Computational Infrastructure

Erik Schnetter

MICRA 2009

København, August, 2009

CENTER FOR COMPUTATION
& TECHNOLOGY

LSU CENTER FOR COMPUTATION & TECHNOLOGY

Provocative Statement:

Our research is not limited by our physics models; it is limited by our computational tools.

CENTER FOR COMPUTATION & TECHNOLOGY

miCra

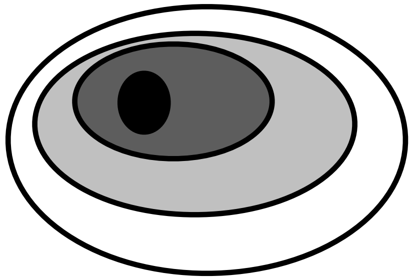

- Example problem:
  Find optical depth in given object

Method 1: determine level sets	good approximation	inefficient in parallel
Method 2: ray by ray	fails if far from spherical symmetry	naturally parallel, easy to implement

... which method will people choose?

Codes

- Learning new physics: takes weeks; implementing new physics: takes months
- We give names to codes (Vulcan, Whisky), but not to papers (Berger&Oliger 1984)
- People leave the field, codes stay around
- <u>Consequence:</u>
  Groups compete not only via ideas and physics, but also via codes

Cactus Software Framework

- Basic idea: Split a code into *components* which can be maintained and distributed separately
- e.g. BSSN, GRMHD, horizons, AMR, time stepping, parallelism, I/O
- Goal: simplify collaboration between different groups and different fields
- See http://www.cactuscode.org/

Cactus in Relativistic Astrophysics

- Three layers of abstraction in a typical code:
- Top: specific physics codes, typically developed by single research groups
- Middle: toolkit, e.g. for numerical relativity, developed by community
- Bottom: computational infrastructure, developed by computer scientists

