SAND2015-0676C

Exceptional service in the national interest

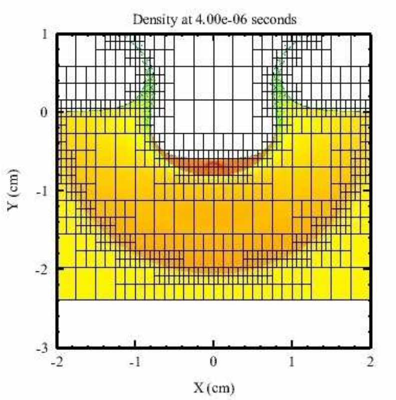

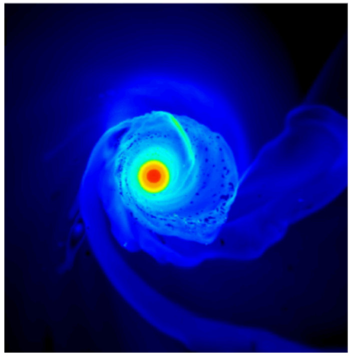

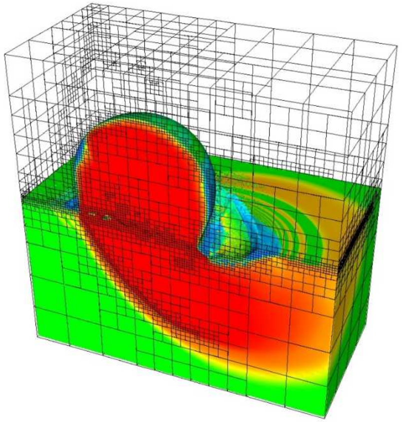

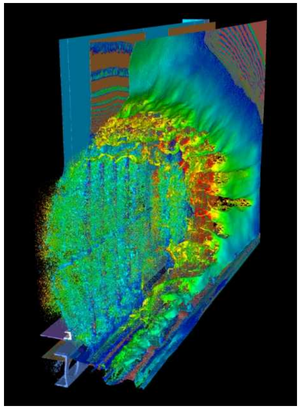

A Decade of Adaptive Mesh Refinement in CTH

Applied Computer Science Meeting

2-5 February 2015

David Crawford

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.

CTH Overview

CTH is a massively-parallel shock-physics code

- Explicit Eulerian shock-physics code (hydrocode)
  - Solves conservation equations of mass, momentum, and energy
  - Up to 98 simultaneous materials
  - Gases, fluids, solids, reactive materials
  - Analytic & Tabular Equation-of-State representations
  - Advanced Strength & Fracture models
  - Adaptive Mesh Refinement (AMR)
- Applications (partial list):
  - Armor, Anti-Armor, Conventional Munitions Design, Blast Effects
  - Planetary Science, Asteroid Impact & Planetary Defense
- CTH licensed to U.S. government agencies and their subcontractors and U.S. academic institutions
  - 600+ users
- www.sandia.gov/CTH

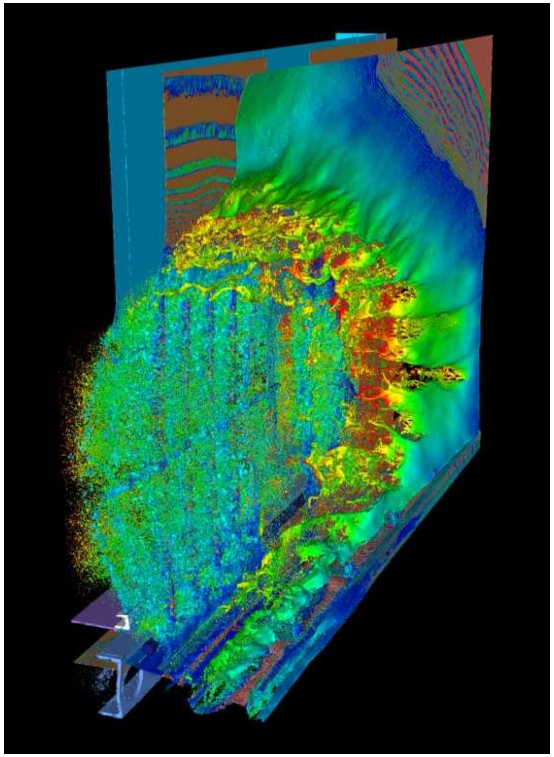

32,000 processor calculation showing nearby blast on aluminum and steel structure

Introduction

- Adaptive mesh refinement (AMR) was added to CTH in 1998-2001
  - Presented at NMH in Edinburgh (2002)
  - In 2002 we could see ~10x performance gain (vs. non-AMR) on the largest problems
- Today we routinely see 10x performance gains…
  - …and an extrapolation using Moore's Law suggests we should see 20-30x today on the largest problems.
  - (about a factor of two for every ten years)
- *However, we occasionally see 200-300x on the largest problems…why?*

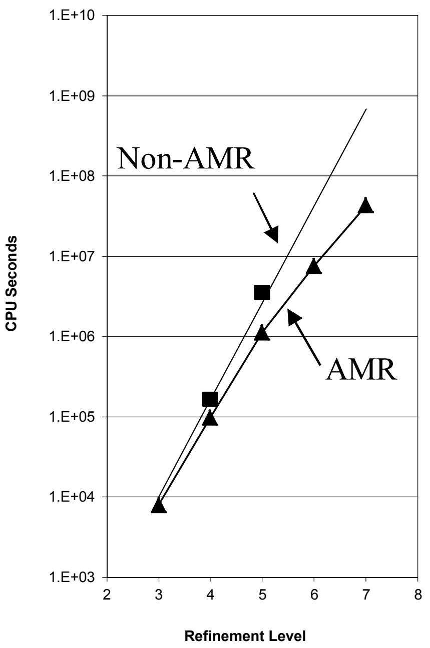

CTH performance on a heroic problem in 2001

Adaptive Mesh Refinement in CTH

we use a simple approach

- Block-based
- Identical logical mesh per block
  - 8x8x8 is typical
  - 10x10x10 with ghost cells
- Isotropic 2:1 refinement
- Single time-step for all blocks
- Load balance on per-block basis
- User-definable refinement indicators
- Problem initialized via iterative refinement/load balance step

