Investigation on Scaling Performance and Mesh Convergence with Sandia's ASC SIERRA/Fuego code for Fire Model Predictions of Heat Flux

Mahesh Rajan, Amalia Black, Stefan Domino

Sandia National Laboratories

Albuquerque, NM 87185, USA

Presented at the Cray Users Group (CUG) Meeting

May 8-11, 2006

Lugano, Switzerland

This work was supported in part by the U.S. Department of Energy. Sandia is a multi program laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States National Nuclear Security Administration and the Department of Energy under contract DE-AC04-94AL85000.

Fuego

Abnormal Environment Analysis

Motivation:

Safety qualification of weapons and components

Advanced Strategic Computing (ASC) program directed effort to establish capability to analyze models of adequate geometric and physics fidelity to supplement the experiment-based qualification

An example of capability class simulations on Sandia's Red Storm/Cray XT3 system

Code Coupling: Fluid Mechanics, Combustion & Heat Transfer

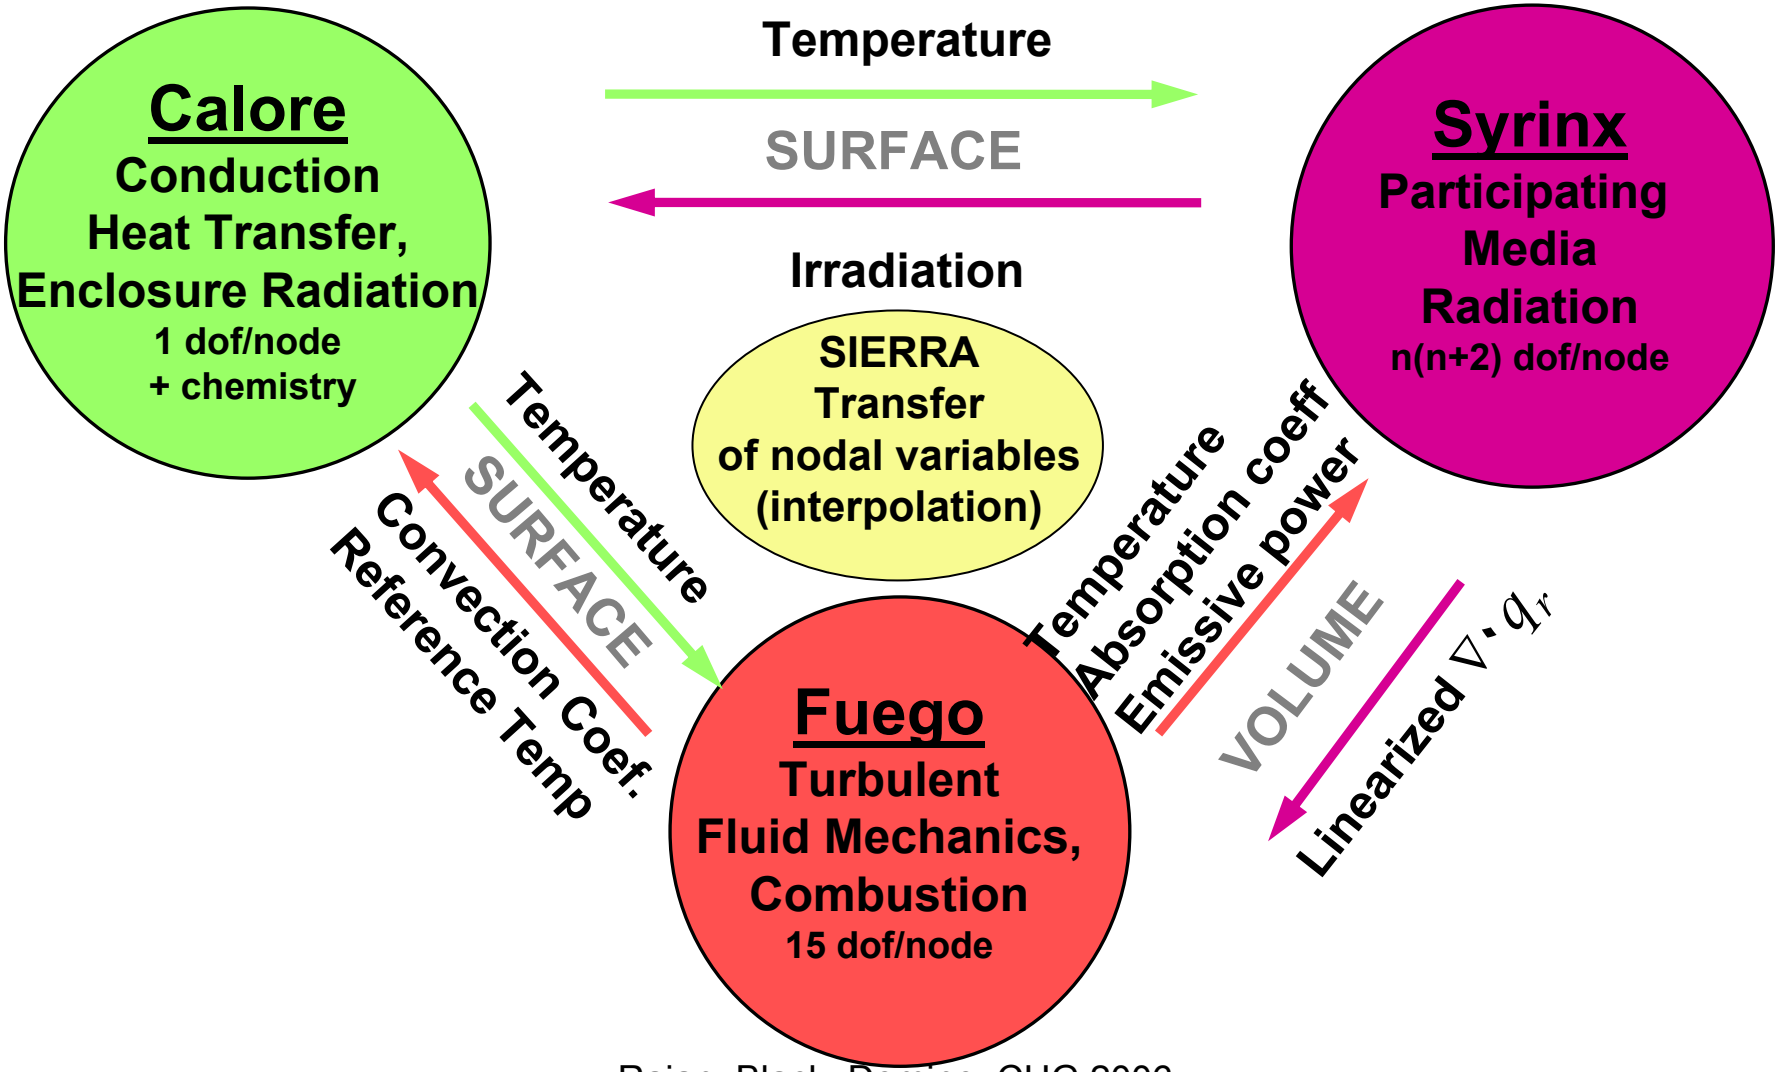

Rajan, Black, Domino; CUG 2006;
5/09/06

Red Storm used for Capability Class Fire Simulation

