Octopus
A scalable AMR tookit for astrophysics

Bryce Adelstein-Lelbach[1], Zach Byerly[3], Dominic Marcello[3]

PI: Hartmut Kaiser[1][2], Co-PI: Geoffrey Clayton[3]

[1]: Center for Computation and Technology
[2]: LSU Department of Computer Science
[3]: LSU Department of Physics

STE||AR

stellar.cct.lsu.edu

Overview of Research

- NSF STAR project: a cross-discipline collaboration between LSU computer scientists and astrophysicists.
  - Primary goal is to facilitate a highly realistic simulation of the merger of two white dwarfs.
    - The study of these binaries is important as they are possible progenitors of a number of astrophysically important objects, such as Type 1a supernovae.

Overview of Research

- Development of new adaptive mesh refinement (AMR) codes utilizing HPX, a framework for message-driven computation, instead of traditional HPC programming mechanisms (MPI, OpenMP, PGAS).
  - Existing unigrid codes are too slow.
    - 0.2 orbits/day running on 1,032 cores.
    - We want to be running hundreds if not thousands of orbits.
  - AMR codes can be many orders of magnitude faster (10^4 – 10^6).
    - Doing AMR with MPI is difficult and can face scalability problems, due to the inherently inhomogeneous nature of AMR.

Numerical Methods

- Our group uses 3D Eulerian hydrodynamics codes:
  - Explicit advection scheme
    - Kurganov and Tadmor, 2000, Journal of Computational Physics, 160, 241
  - Finite-volume method
  - Adaptive mesh refinement
    - Multigrid method (for solving Poisson’s equations): Martin and Cartwright, 1996
    - Interpolation (PPM): Colella and Woodward, 1984, Journal of Computational Physics
  - Angular momentum conservation
- Other references:
  - Dominic Marcello’s Ph.D. thesis