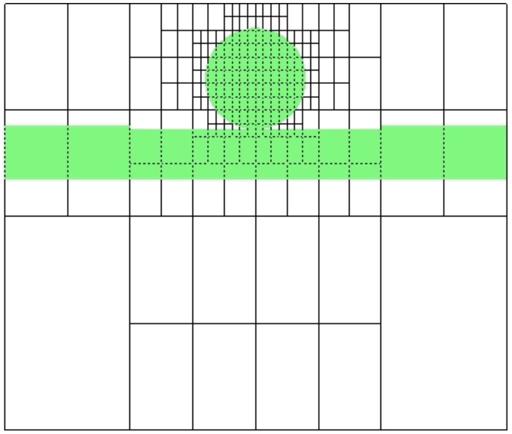

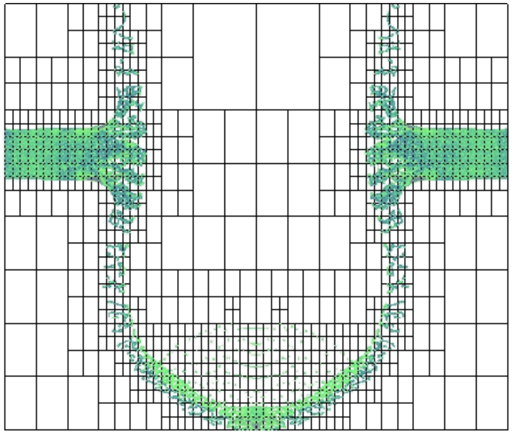

Achieving Parallel Performance *with our simple approach*

- Parallel process block refinement and unrefinement as much as possible by supercycling
  - Every 3 cycles for refinement
  - Every 6 cycles for unrefinement
- Perform load balancing only when disparity is more than 10%.
  - Smaller tolerance when memory resources are tight

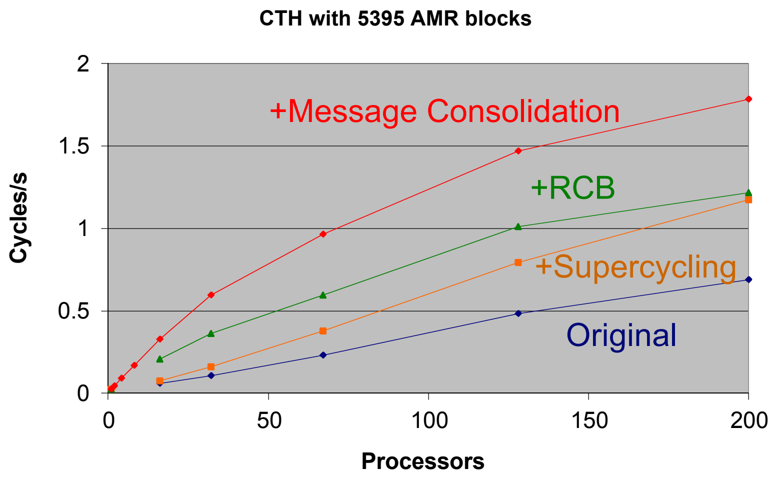

- Use Recursive-Coordinate-Bisection (RCB) algorithm to minimize off-processor communications.
- Consolidate message passing to reduce latency

AMR since 2002

- Surprisingly little has changed:
  - Added the ability to refine based on advanced material constitutive properties such as “damage”.
  - Added a spherical region to indicators.
  - Fixed some minor bugs appearing at boundaries
  - Increased number of allowable indicators (to 100 from 10)
- We provided sufficient flexibility in the original design
  - The “style of use” has changed
  - New guidelines for indicator design have evolved
- Flexible, user-defined indicators appear to be the most important aspect of AMR success in CTH

User-Defined Indicators

Where/When Refinement Will Happen

Indicator consists of:

1) Filters for materials, time, resolution or space
2) Operator
   val, abs, diff, grad
3) Database Field
4) Threshold(s)

```
indicator
    mat 2
    maxl -3
    p1 = 0, 5, 5
    p2 = 5, 15, 30
    val vmag
    refabove 10
endi
```

