High Performance, Massively Parallel Computing with FLASH

Dongwook Lee
Applied Mathematics & Statistics
University of California, Santa Cruz

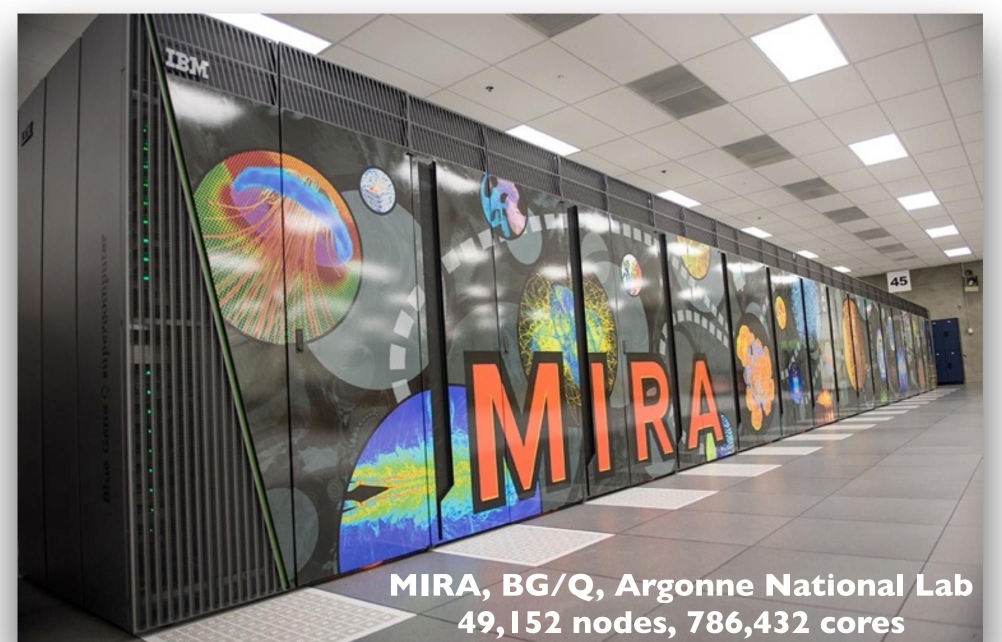

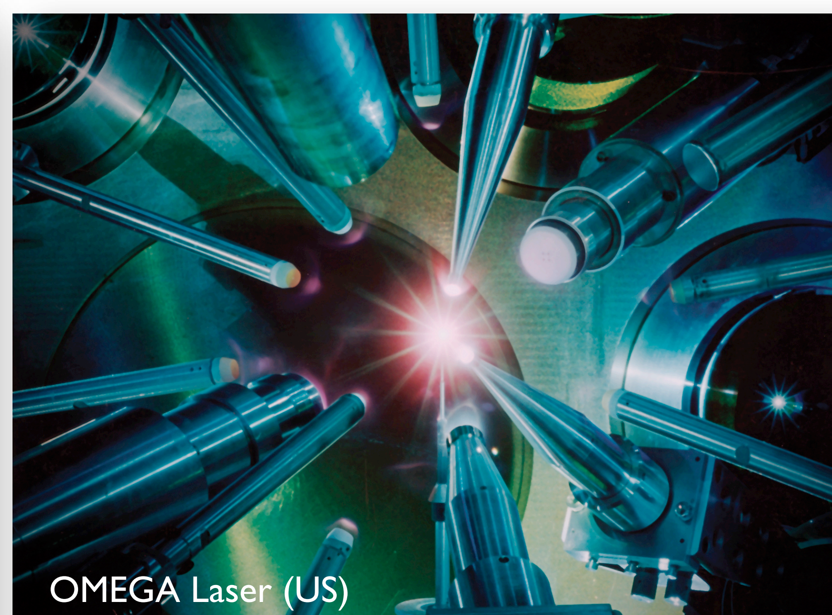

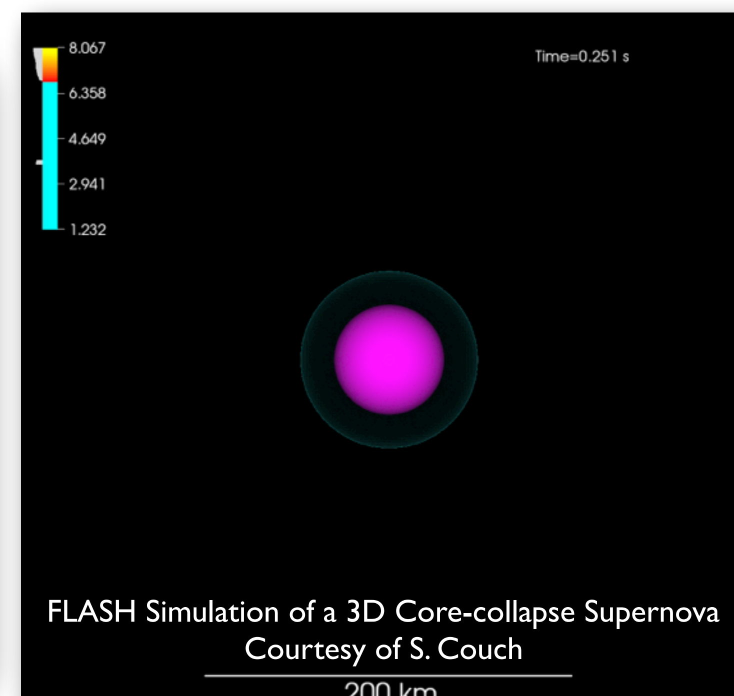

AMS 209, Fall, 2015

Parallel Computing in1917

Baskin Engineering UC SANTA CRUZ

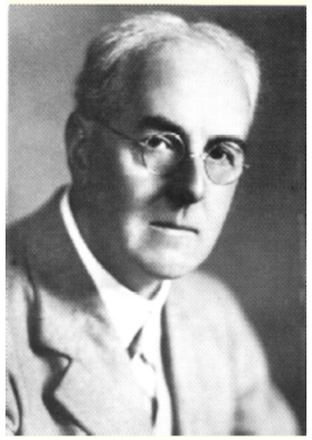

Richardson (1881-1953) introduced a principle of parallel computation using men power in order to conduct weather predictions

=

"High-Performance" Scientific Computation

HPC in 21st Century

- Goal: To solve large problems in science, engineering, or business using multi physics, multi scale simulations on large scale computing platforms
- Hardware: Currently petaflops/s (going to exaflops/s), heterogeneous architectures, inter-node and intra-node parallelism
- Software: MPI, OpenMP (fine vs. coarse grained), discretizers, partitioners, solvers and integrators, grid mesh, grid adaptivity, load balancing, UQ, data compression, config system, compilers, debuggers, profilers, translators, performance optimizers, I/O, workflow controllers, visualization systems, etc…
- Hardware & software: concurrent programming models, co-Design, fault monitoring, etc…

Flop/s Rate Increase since 1988

- ACM Gordon Bell Prize
  - 1 Gigaflop/s — structural simulation, 1988
  - 1 Teraflop/s — compressible turbulence, 1998
  - 1 Petaflop/s — molecular dynamics simulation, 2008

- Total of 6 orders of magnitude of improvements for individual cores over the two decades:
  - computer engineering (only two orders of magnitude)
  - concurrency (the rest four orders of magnitude)

High Performance Computing (HPC)

- This tension between computation & memory brings a paradigm shift in numerical algorithms for HPC
- To enable scientific computing on HPC architectures:
  - efficient parallel computing, (e.g., data parallelism, task parallelism, MPI, multi-threading, GPU accelerator, etc.)
  - better numerical algorithms for HPC

