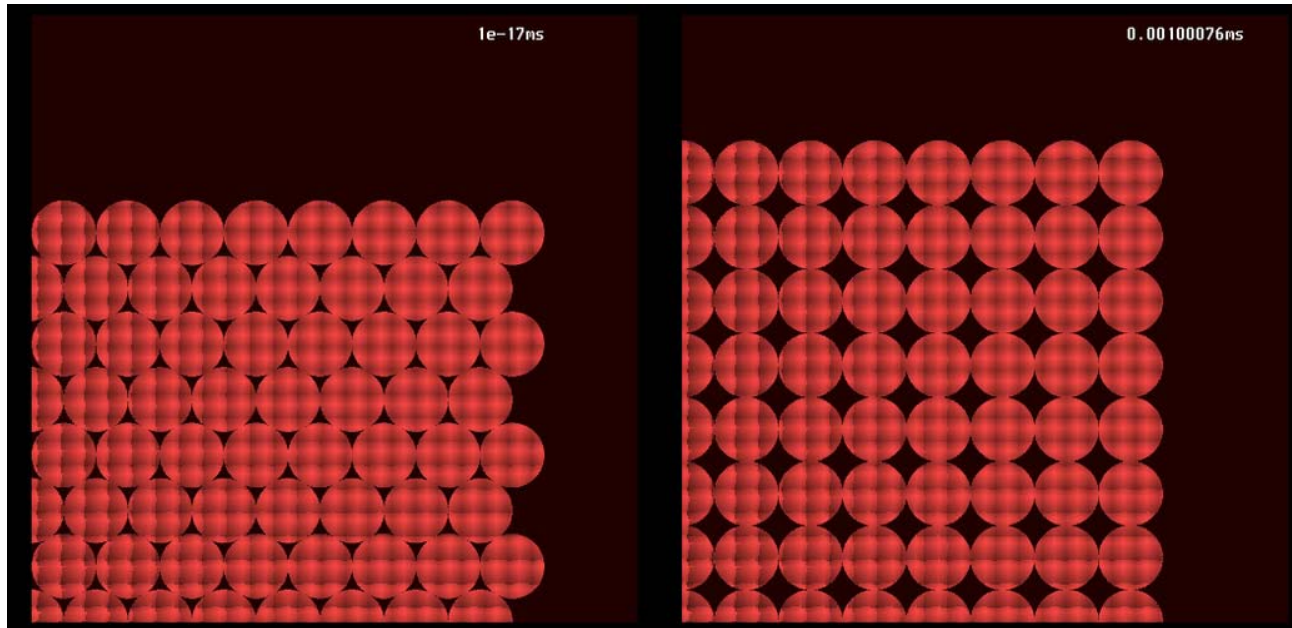
$\rho = \rho_0 \times (P/K + (1 + (P/K)^2)^{1/2})$ ,  $K = 11.4$  GPa  
and  $\rho_0 = 1840$  kg/m<sup>3</sup>

Gives good agreement with HMX explosive  
activation energy [German et al.]