Hypothesized Formation of the Moon by Giant Impact early in Earth's history

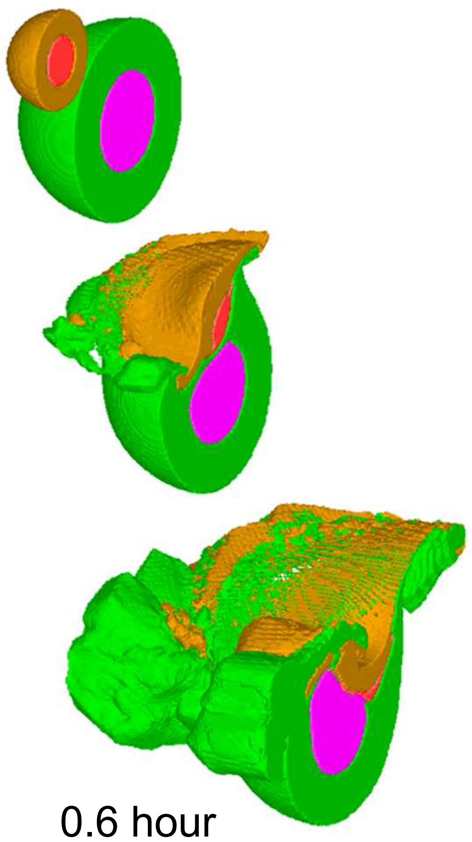

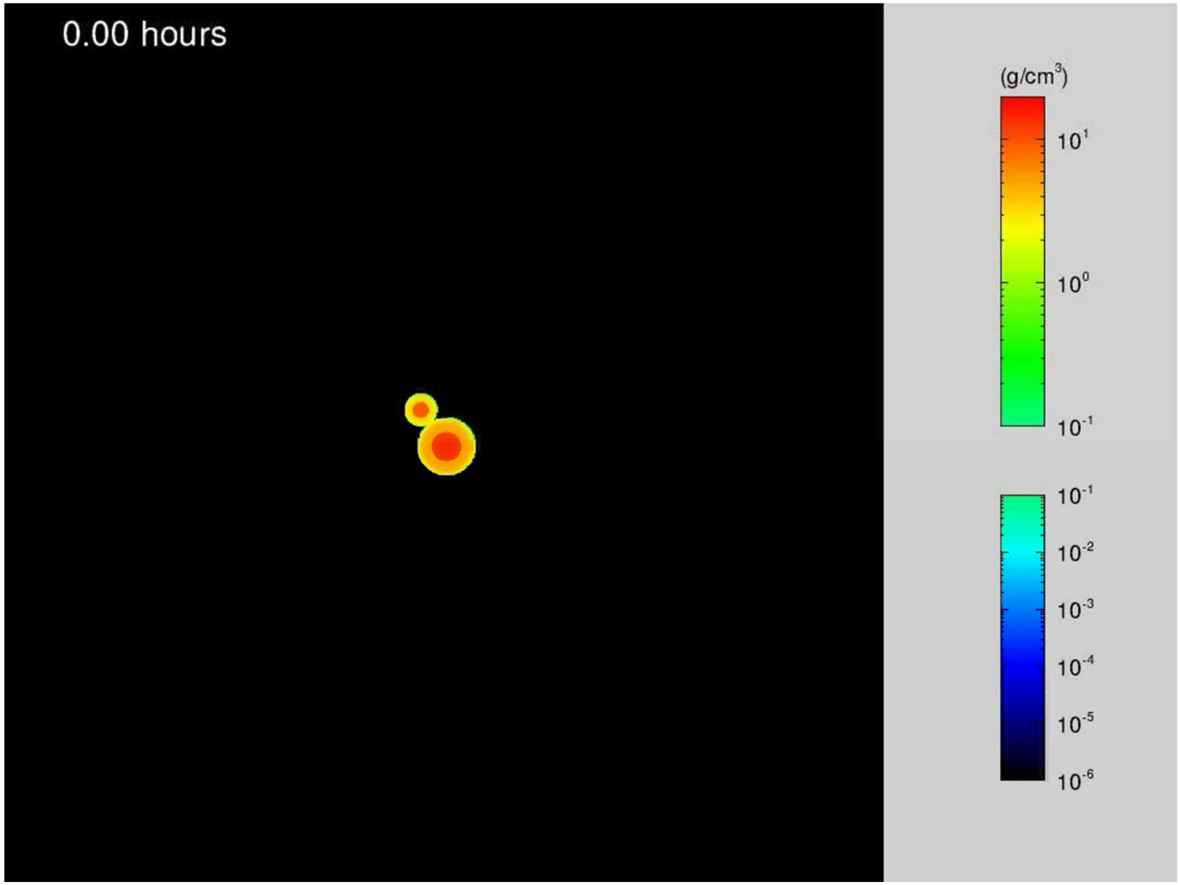

0.6-hour CTH simulation with central gravity, 3-million zones, Melosh and Kipp (1988)

54-hour, AMR-CTH simulation with self-gravity, 40 million zones, equivalent to 20 billion zones without AMR (2011)

- 500x memory gain, 200-300x performance gain

Giant Impacts

SPH simulations

- SPH has long history for these simulations
- Equal mass per particle
- $10^6 - 10^7$ particles provide adequate resolution