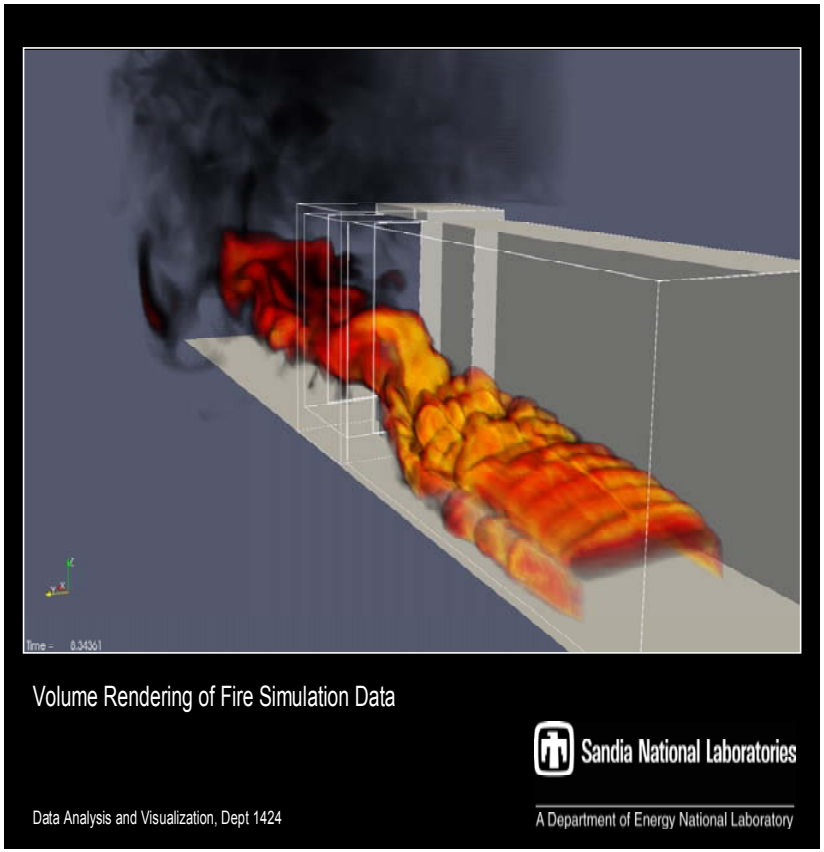

- Largest Fire Simulations (150 Million DOF
- Coupled calculations Fuego/Calore/Syrinx on 1024, 2048 and (limited number) 5200 processor Red Storm to demonstrate fire and thermal response calculations relevant to qualification of the W76 and W80
- Runs over several days with restart
- 400GB of data transferred to Viz cluster for parallel volume rendering using ParaView

Mesh Convergence Investigation

- **Coupled fire/thermal response predictions for a weapon-like calorimeter is validated for a quiescent fire representative of a transportation accident scenario.**
- **The model constructed was used to compare numerical predictions against experimental data.**
- **Temperature measurements were used to validate the coupled Fuego/Syrinx/Calore predictions**
- **A mesh convergence study using coarse (599K nodes), medium (1.1M nodes) and fine (2.4M nodes) meshes for the Fuego/Syrinx calculations was conducted.**
- **Run on 256, 512 and 1024 processors each for ~3, ~7, ~15 days**