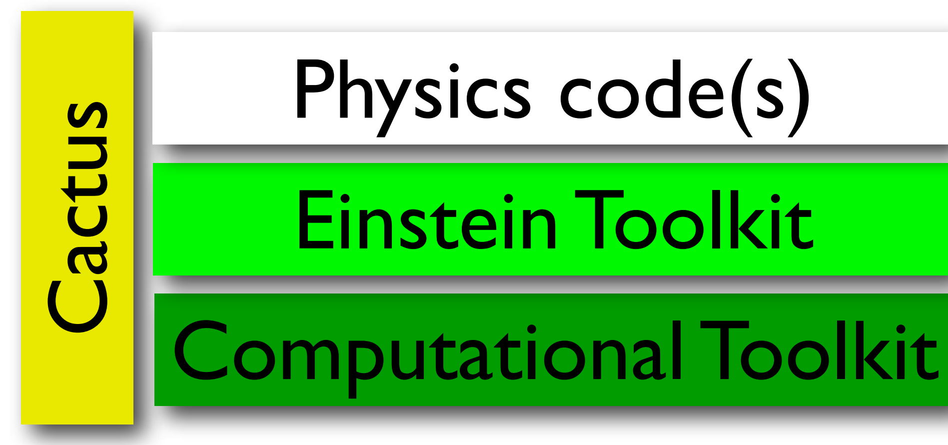

Thorn Structure

Inside view of a plug-in module, or thorn for Cactus

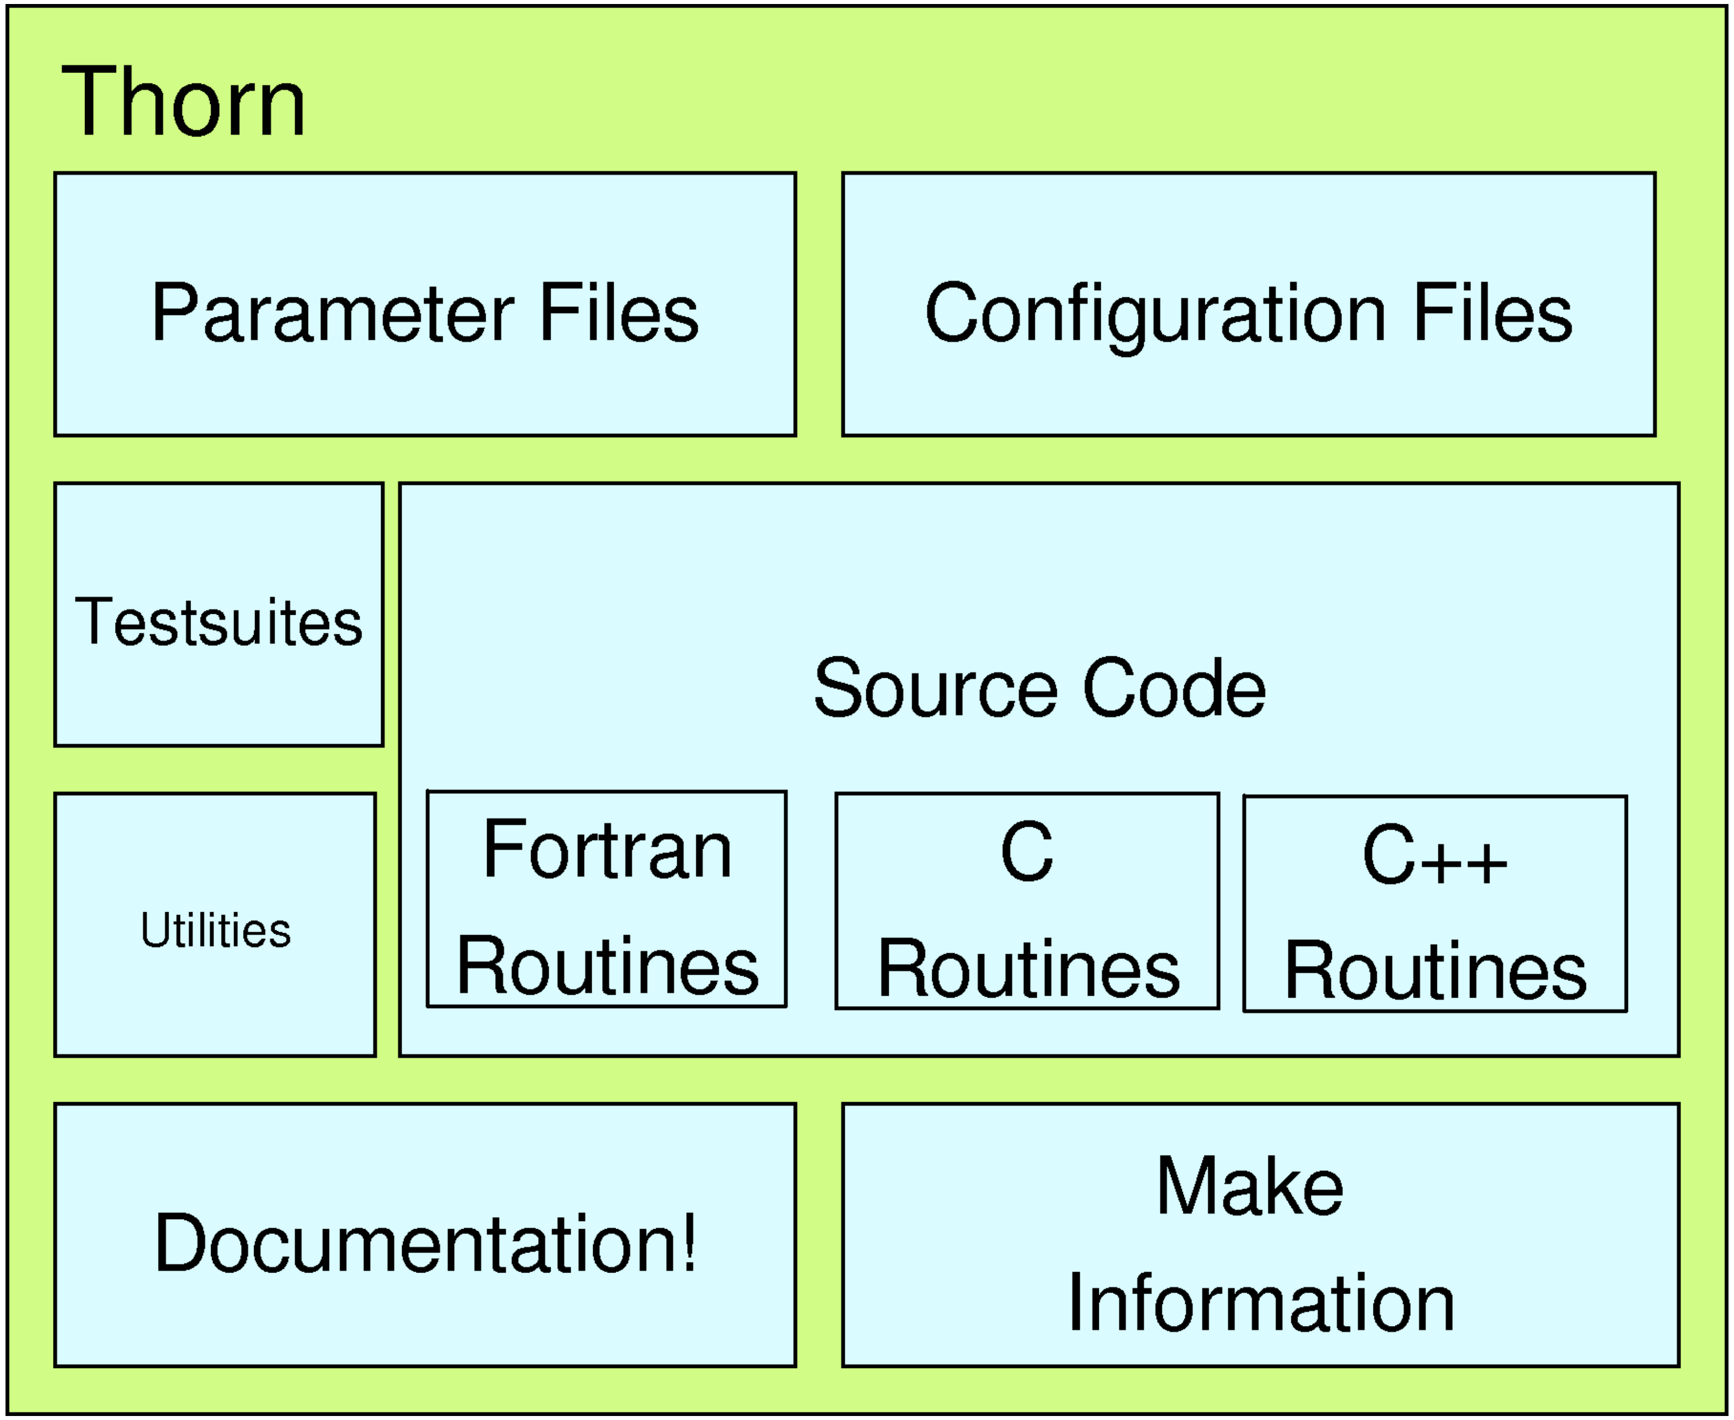

Carpet: Scalable Adaptive Mesh Refinement

LSU
CENTER FOR COMPUTATION
& TECHNOLOGY

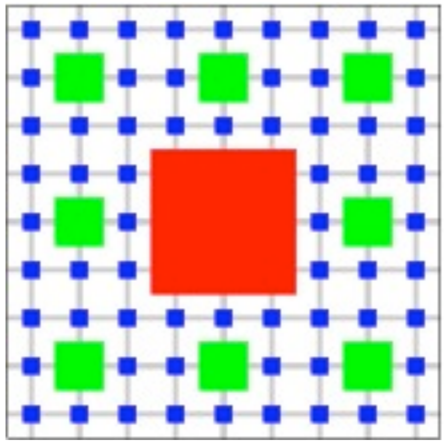

- Berger-Oliger adaptive mesh refinement (AMR) with subcycling in time
- Higher order methods require up to 5 ghost zones (may lead to a memory overhead of more than a factor of 2)
- Hybrid parallelisation
- AMR tracks physics features, refining around black holes or neutron stars
- See http://www.carpetcode.org/