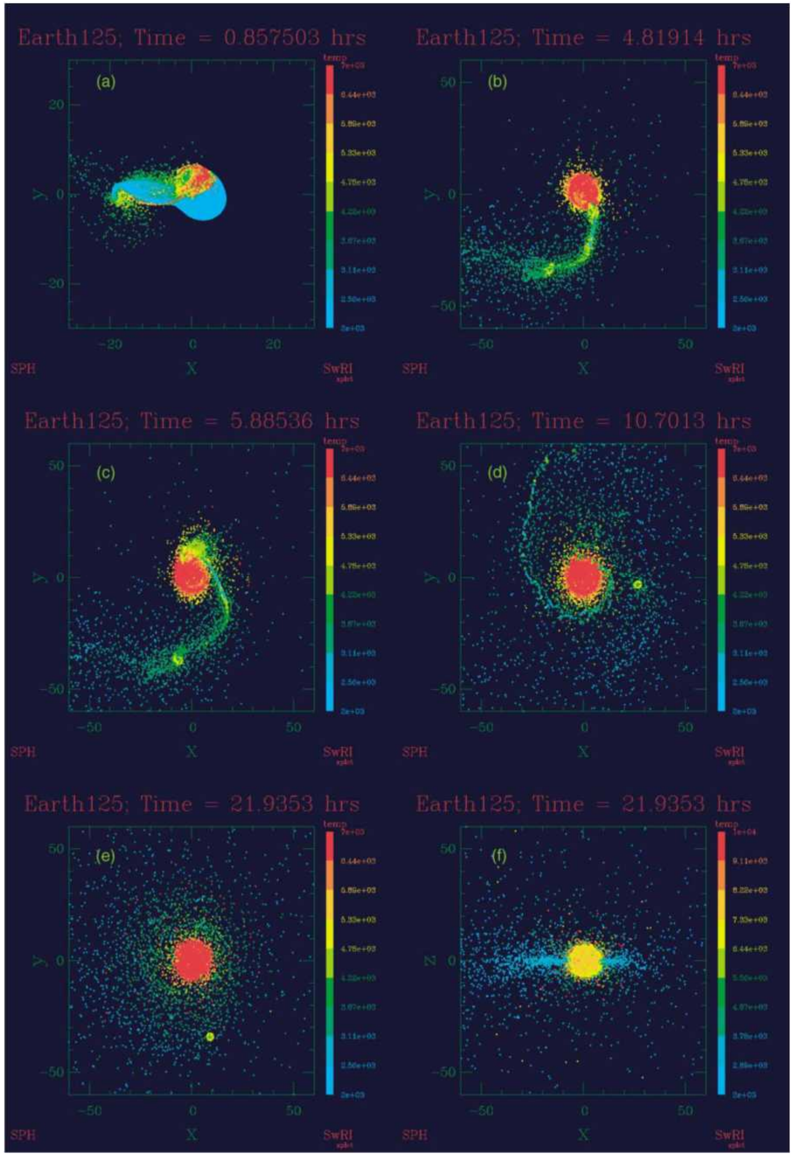

R. M. Canup, *Simulations of a late lunar-forming impact,* **Icarus: 168** *pp.* 433–456, 2004.

AMR indicators for equal mass approximation

Sandia National Laboratories

```
indicator
   val density
   refabove 1e-2
endi

indicator
   maxl -1
   val density
   refabove 1.25e-3
endi

…
```

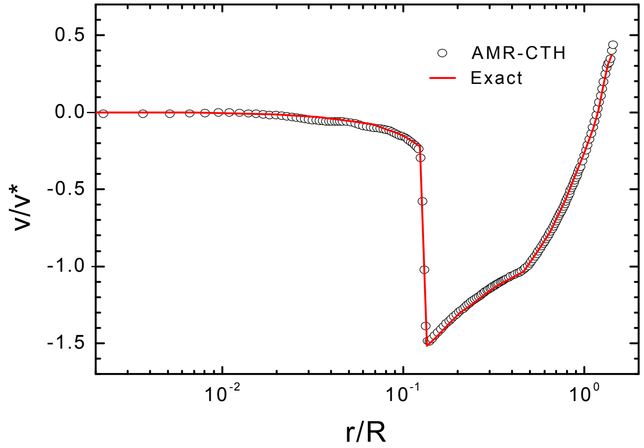

Verification: adiabatic collapse of an initially isothermal spherical gas cloud.

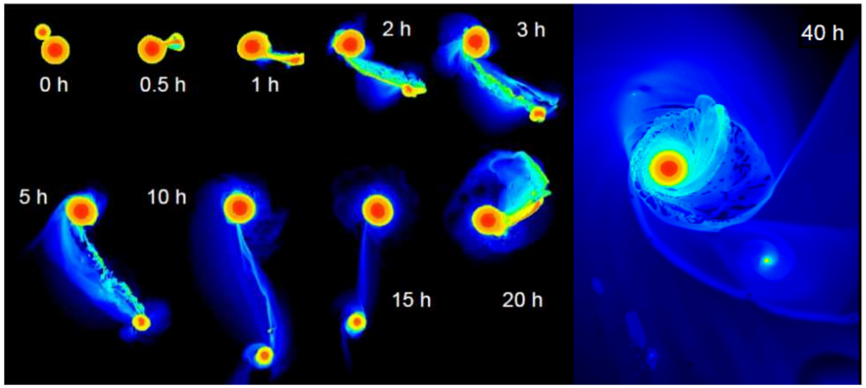

$10^6 – 10^7$ AMR zones provide adequate resolution for giant impact simulations

AMR-CTH vs. SPH for Giant Impacts

- We're seeing good agreement in answers
- *…and comparable performance…*

- R. M. Canup, A. C. Barr and D. A. Crawford, *Lunar-forming impacts: High-resolution SPH and AMR-CTH simulations,* **Icarus** (submitted)

An overlooked AMR scaling paradigm?

- Our 2002 view of AMR scaling:
  - High resolution mesh concentrated along sub-dimensional regions (curves in 2-D, surfaces in 3-D)