- **The numerical results (coarse, medium and fine) used a time-filtered Navier Stokes (TFNS) turbulence treatment which produces a time varying solution. The TFNS results were time averaged over the 10-30 second interval in order to produce average vertical velocities**

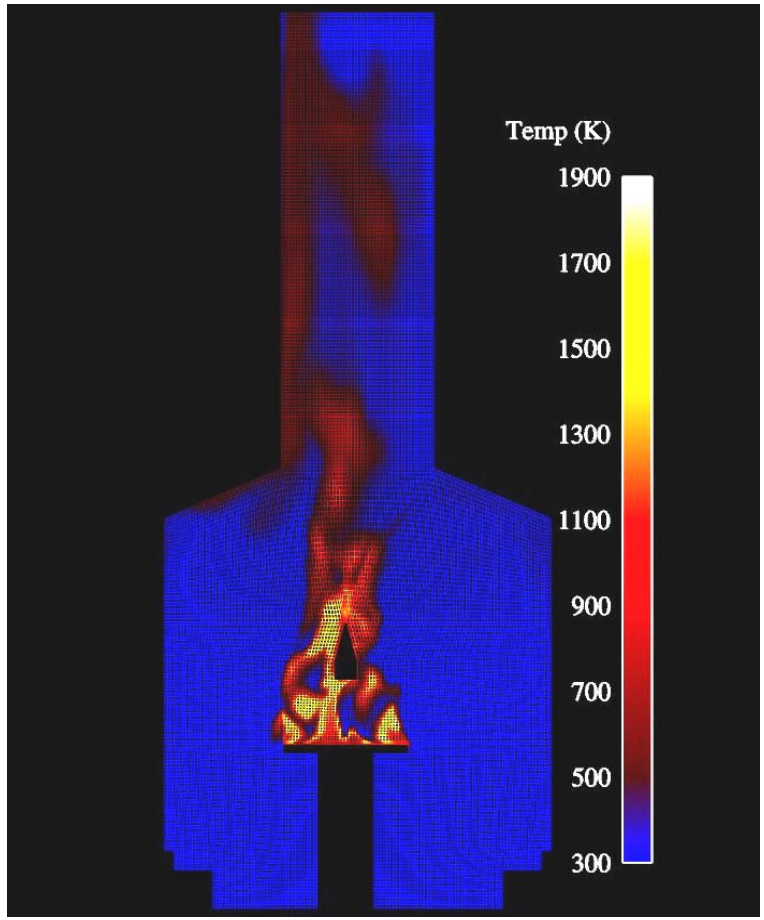

FLAME Facility – Coarse Mesh Along the Burner Centerline with Temperature (K) Contours Overlaid

All three Meshes Shows Similar Vertical Velocity

- **The vertical distance shown is along the centerline between the burner surface and the base of the calorimeter.**
- **All three mesh results show similar values up to 2.8 m and reach peak velocity values at the same vertical location (3.31 m)**
- **Beyond a distance of 2.8 m, the coarse mesh result and medium mesh result are very similar and only differ by up to 8%; whereas, the fine mesh result differs by up to 25% in vertical velocity..**
- **Without the fine mesh result, which was only possible because of the availability of large computing resources, the numerical solution may have appeared to be converging.**
- **A finer mesh result is still necessary to confirm the grid independence of the solution but this will require substantially more computational resources and execution time.**