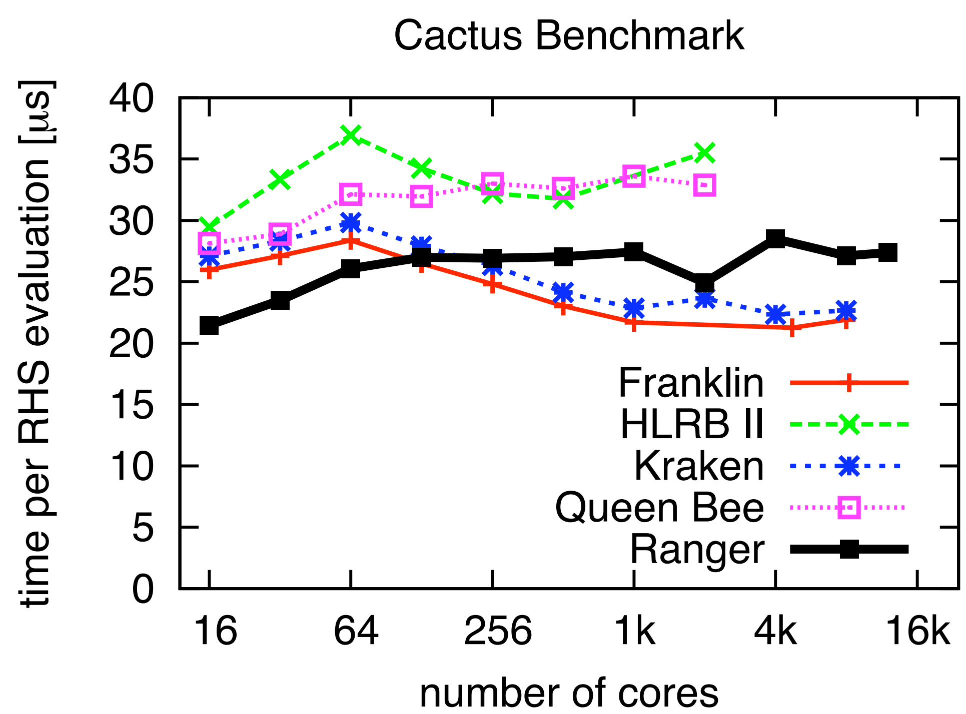

Weak scaling benchmark,
9 levels of mesh refinement,
very good parallel scaling

[L. Rezzolla, R. Kähler]