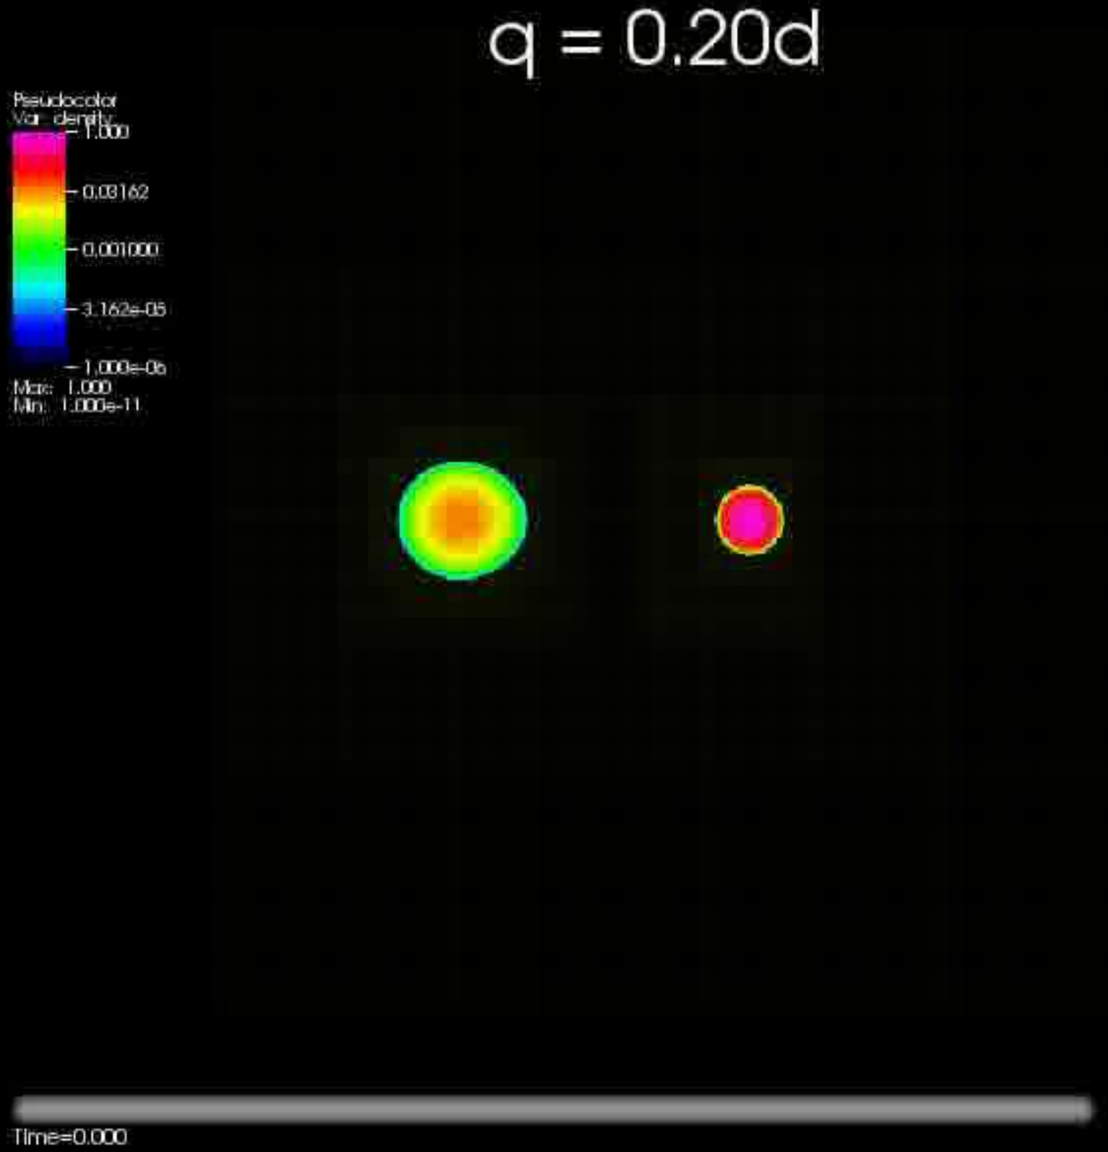

Source: Dominic Marcello, LSU Department of Physics

What’s HPX?

- A general purpose C++ runtime system for parallel and distributed applications of any scale.
- The HPX paradigm prefers:
  - **Asynchronous communication** to hide latencies and contention instead of avoiding them.
  - **Fine-grained parallelism** and an **active global address space** to enable dynamic and heuristic load balancing instead of statically partitioning work.
  - **Local, dependency-driven synchronization** instead of explicit global barriers.
  - Sending **work to data** instead of sending data to work.

Hiding Latency and Contention
(pull model)

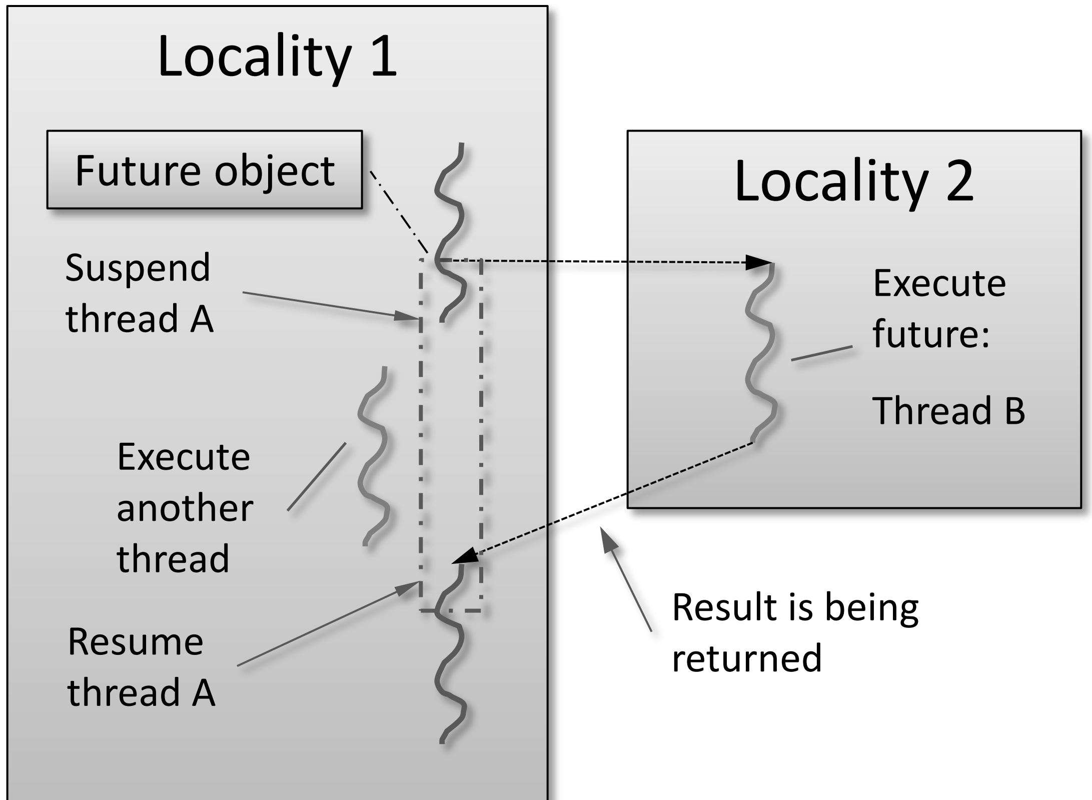

Hiding Latency and Contention
(push model)