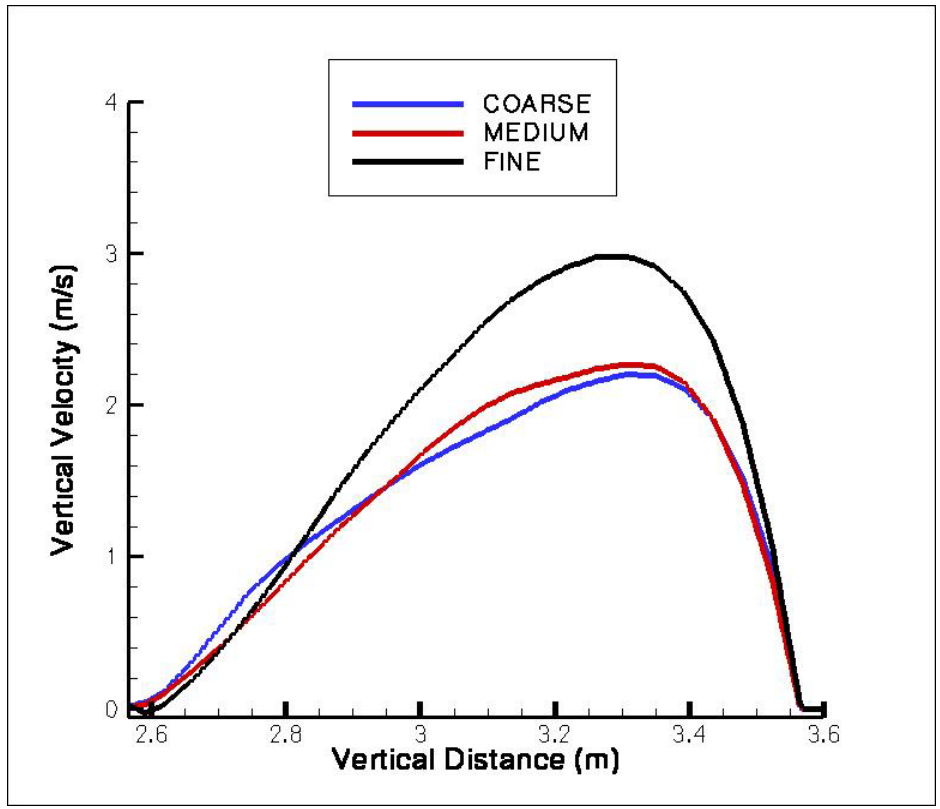

Comparison of Three Mesh Results for Average Vertical Velocity along the Burner Centerline

Performance Scaling Analysis

- Leverage the four meshes used in mesh convergence investigation; ‘Strong Scaling’ runs with each mesh
- The objectives:
  - determine the optimal number of elements per node
  - analyze SIERRA Framework scaling separate from the linear solver scaling
  - Identify scaling inhibitors.

Optimal Processor configuration based on knee of the run time curve:

Coarse=128, Medium=256, Fine=512 and Very-Fine=1024.