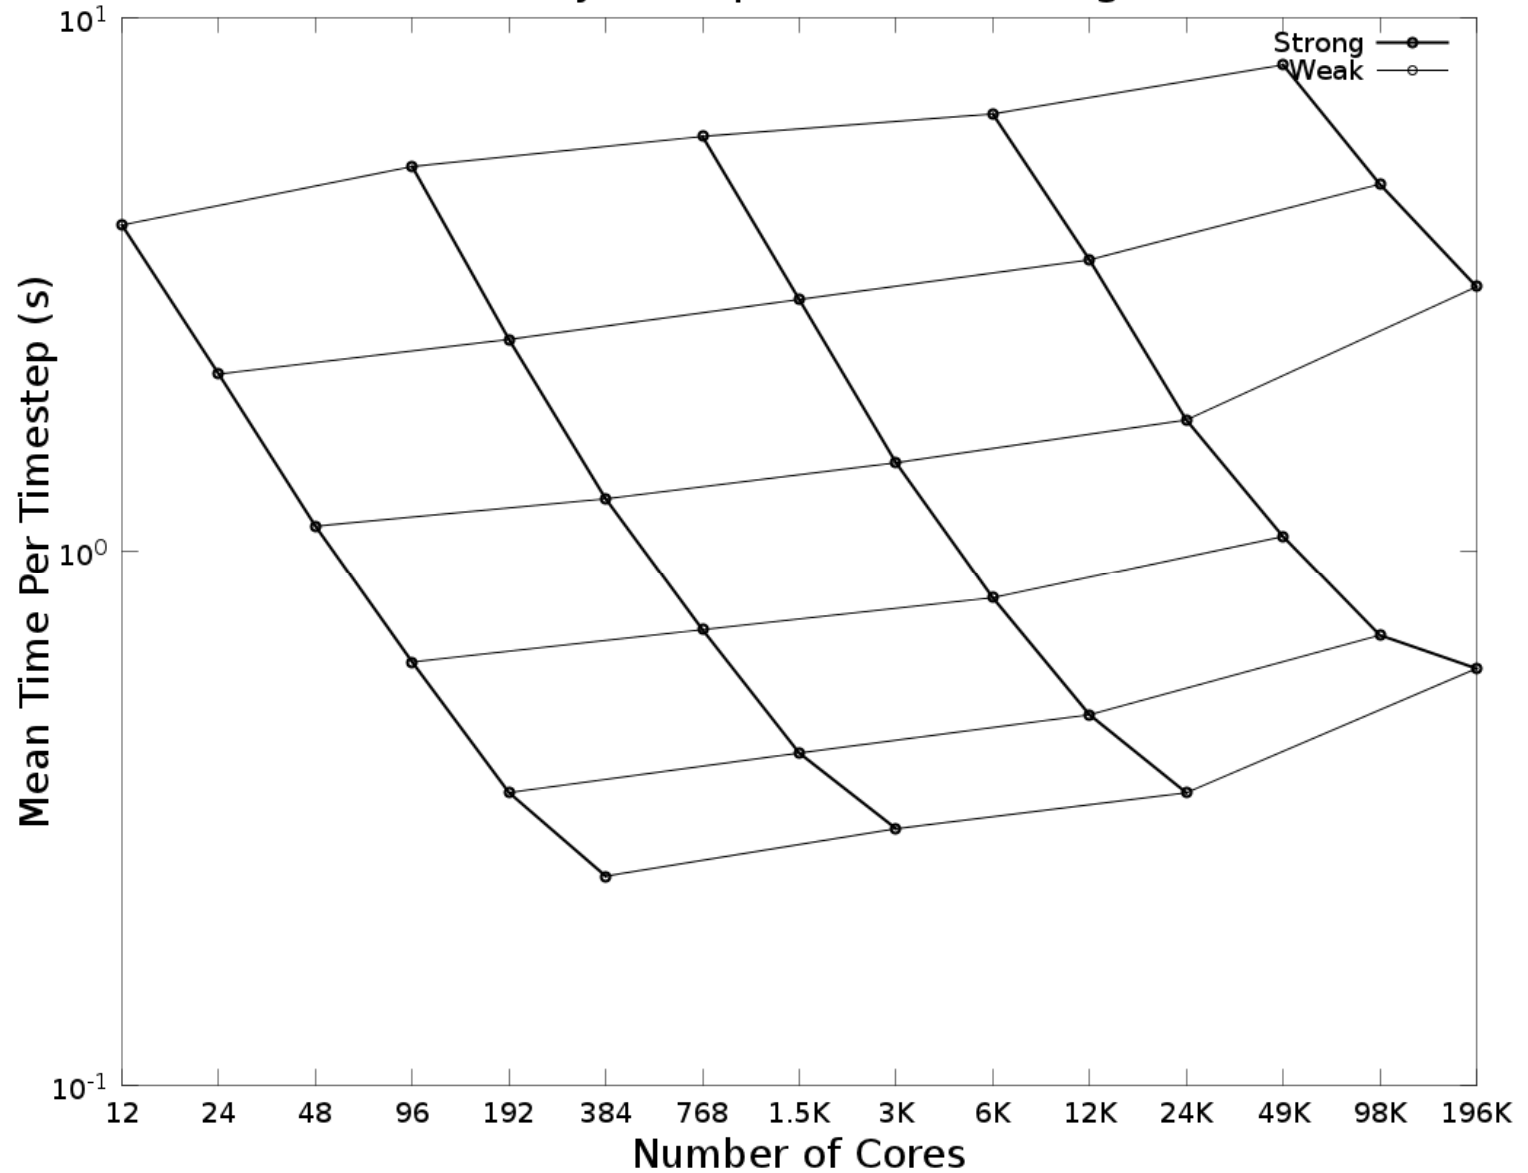




# Loose and Tightly packed Containers



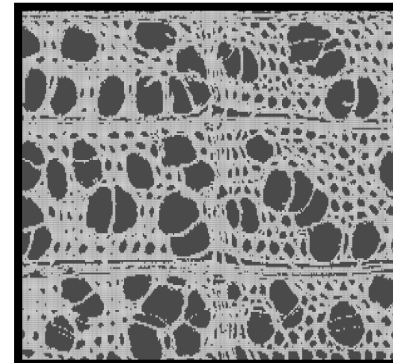
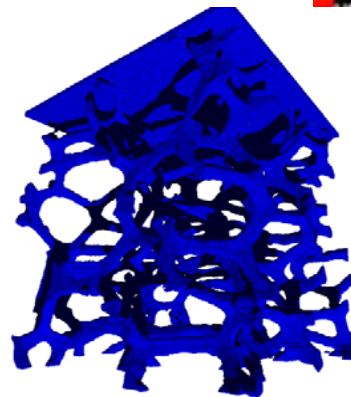