Astrophysics Application

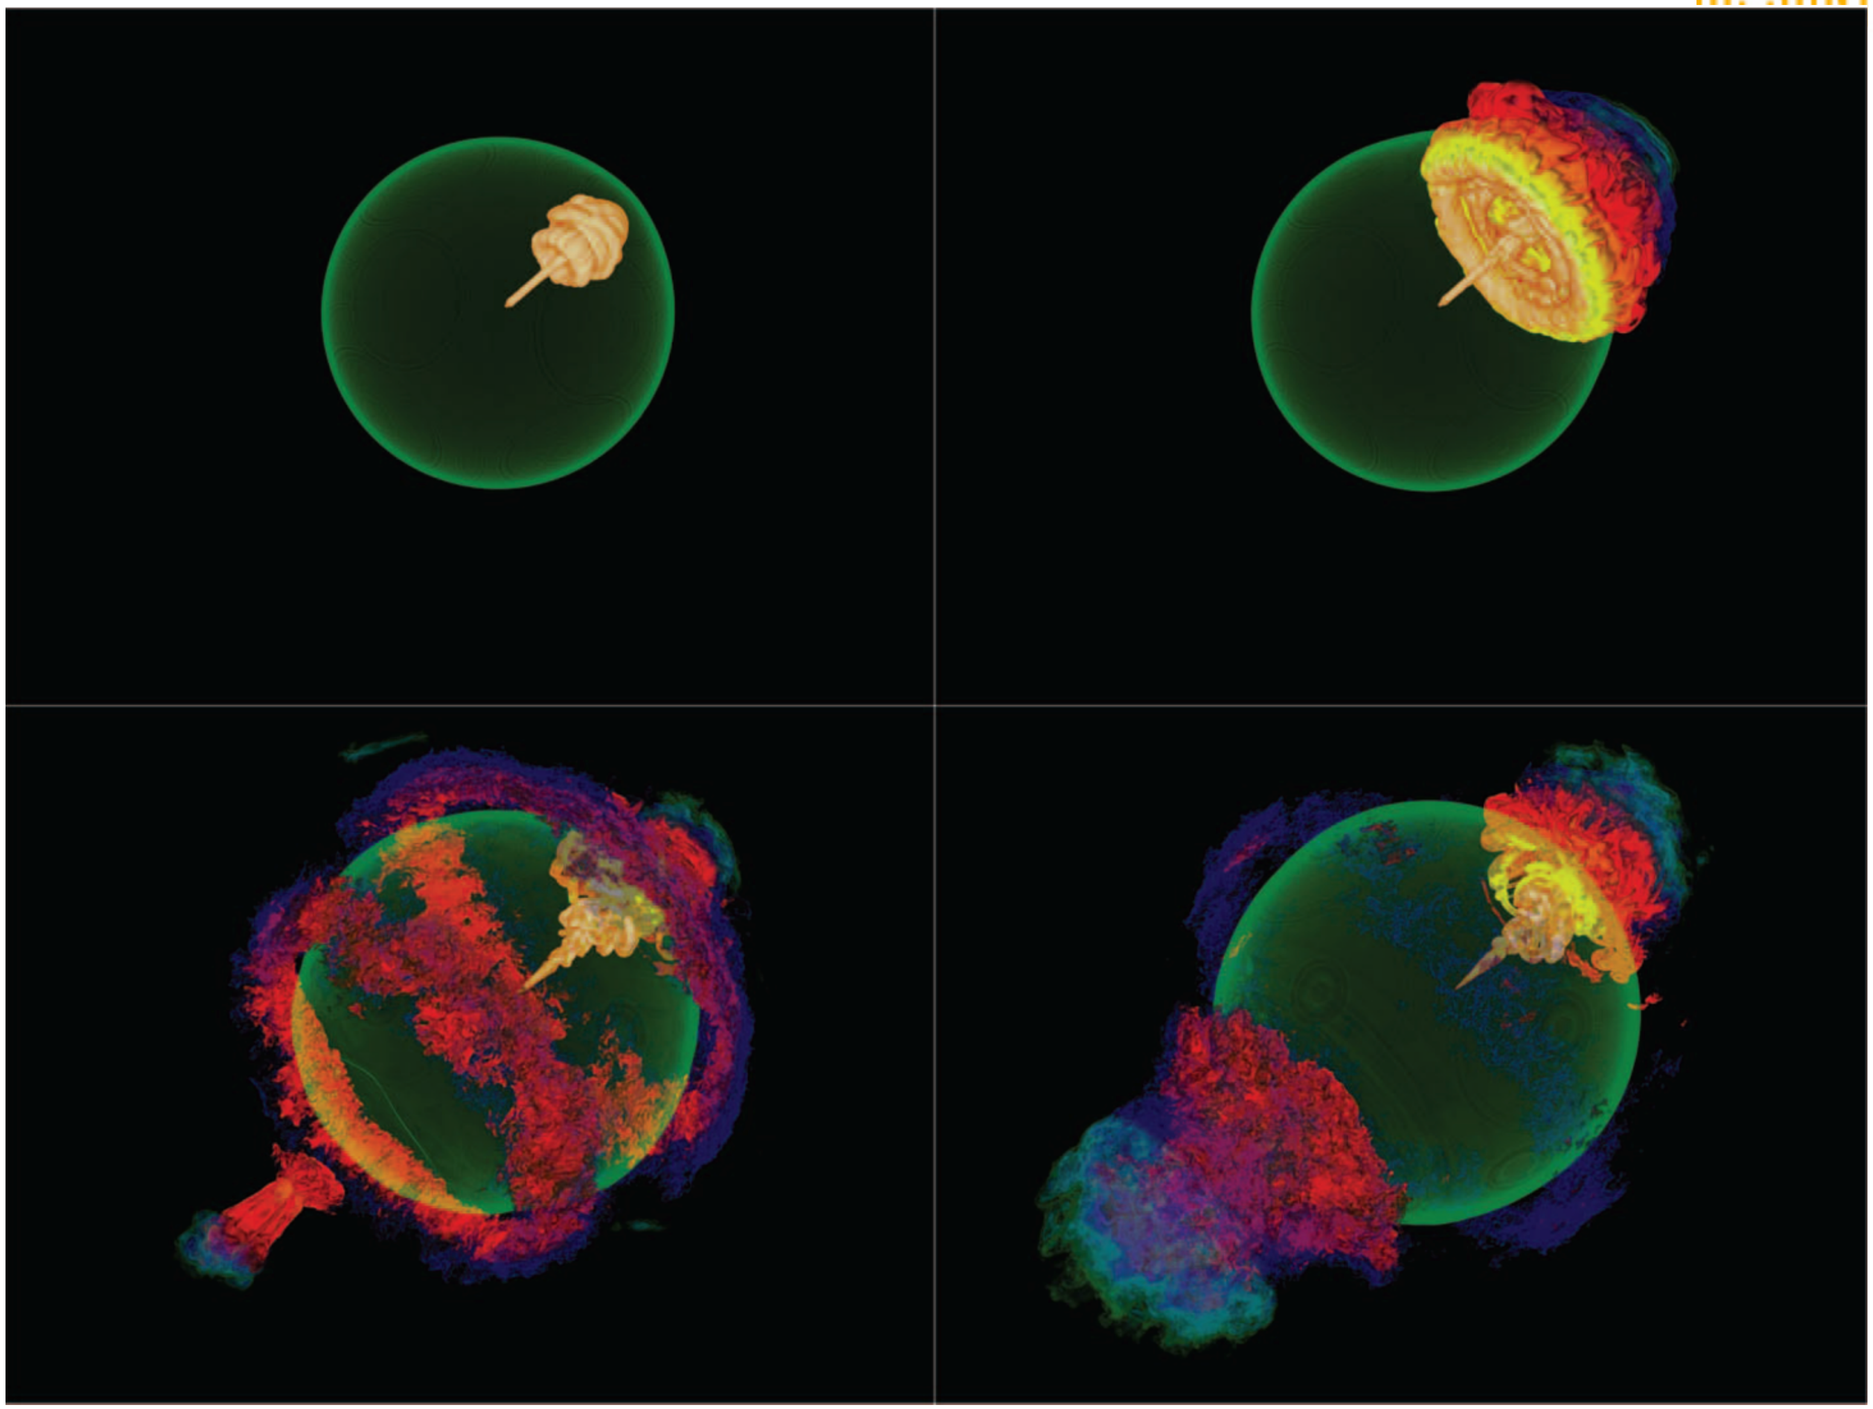

Gravitationally Confined Detonation of Ia SN:
Calder et al (2003); Calder & Lamb (2004); Townsley et al (2007); Plewa (2007); Plewa and Kasen (2007); Jordan et al (2008); Meakin et al (2009); Jordan et al (2012)