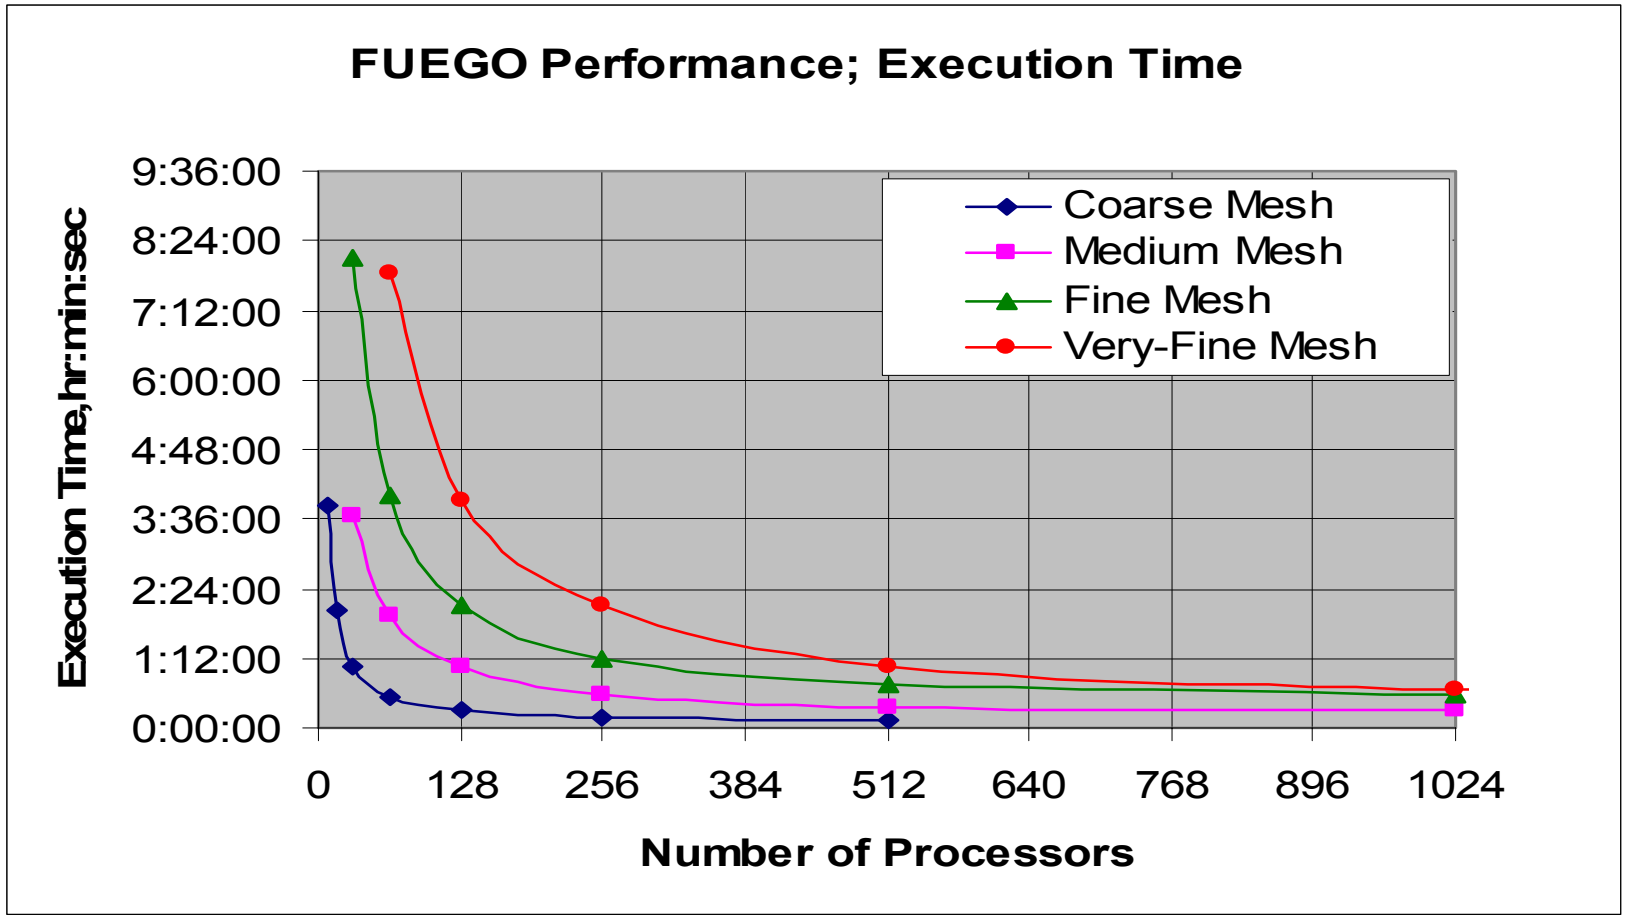

Rajan, Black, Domino; CUG 2006;
5/09/06

The Important Fluid Region Computation shows 74% efficiency at 1024 processors with the Very-Fine Mesh