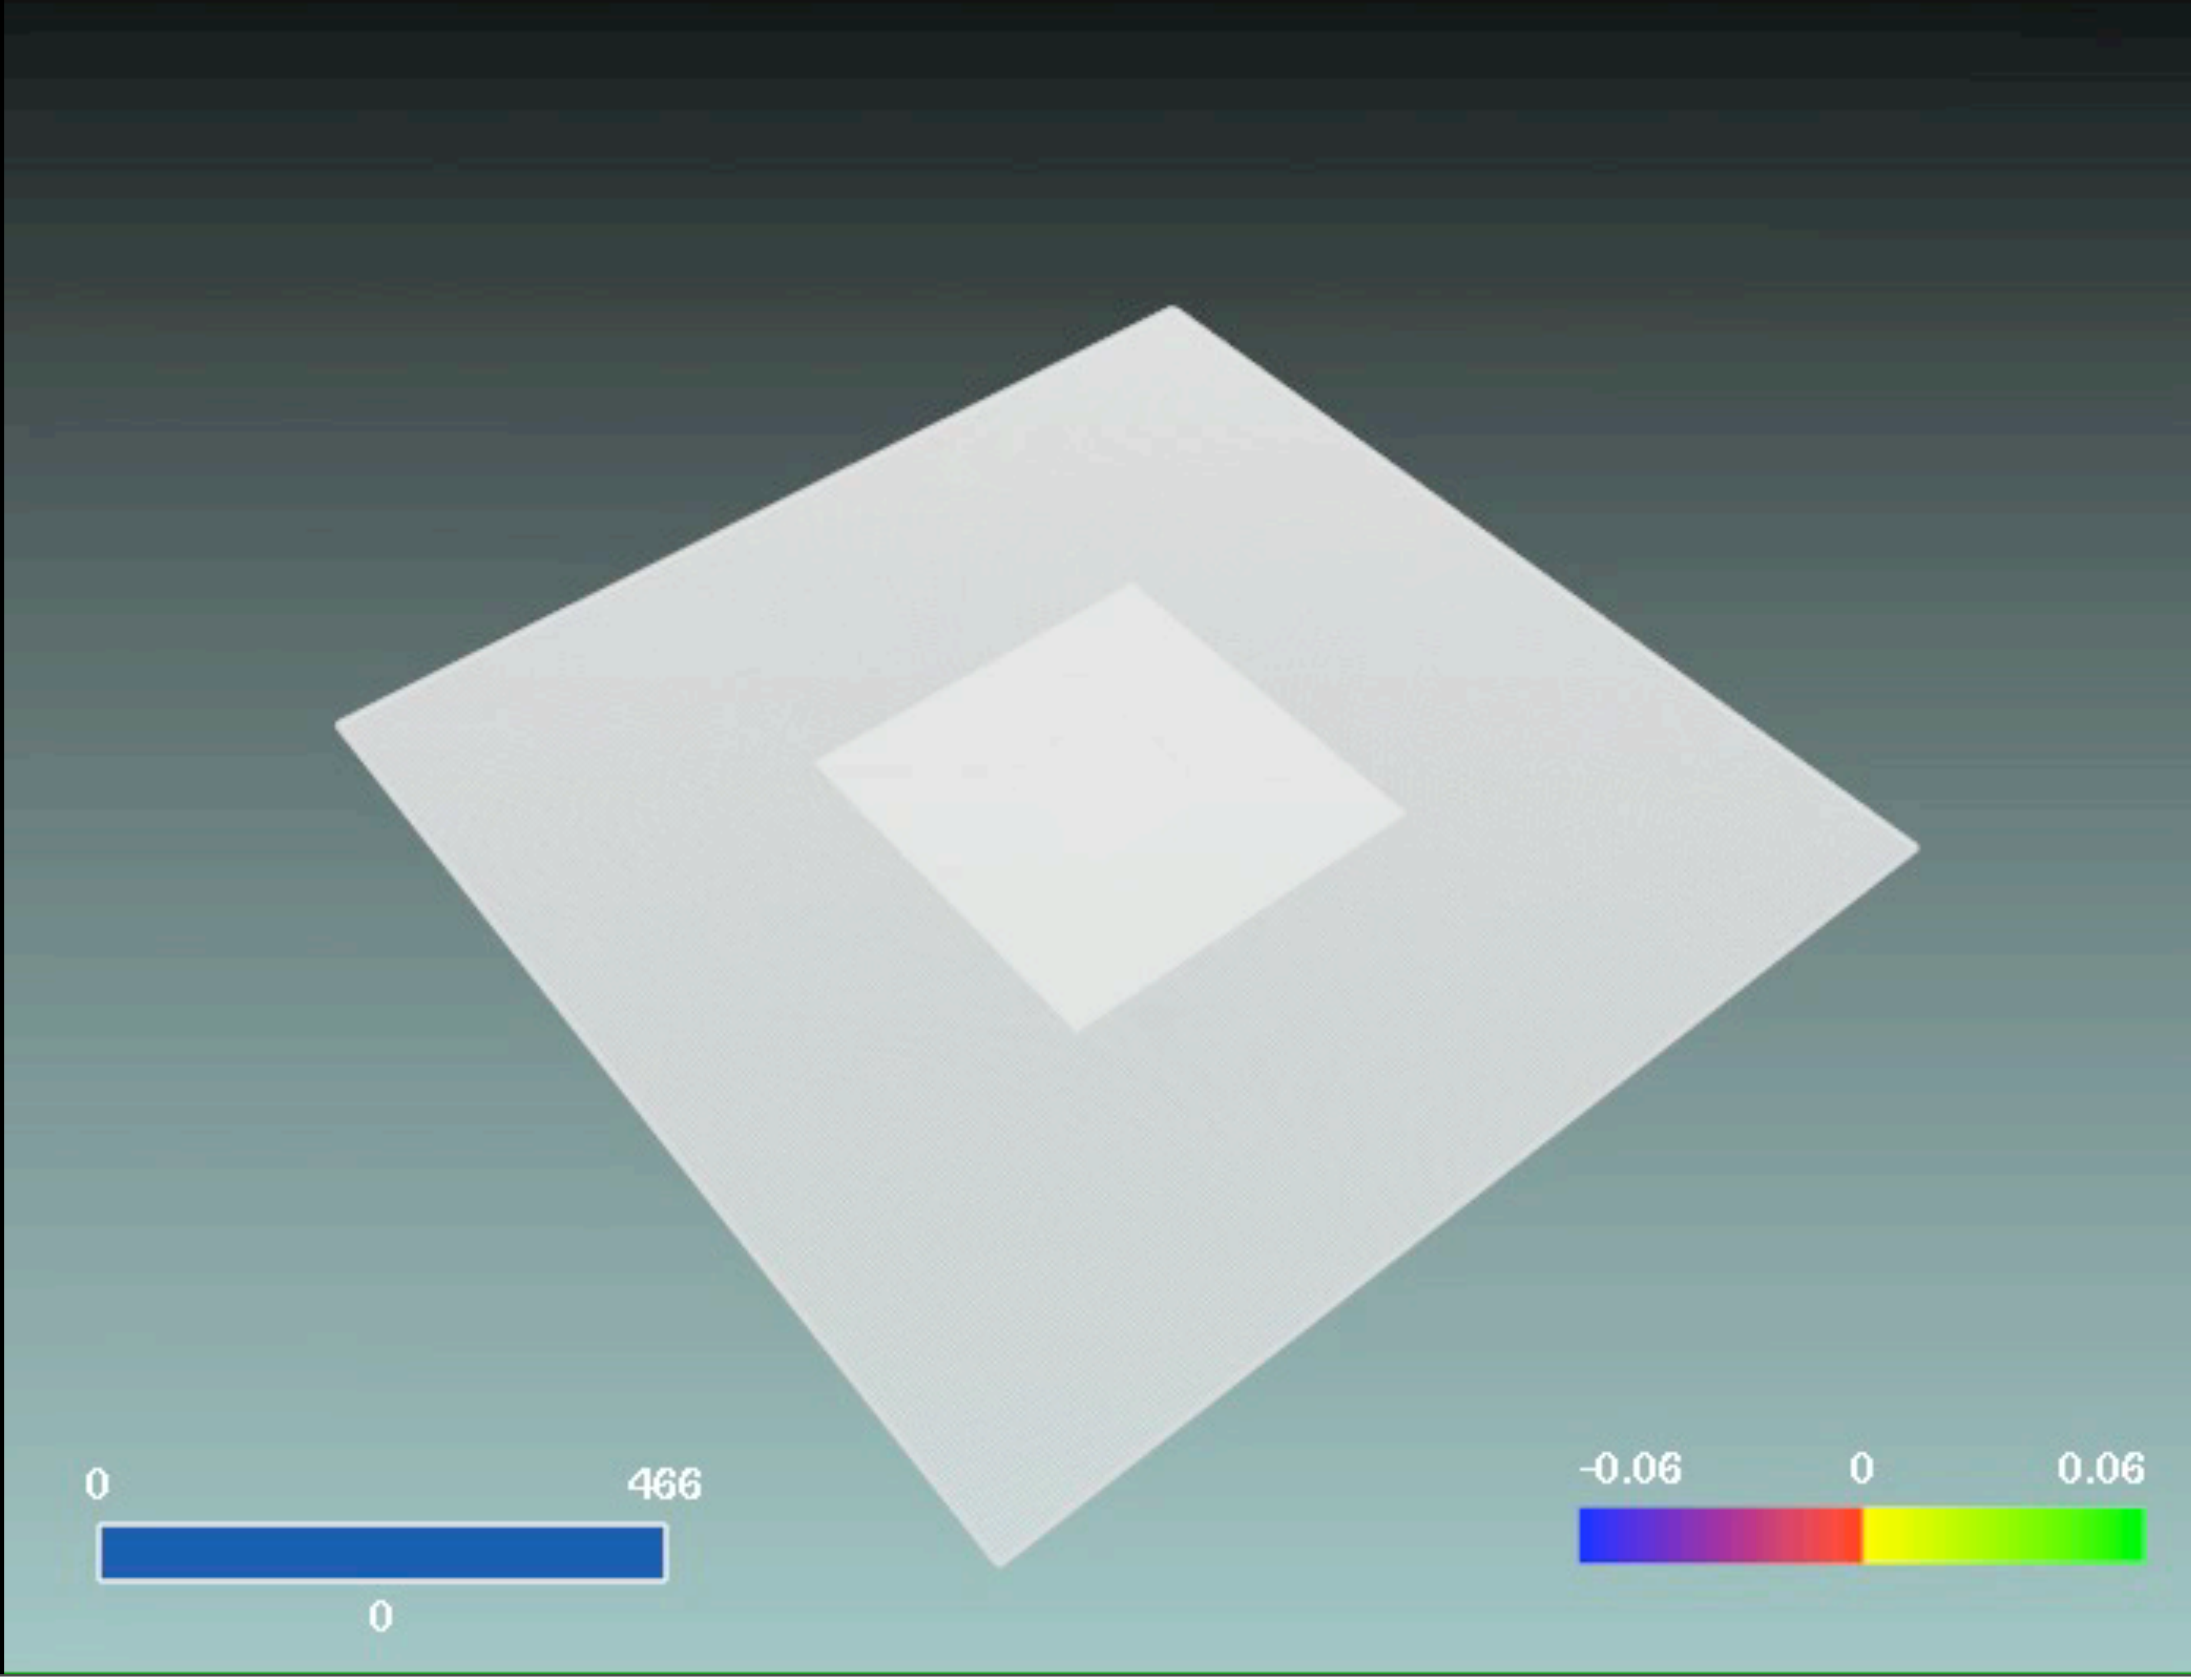

Numbers

- Large scale differences and moving objects require adaptive mesh refinement (AMR)
[typical: L=1000, h=0.02, using 9 refinement levels]