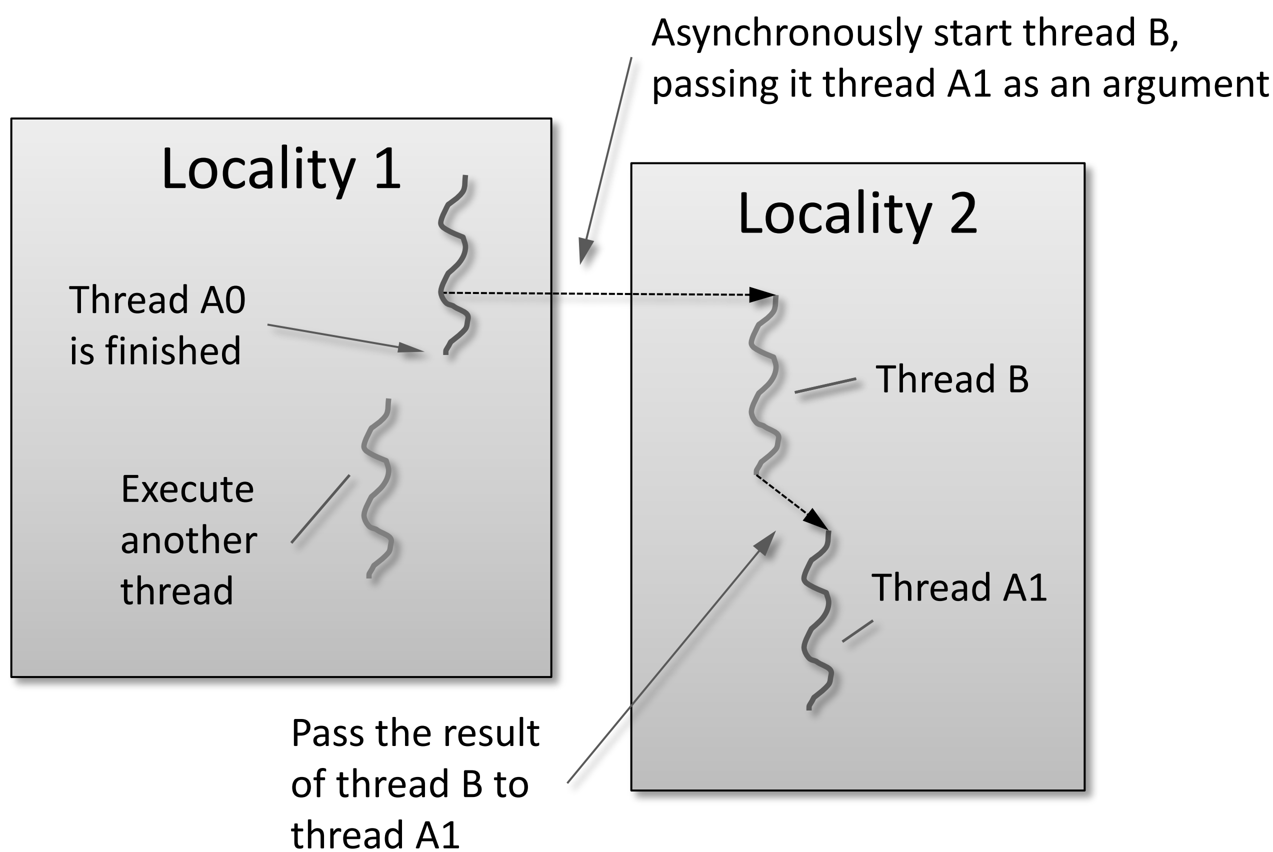

Asynchronous vs Synchronous

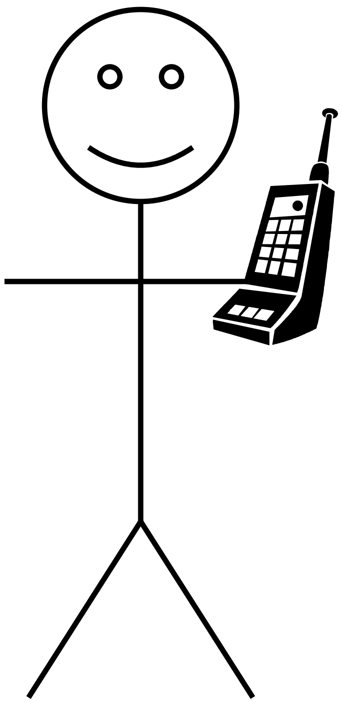

A phone call is a form of synchronous communication.