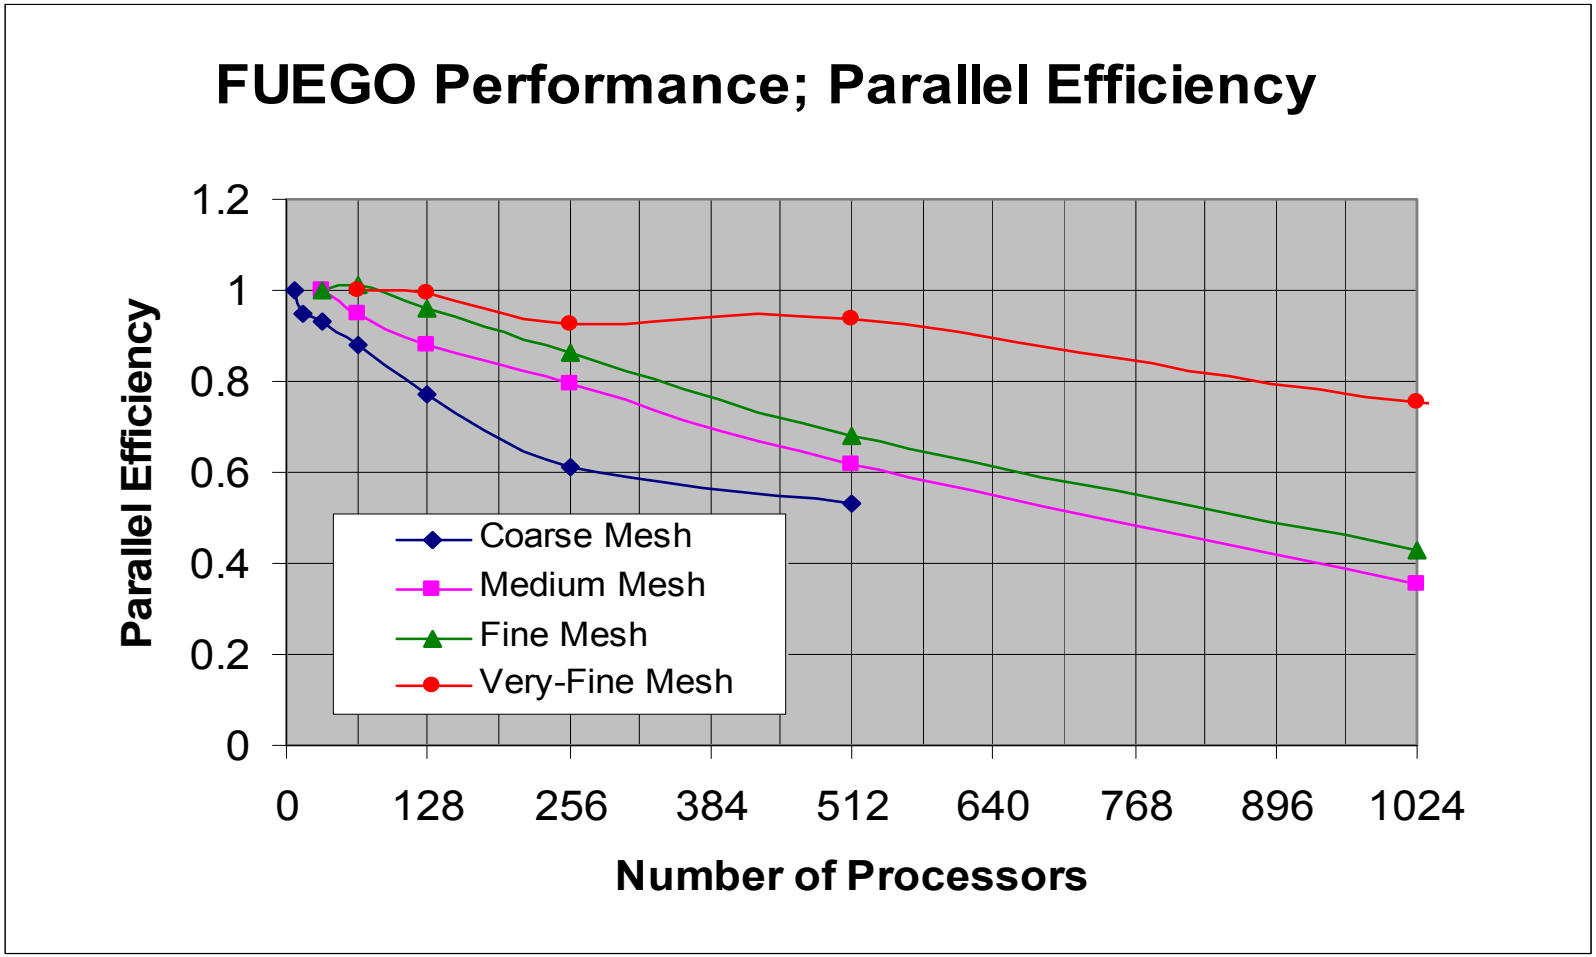

Rajan, Black, Domino; CUG 2006;
5/09/06

Perfect Scaling Observed for Matrix Assembly; Matrix Solve Needs to be Improved

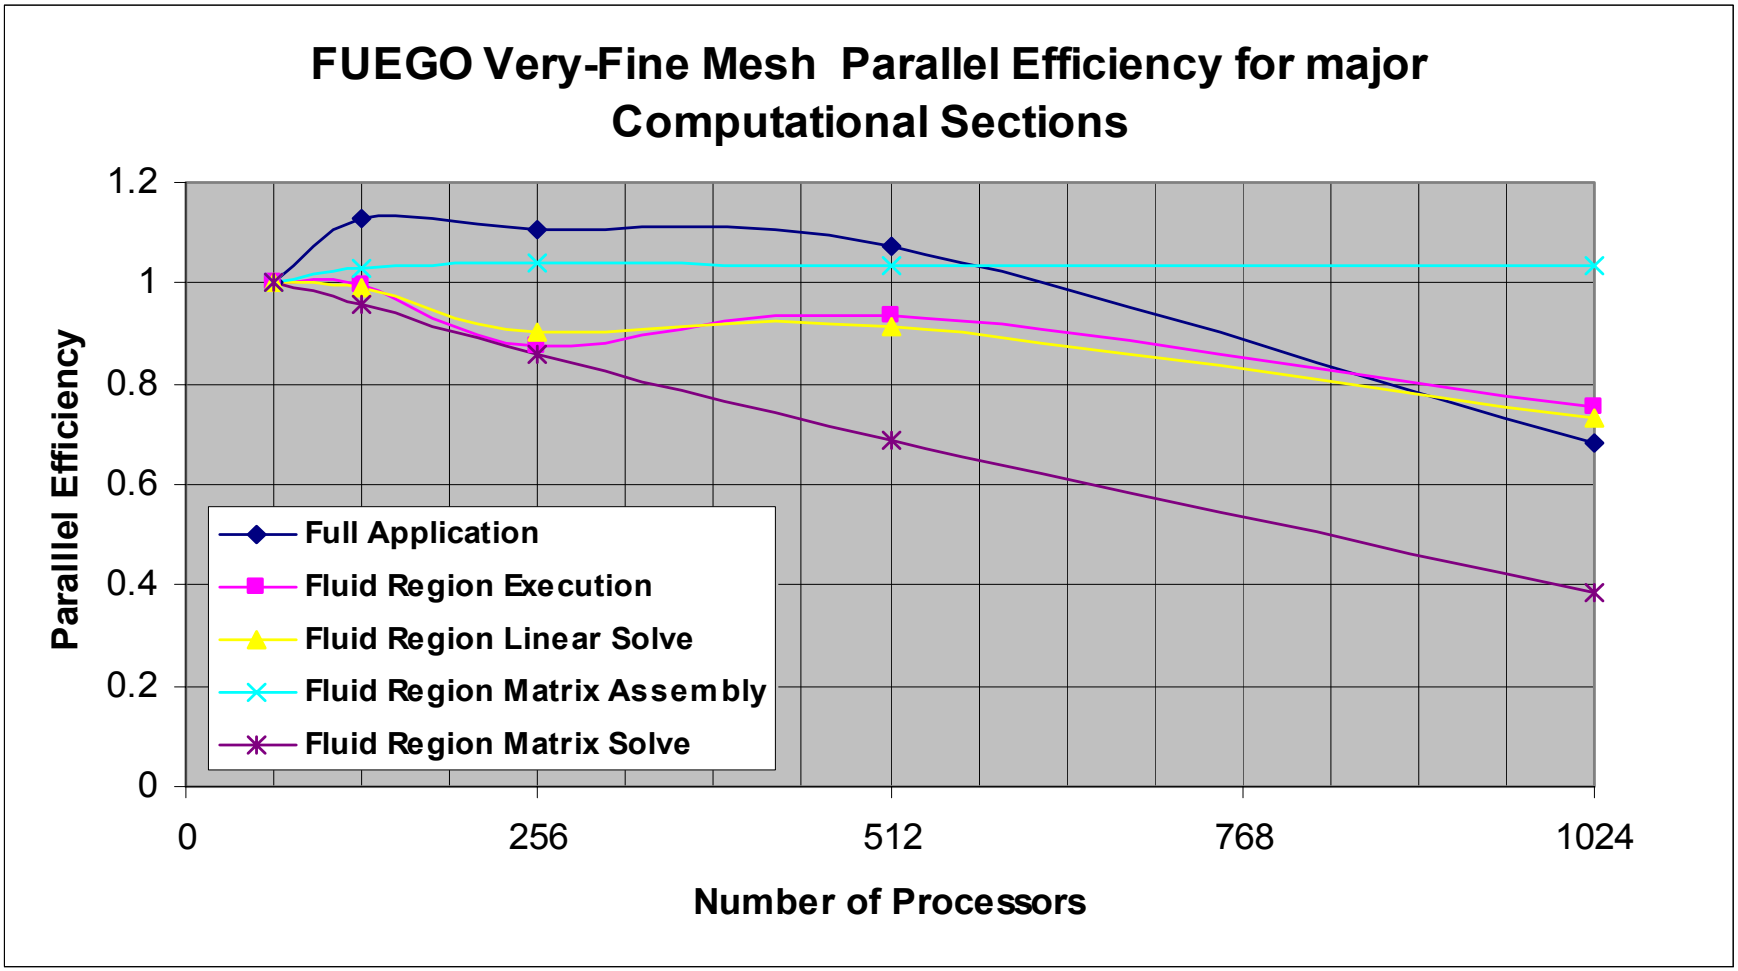

Rajan, Black, Domino; CUG 2006; 5/09/06

Analyze Performance Inhibitors

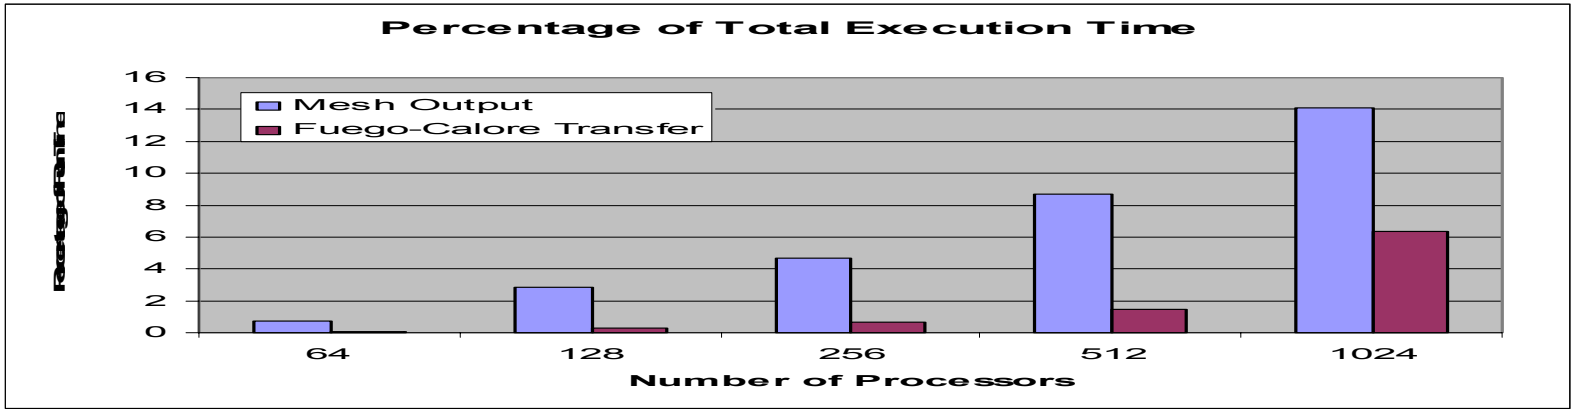

I/O Overhead:

- **Used Lustre file system**
- **Mesh Results output directed to 32 I/O servers and 64 storage targets (stripe count=64, stripe size=1MB)**
- **Output consisted of 4 files per processor with a total output of approximately 1.6GB per output.**
- **Ideally percentage time in Mesh Output should be flat**
- **Solution: When using 1 file per processor use 'lfs setstripe' to set the stripe count to 1; In Progress!**

Fuego-Calore Transfer Overhead:

- **The reasons for this may likely be related to the disparity in element count between the fluids and conduction region. Specifically, at 1024 processors the calore mesh has only O(50*nodesPerElement) sends per processor.**

CONCLUSIONS

- Fire simulations for mesh convergence investigations with runs lasting several days to resolve the fire to 30 seconds of simulation time have been routinely carried out on Red Storm
- Further refined mesh runs are planned and would not be possible without a capability class system like Red Storm
- Performance analysis indicates FUEGO has impressive scaling to 1024 processor for the largest mesh.
  - Near perfect scaling of the Matrix Setup for the solve
  - Linear Solver scaling could be improved and is to be investigated
- For applications like FUEGO where mesh output is one file per processor, Lustre file with stripe count of 1-8 may prevent I/O overhead growth
- Additional scaling inhibitors have been identified that may impact planned large (> 4000) processor count runs. Further instrumentation is needed to locate and remove them.