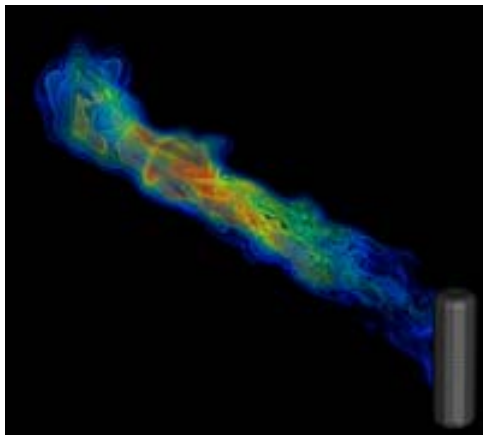
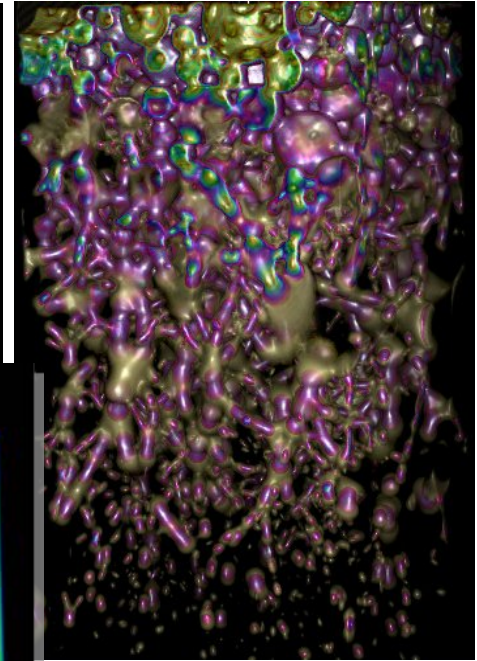
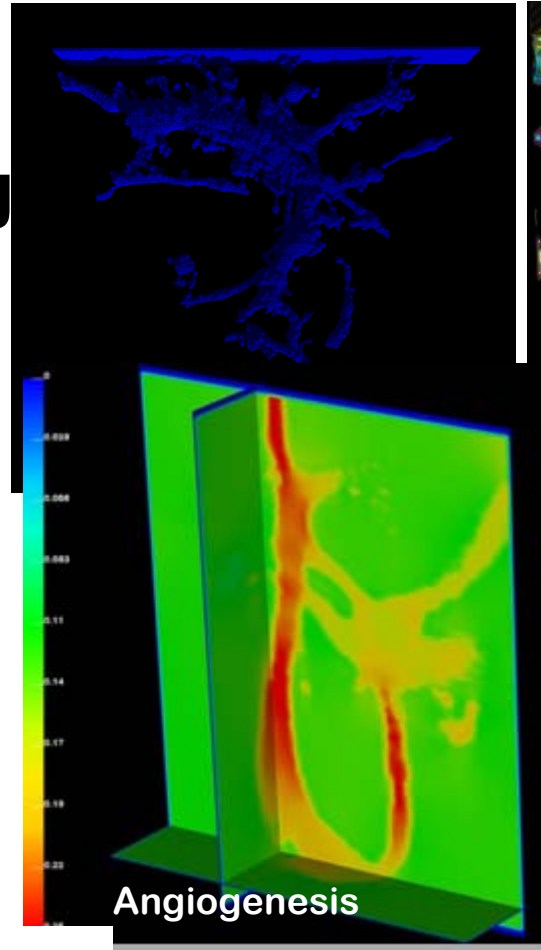
# Array of Explosives: Scaling