Texting is a form of asynchronous communication. A text message is a future.

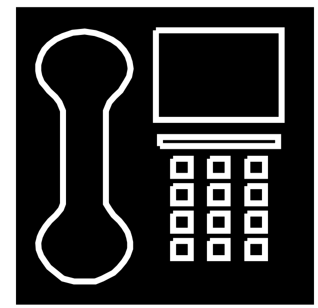

What's Octopus?

- A general purpose HPX AMR framework.
  - Based heavily on ideas drawn from an existing LSU OpenMP AMR code.
    - Octree-based AMR.
  - Primarily designed for high-resolution, high-accuracy astrophysics hydrodynamics simulations.
- Octopus design:
  - **Multi-tiered software architecture** to maintain abstraction while supporting domain-specific physics.
  - **Policy-driven** genericity.
  - Powerful **optimizations** applied to the generic layers:
    - Timestep size prediction.
    - Time-based refinement.
    - Localization of dependencies
    - Eager computation of fluxes.

Octopus Architecture

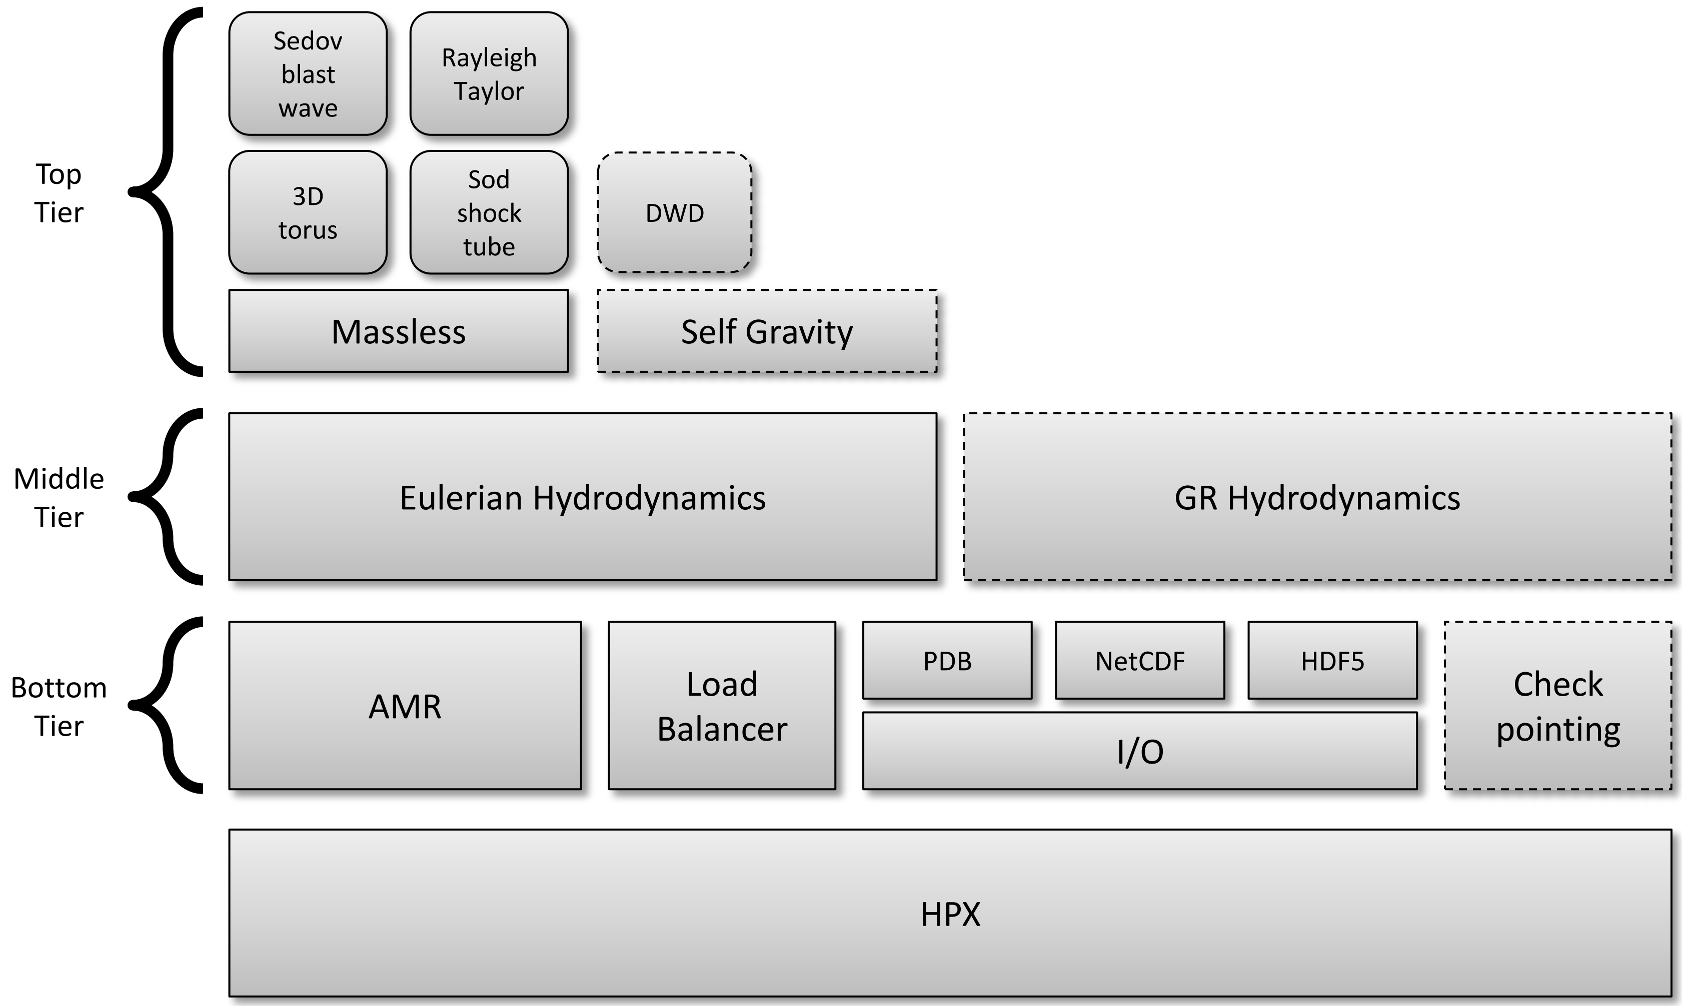

Universal Scalability Law

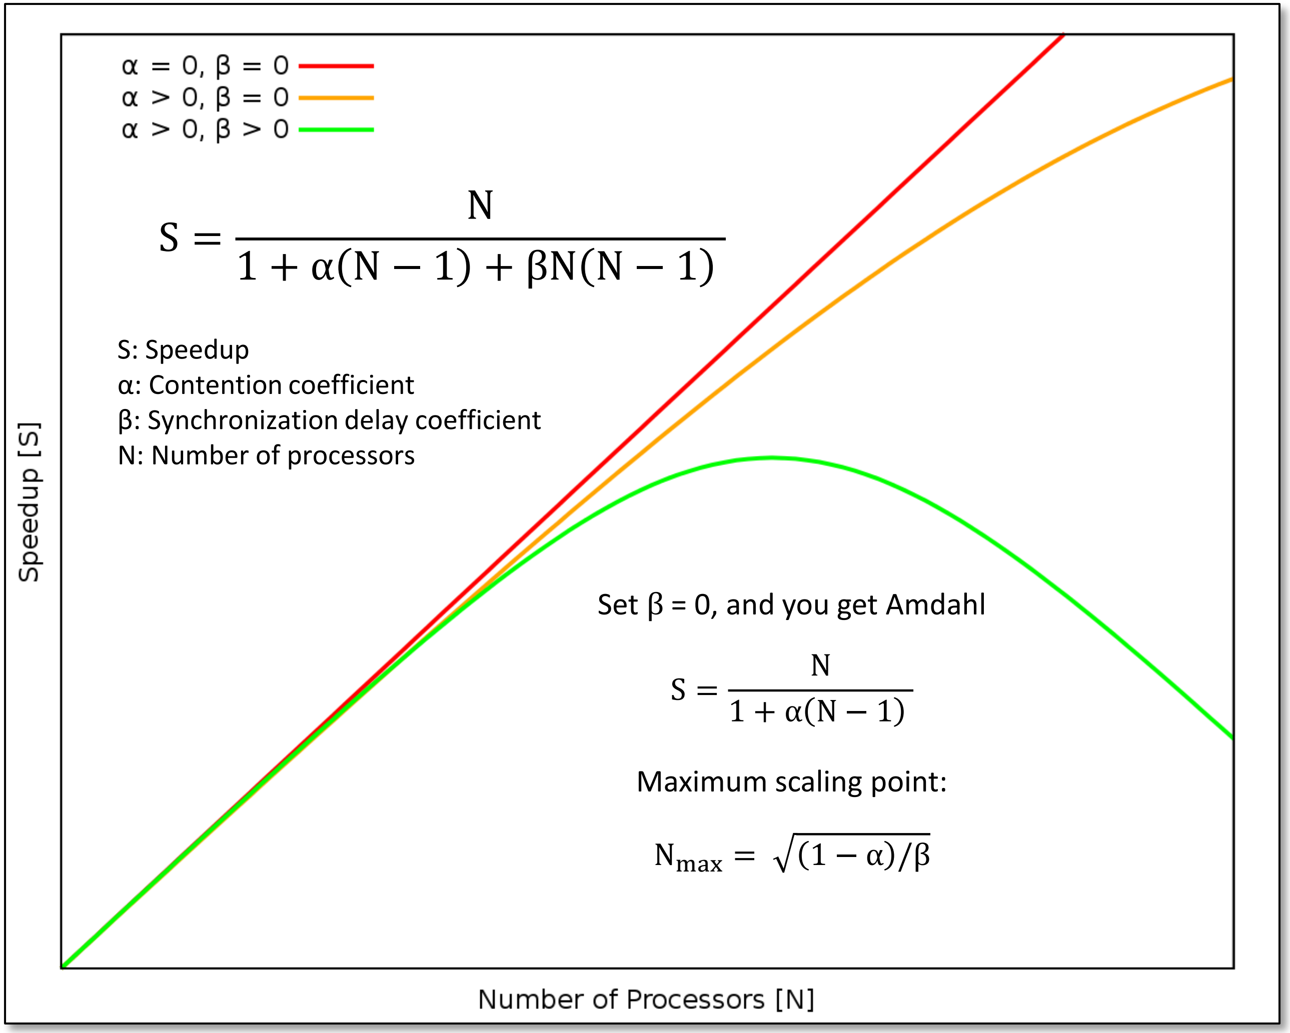