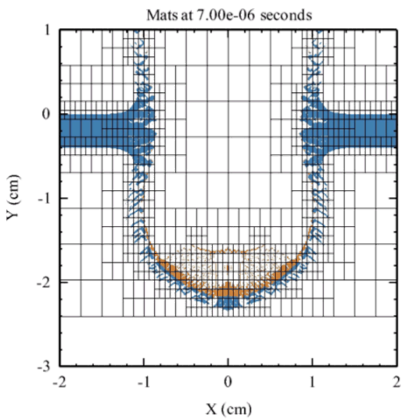

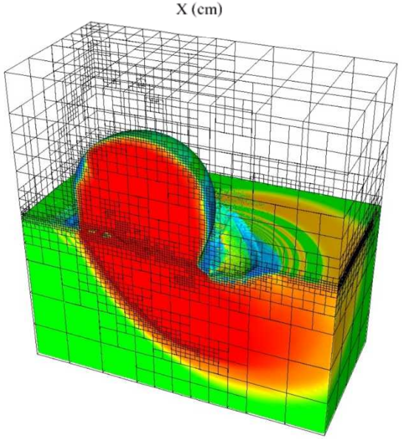

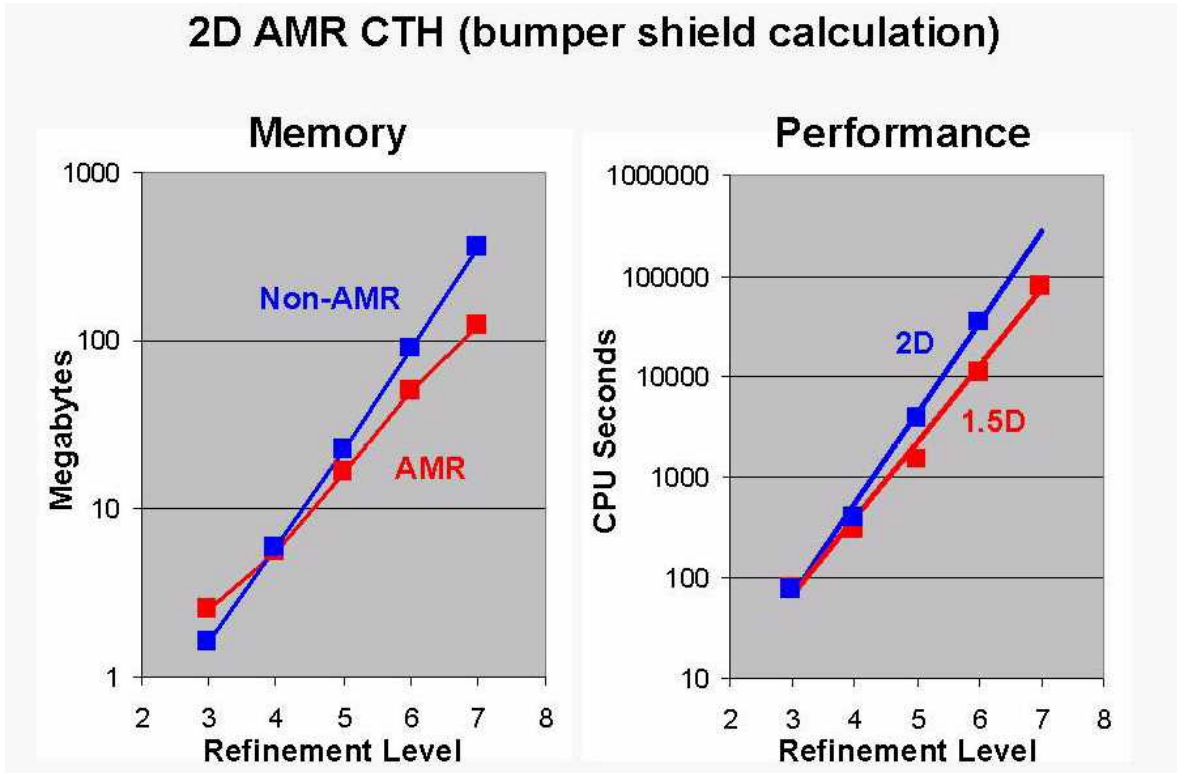

An overlooked AMR scaling paradigm (cont)?

- Indicators based on:
  - density or
  - presence of particular materials or
  - with spatial or temporal filters
- Can provide dramatic performance gains in many circumstances
- Without requiring detailed understanding of indicators tied to the underlying physics
- Many users apply the above strategy…*but how accurate is it?*

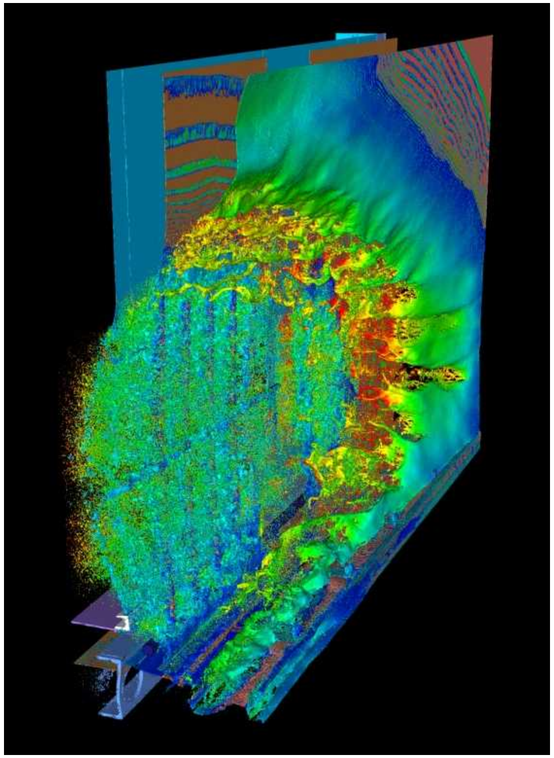

32,000 processor calculation showing nearby blast on aluminum and steel structure

Modified Sedov Blast Wave

- Approximation of blast/structure problem
  - High P gas sphere
  - Low P background gas
  - Modified to impinge shock on a thin spherical shell
  - Measure momentum of shell
- Use AMR to resolve different regions of the problem:
  - High P gas
  - Background gas
  - Thin shell