**Five cases  $64^3$ ,  $128^3$ ,  $256^3$   $512^3$   $1024^3$  meshes**

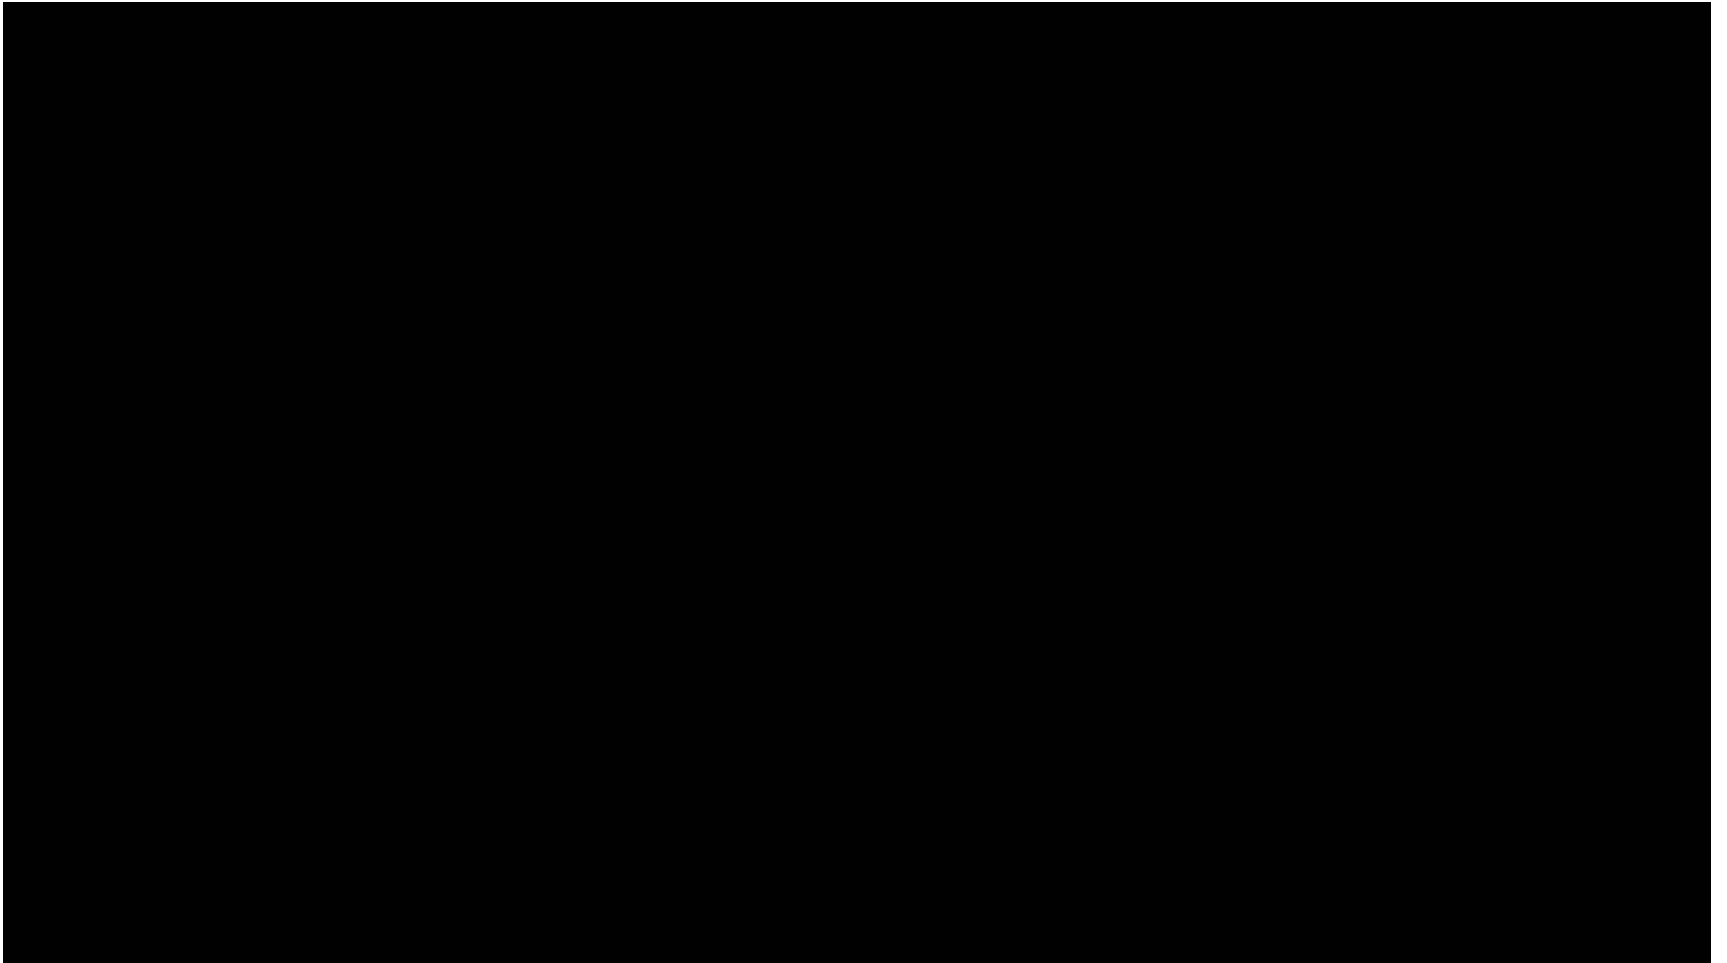
# Uintah Applications

- Flare Simulation
- Angiogenesis
- Vocal chord modeling
- Rocket stage separation
- Heart injury modeling
- Foam properties
- Granular flow



# Shape Charge Simulation Guilkey Harman and Brannon

Metal casing copper liner deforms into  
penetrating hypersonic jet



# The Algorithms will have to change

Consider stencil pde dx,dy,dz and dt with order p

Key metric is work per digits of accuracy in quantity of interest?

$$\mathbf{work} = K p n_x n_y n_z n_t ,$$

K = stencil work constant

**Pts in space and time**

Reducing the error by a factor of  $f$

$$\mathbf{Work}_{\text{new}} = \mathbf{work}_{\text{old}} f^{\left(\frac{4}{p}\right)}$$

**If  $p=1$  and  $f =2$  first order, doubling accuracy requires 16x work**

**If  $p=4$  and  $f =2$  fourth order, doubling accuracy requires 2x work**

# The Algorithms will have to change

High order methods are widely used in some areas but...