- Long time evolutions and desired accuracy require high order methods (4th order or higher)
- Multi-block methods: Much more efficient far away from source ["spherical" grids: O(L) vs. $O(L^3)$] can have causally disconnected outer boundaries

LSU
CENTER FOR COMPUTATION
& TECHNOLOGY

Multi-Patch Systems

BH:

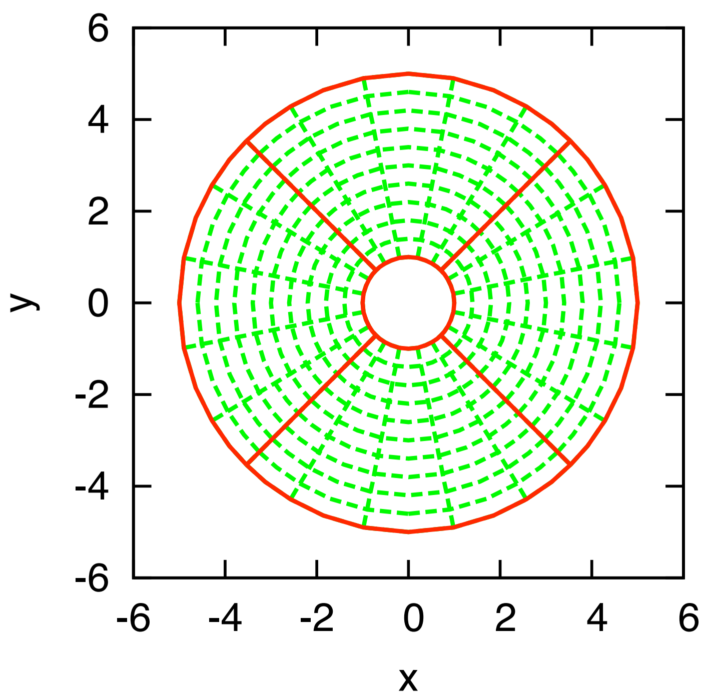

NS:

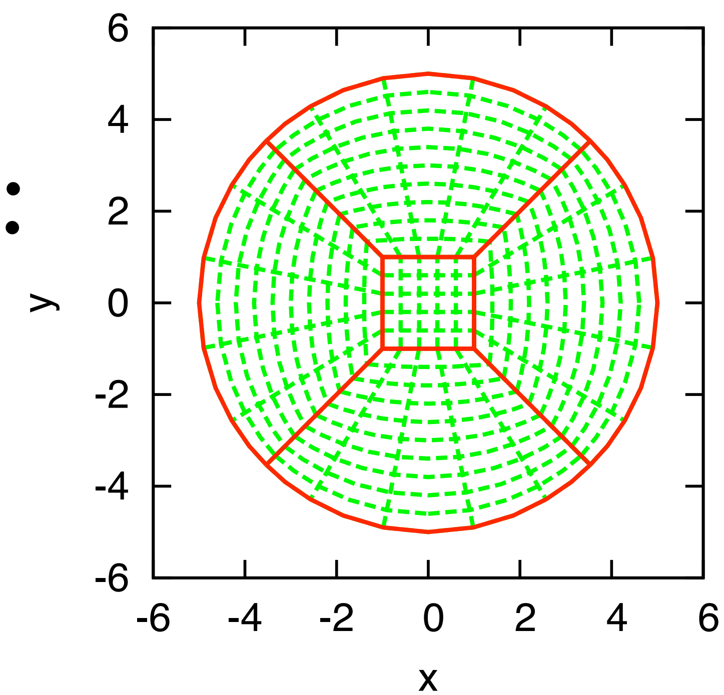

Binary:

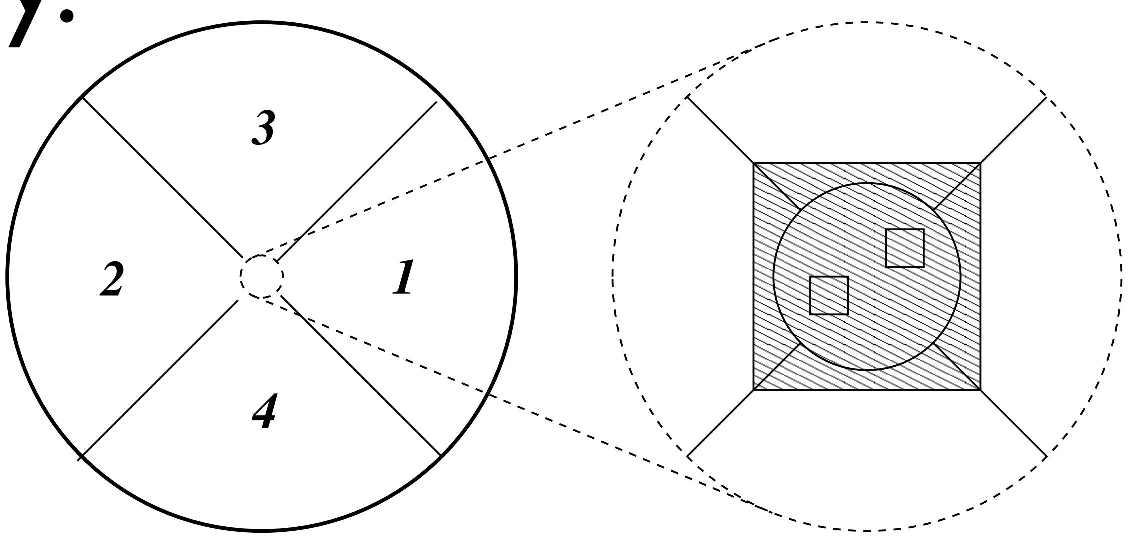

BH with wave:

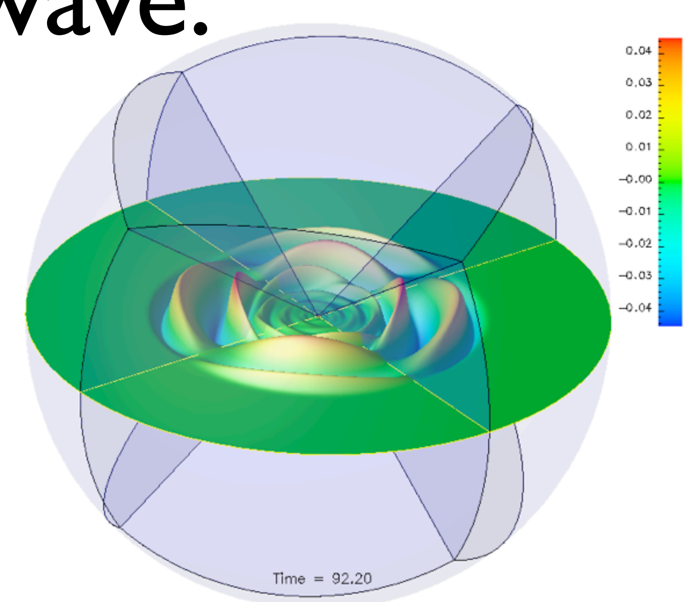

Einstein Toolkit: Free, Public Components

- Spacetime evolution: McLachlan
- GR hydro: Whisky
- horizons
- exact solutions (testing)
- AMR: Carpet
- Note: These are made available by different groups, not just LSU
- Additional components may be available for those who ask

LSU
CENTER FOR COMPUTATION & TECHNOLOGY

Summary

- Codes are only tools, it is strange that they are so important in the daily routine
- The Cactus framework makes possible collaboration between competitors
- McLachlan, Whisky, Carpet, Cactus (and friends) form a public, basic code for GR hydro simulations: *Einstein Toolkit*