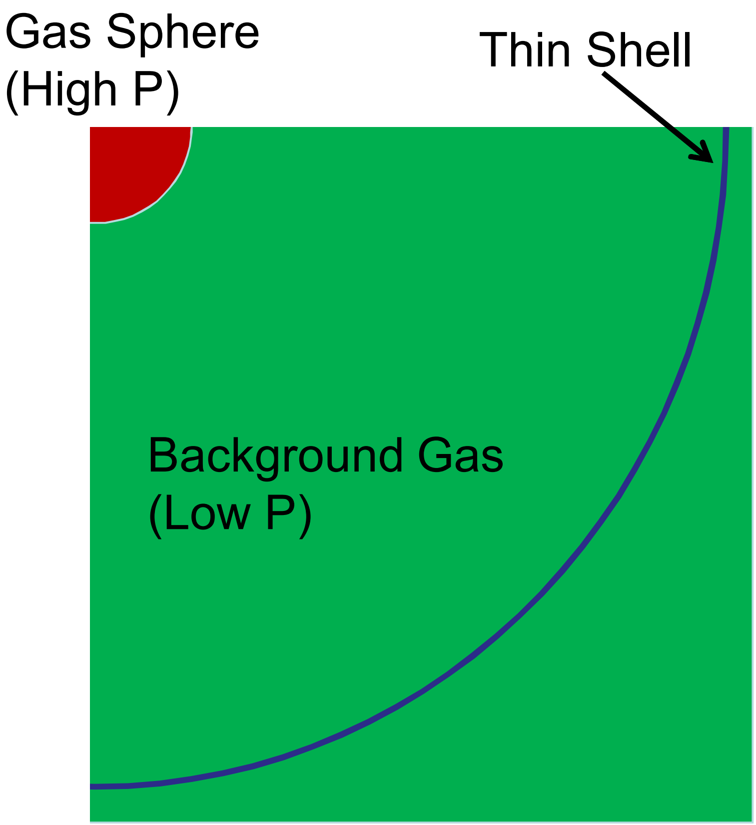

Modified Sedov Blast Wave (cont.)

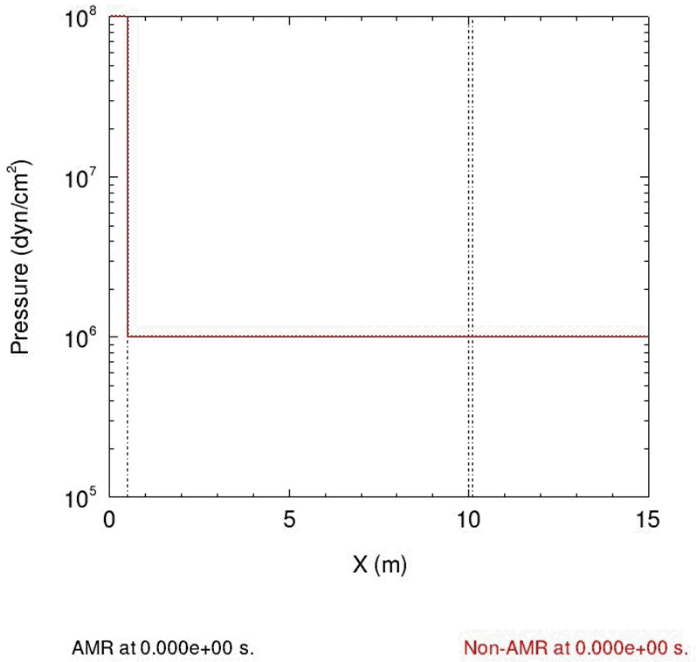

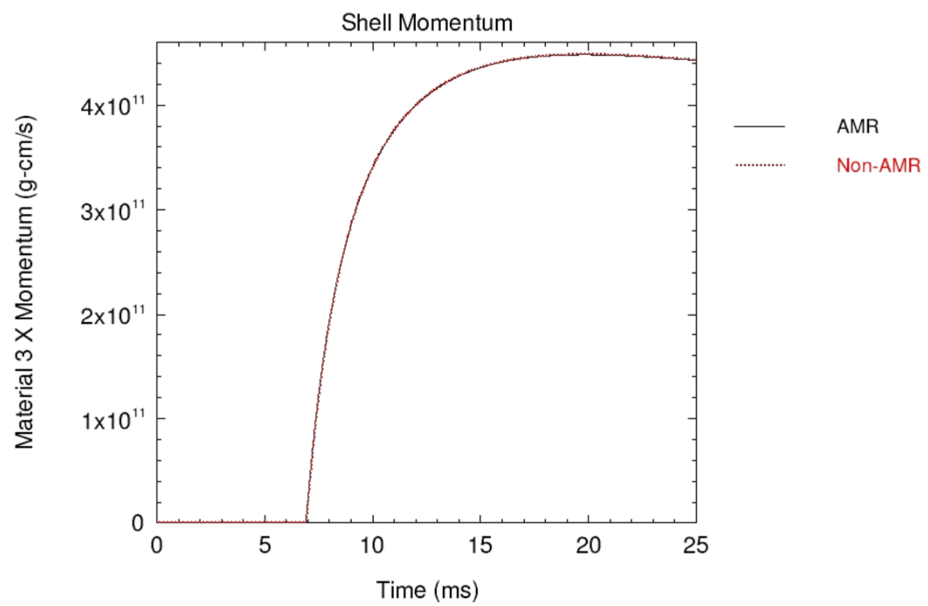

- 1 m sphere
- 10 cm shell
- 10 m standoff

- Uniform Resolution
- AMR & Non-AMR using same effective grid (1 cm)

Modified Sedov Blast Wave (cont.)