Scaling of OpenMP 3D Eulerian Code vs Octopus 3D Eulerian Code

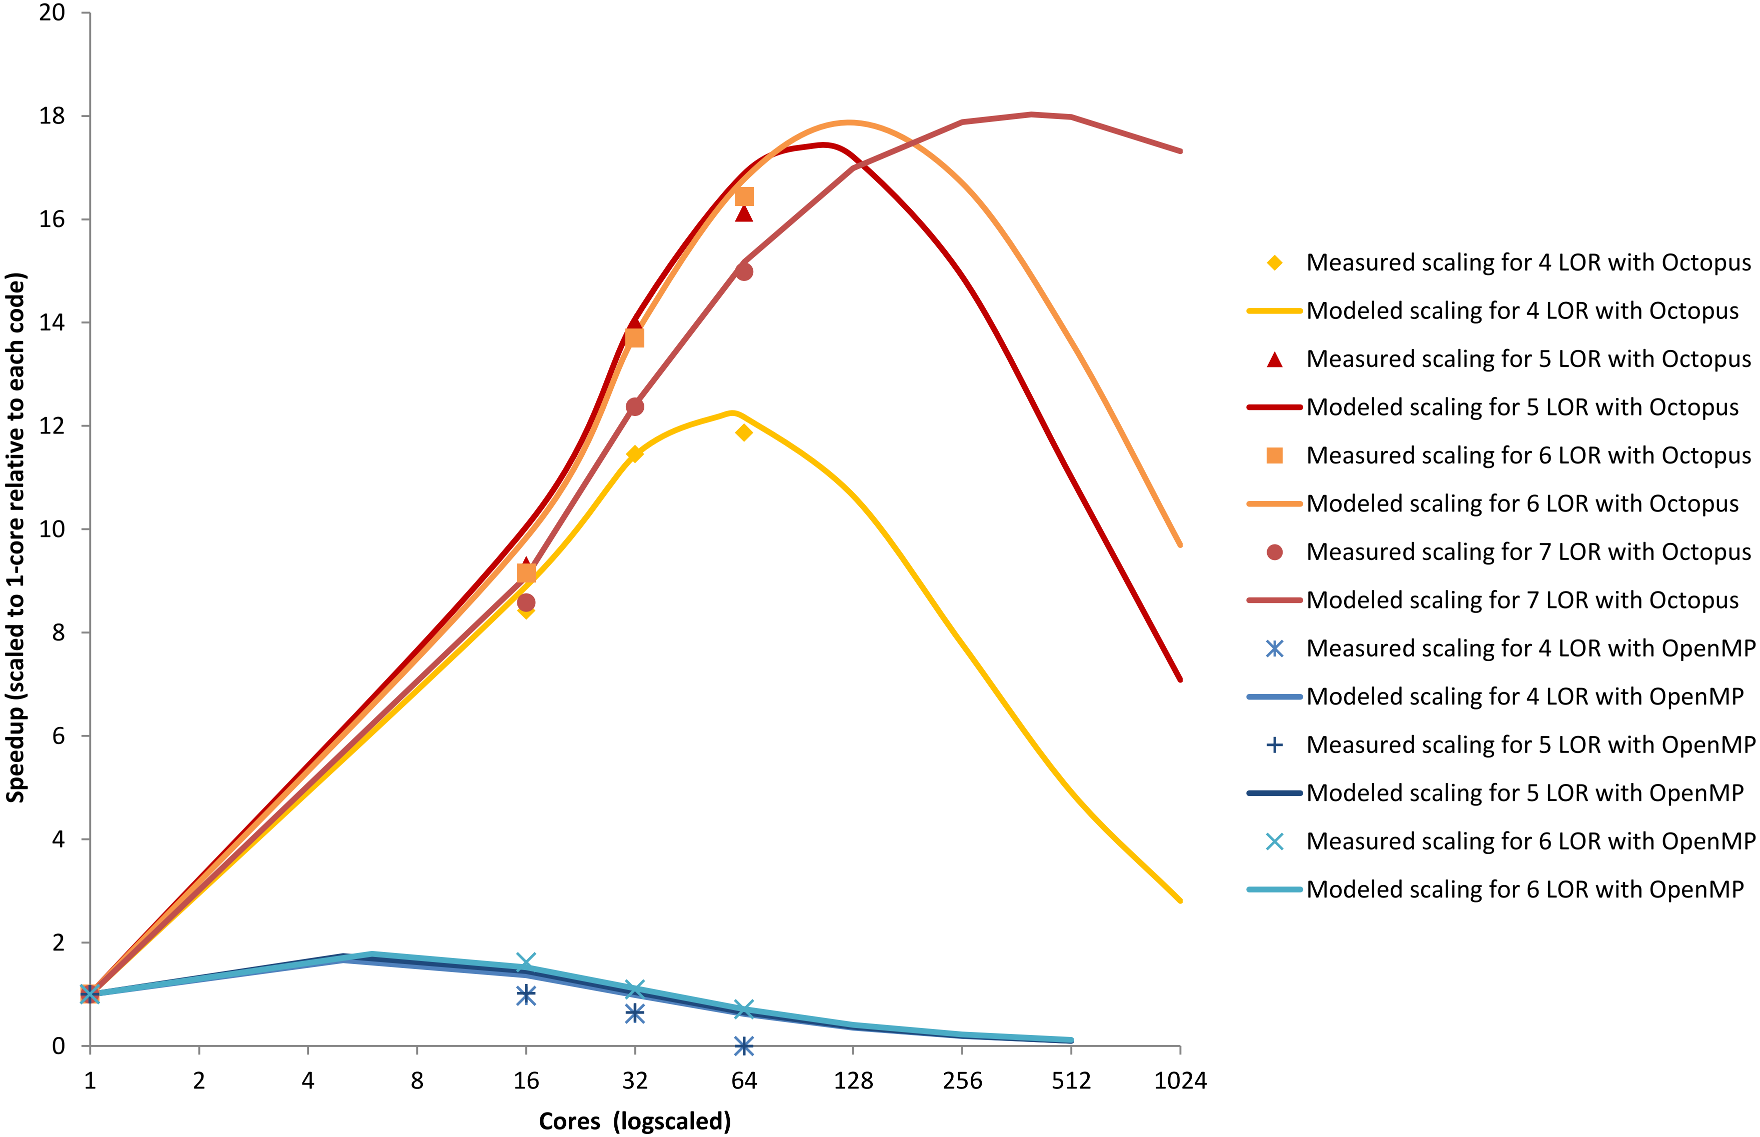

USL Modeling

Octopus			
LOR	α	β	N_{max}
4	0.04840	3.00E-04	56
5	0.03790	1.00E-04	98
6	0.04084	6.00E-05	126
7	0.05069	6.00E-06	397

OpenMP			
LOR	α	β	N_{max}
4	0.40890	1.89E-02	5
5	0.37560	1.83E-02	5
6	0.37470	1.62E-02	6