Solution of hard problems sometimes suggests that they offer fewer advantages—

**Non-oscillatory approximation of a front with order  $p$  requires  $(p+1)$  points in the front.**

**Mesh refinement needs to be combined with high-order methods**

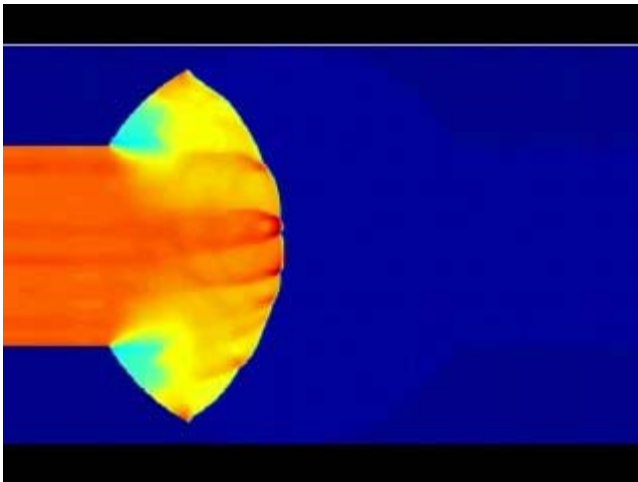
**Error in quantity of interest requires adjoint-based remeshing methods **not just gradient-based****

**Scalable Linear algebra solvers for exascale are a daunting challenge but one that is being tackled....**

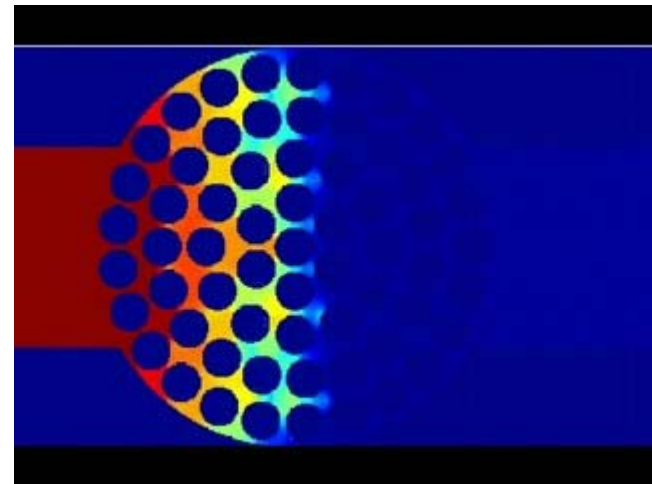
# Summary

- Uintah version of silver model works
- Silver model is not a silver bullet
- Uintah runs on 196K cores in AMR and fixed grid modes
- Scalability aspects of Uintah investigated
- Work needed to move to full DDT model
- Need to start work on detonation reduction approaches

Detonation Diffracting  
into cylindrical volume



Detonation disrupted by  
cylindrical rods



Source DDT experts