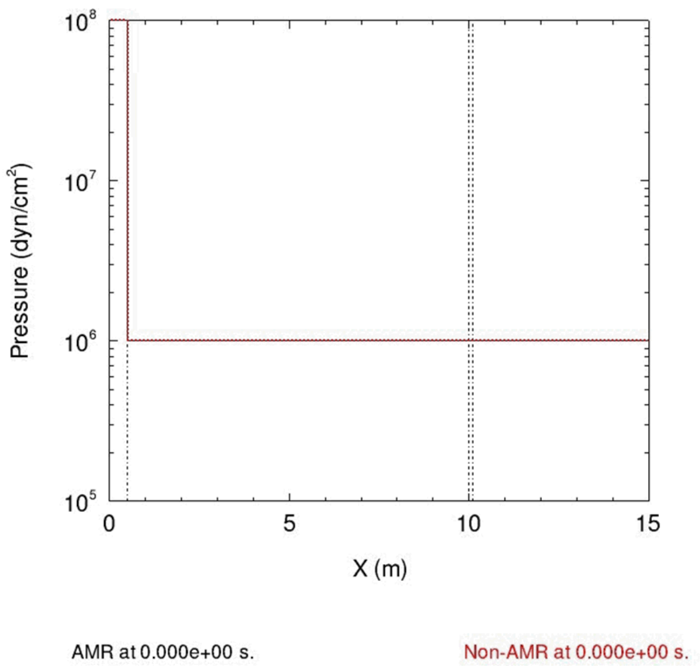

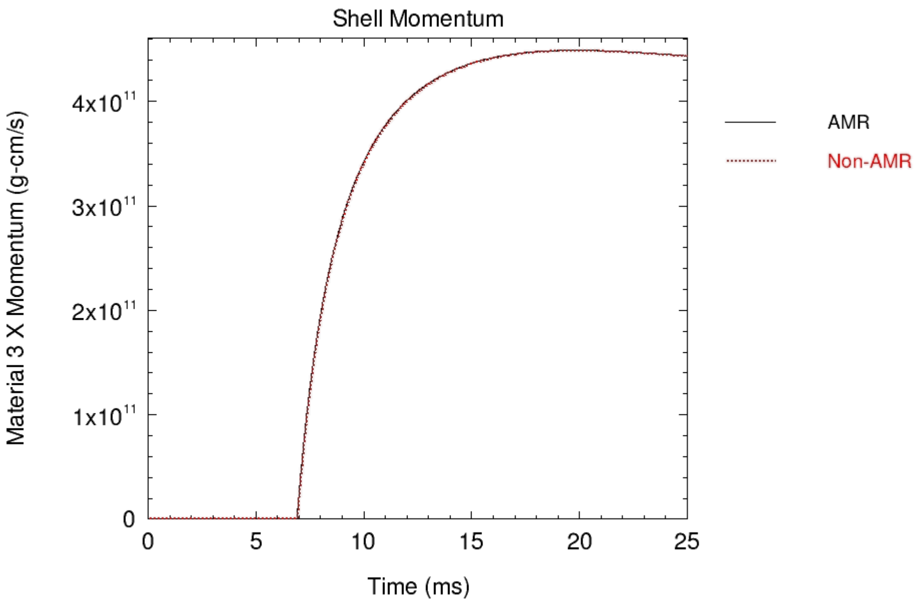

- 1 m sphere
- 10 cm shell
- 10 m standoff

- Non-AMR using uniform 1-cm grid
- AMR:
  - 0.5 cm grid on sphere
  - 1-cm grid on shell and background

Modified Sedov Blast Wave (cont.)