Grain Size

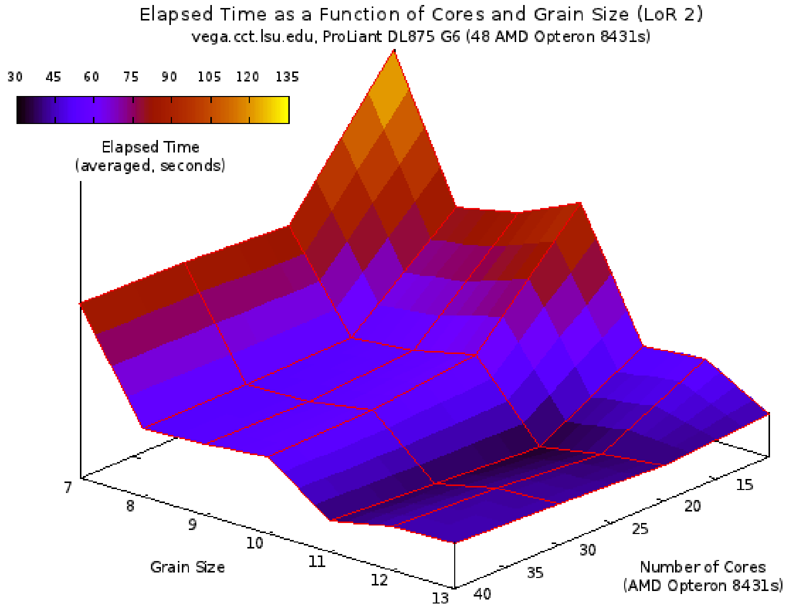
Elapsed Time as a Function of Cores and Grain Size (LoR 2)
vega.cct.lsu.edu, ProLiant DL875 G6 (48 AMD Opteron 8431s)
30 45 60 75 90 105 120 135
Elapsed Time
(averaged, seconds)
7
8
9
10
11
12
13
Grain Size
15
20
25
30
35
40
Number of Cores
(AMD Opteron 8431s)

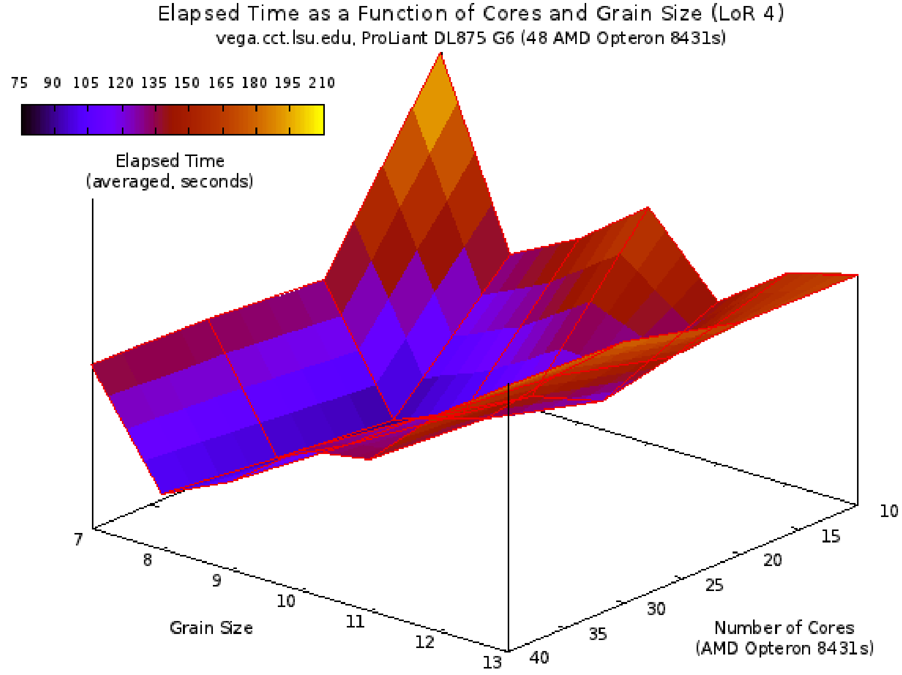
Elapsed Time as a Function of Cores and Grain Size (LoR 4)
vega.cct.lsu.edu, ProLiant DL875 G6 (48 AMD Opteron 8431s)
75 90 105 120 135 150 165 180 195 210
Elapsed Time
(averaged, seconds)
7
8
9
10
11
12
13
Grain Size
10
15
20
25
30
35
40
Number of Cores
(AMD Opteron 8431s)

STE||AR

stellar.cct.lsu.edu