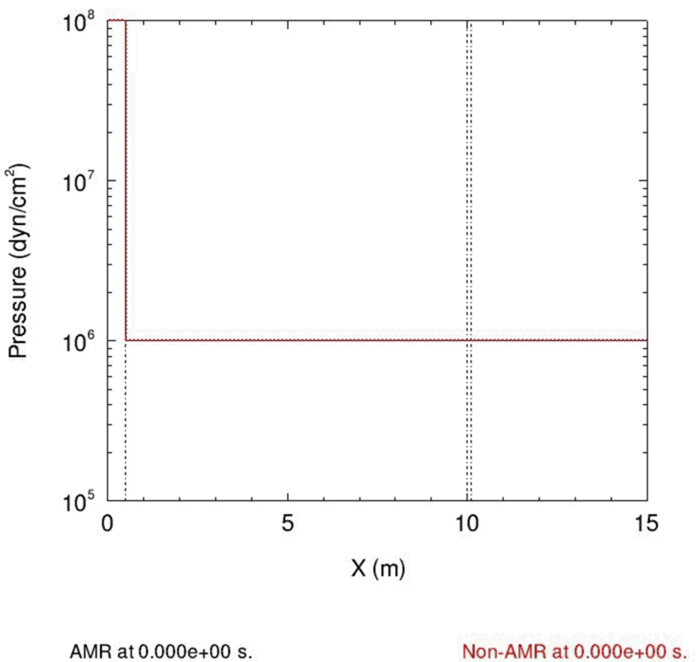

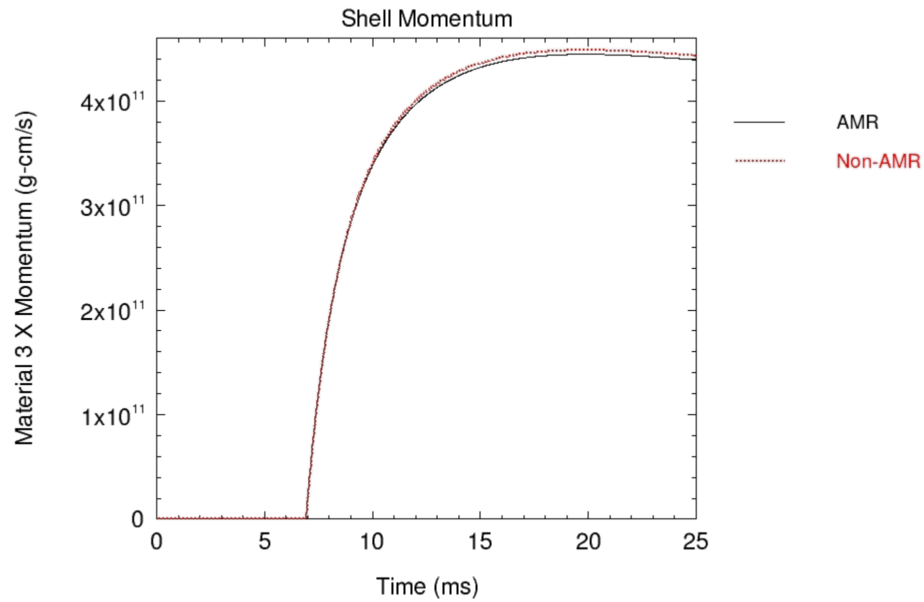

- 1 m sphere
- 10 cm shell
- 10 m standoff

- Non-AMR using uniform 1-cm grid
- AMR:
  - 0.5 cm grid on sphere and shell
  - 1-cm grid on background

Modified Sedov Blast Wave (cont.)

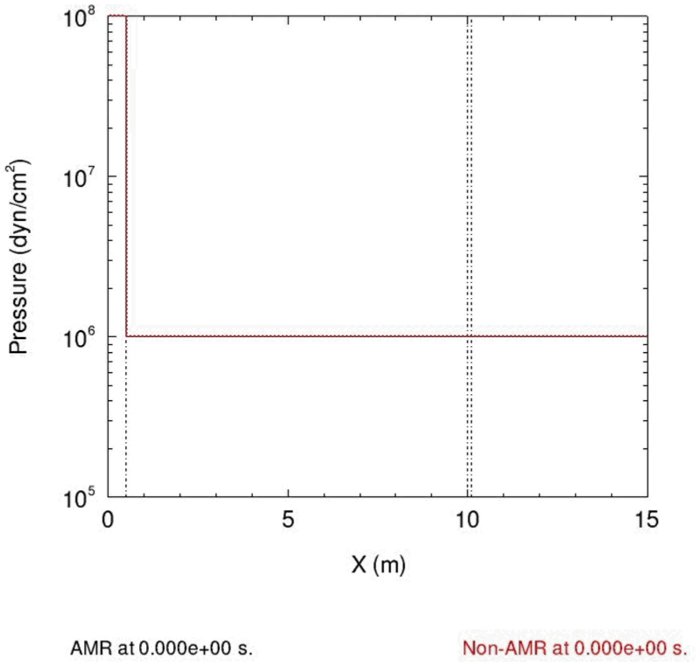

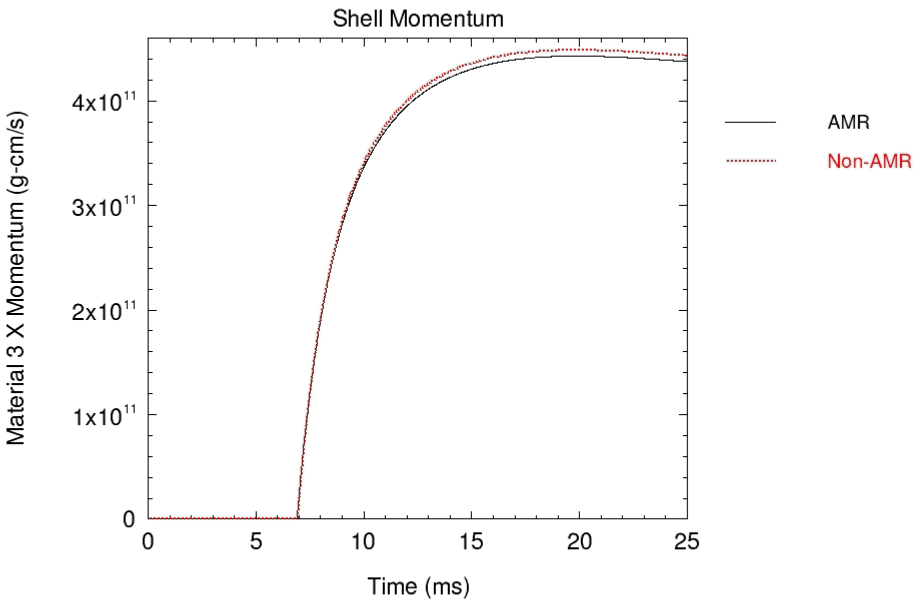

- 1 m sphere
- 10 cm shell
- 10 m standoff

- Non-AMR using uniform 1-cm grid
- AMR:
  - 0.25 cm grid on sphere and shell
  - 1-cm grid on background

Modified Sedov Blast Wave

conclusions

- Shock reflections occur when transitioning from low to high resolution mesh
  - Typical error of momentum *delivered to the plate* is less than 3%
  - Provided background mesh is well resolved to begin with
- Active area of research
  - Higher order refinement schemes
- Problem dependent
  - Importance of user validation

Conclusions

- AMR-CTH has had a successful decade
  - Increasingly used for production computing on large 3-D problems
- Order of magnitude performance advantage is routinely seen
  - In some cases 200-300x performance advantage is seen
- User-defined indicators are an important aspect of AMR-CTH
  - Our users helped us find some of these dramatic performance gains
  - Accuracy can be an issue if naïve indicators used
  - We’ve always known propagating shocks across resolution boundaries can cause reflections
  - Improving accuracy across resolution boundaries is an ongoing area of research