https://vimeo.com/40696524

Astrophysics Application

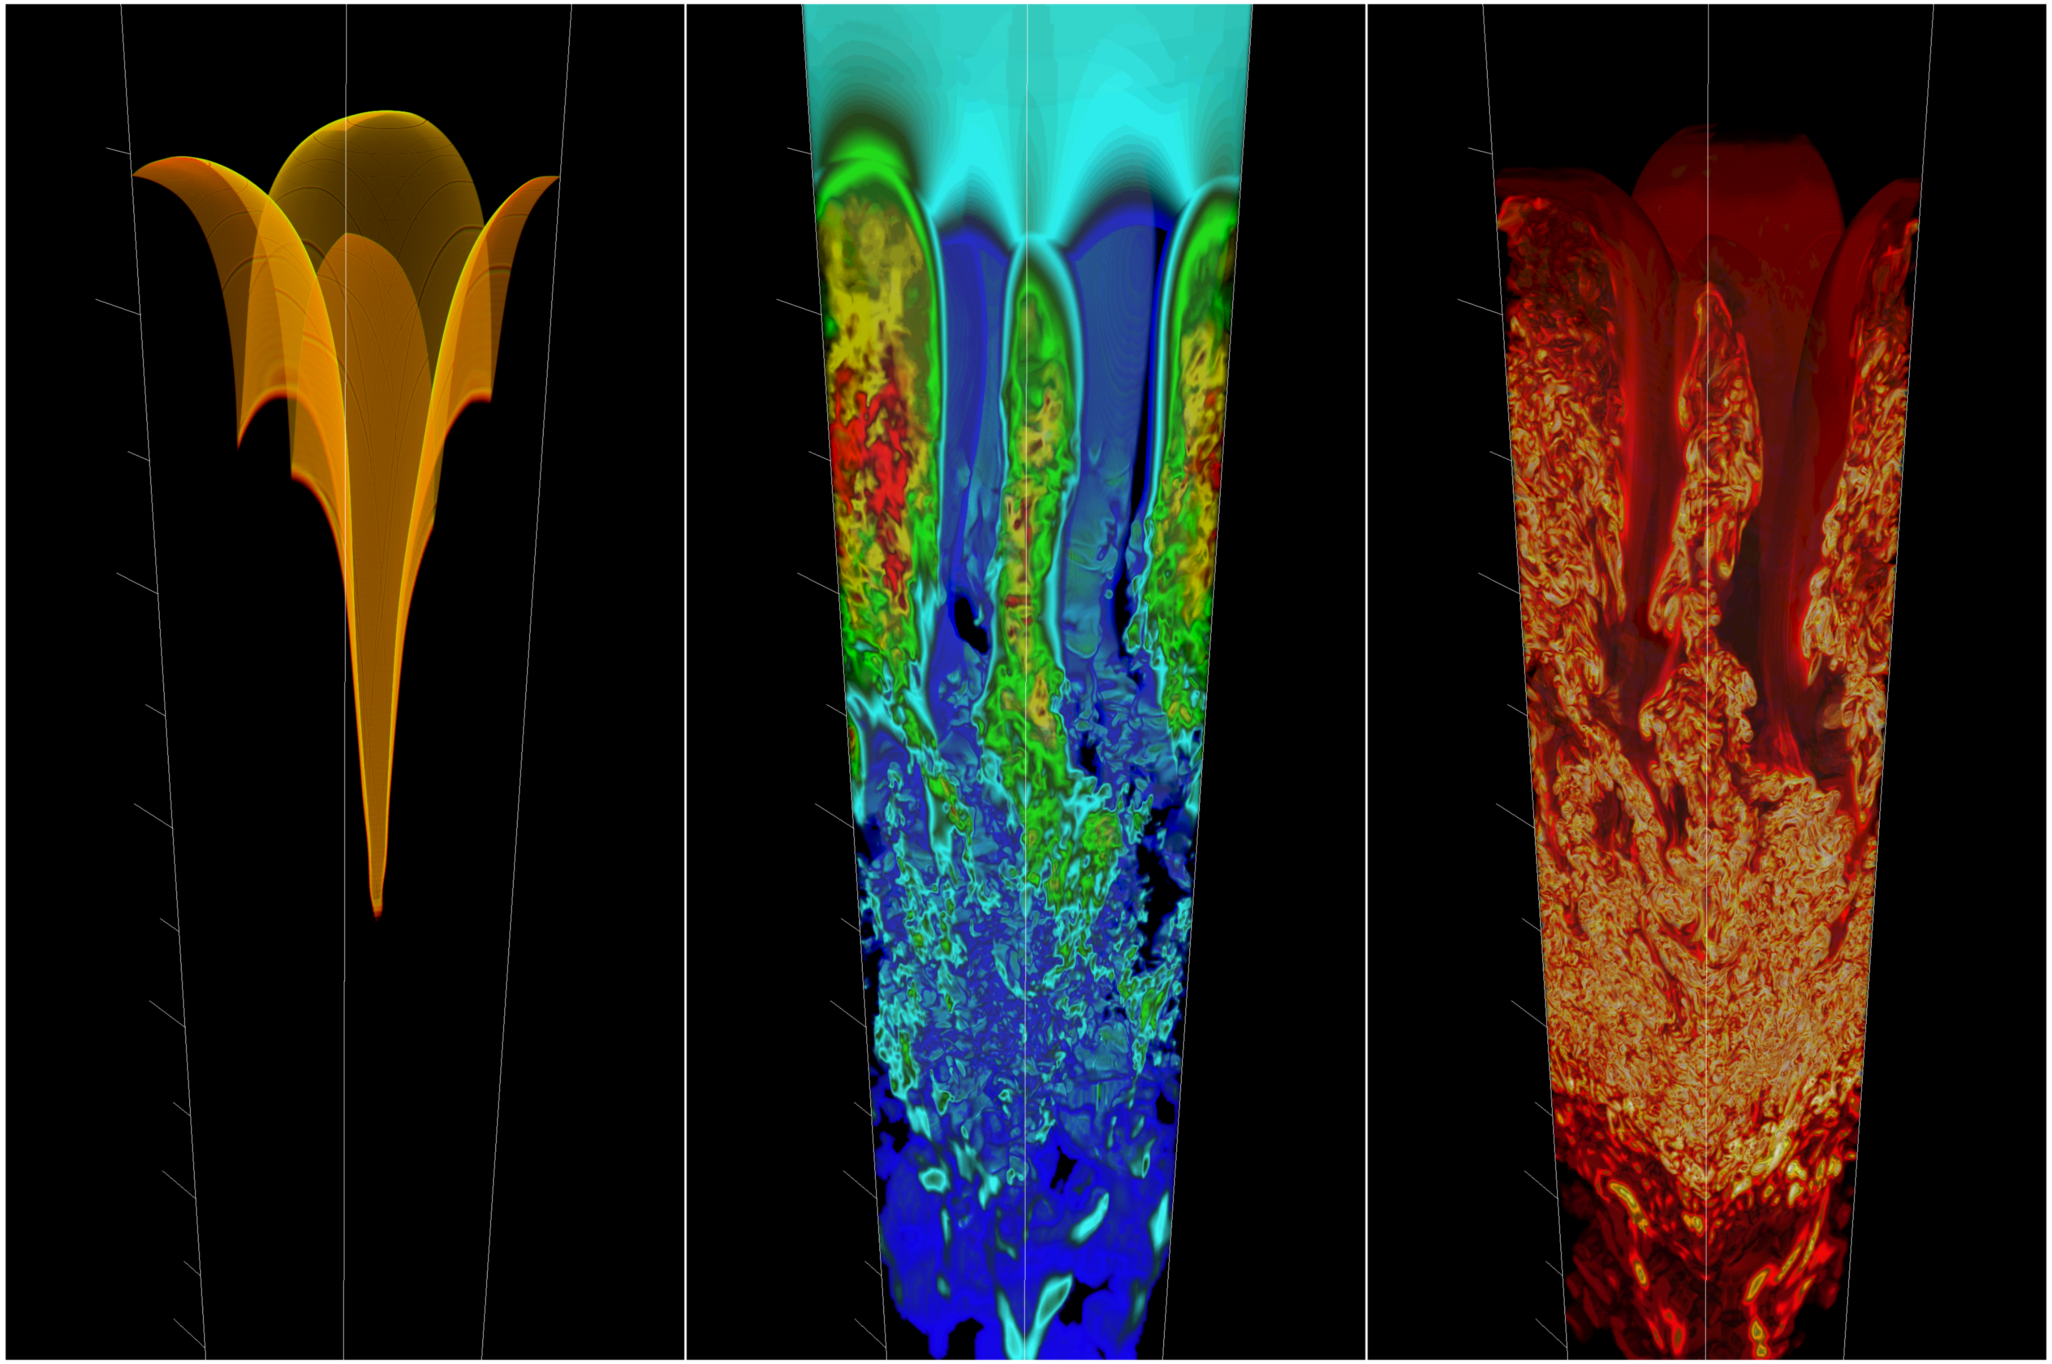

Large-scale FLASH simulations of Buoyancy-driven Turbulent Nuclear Combustion

https://vimeo.com/40691923

Laboratory Astrophysics: HEDP

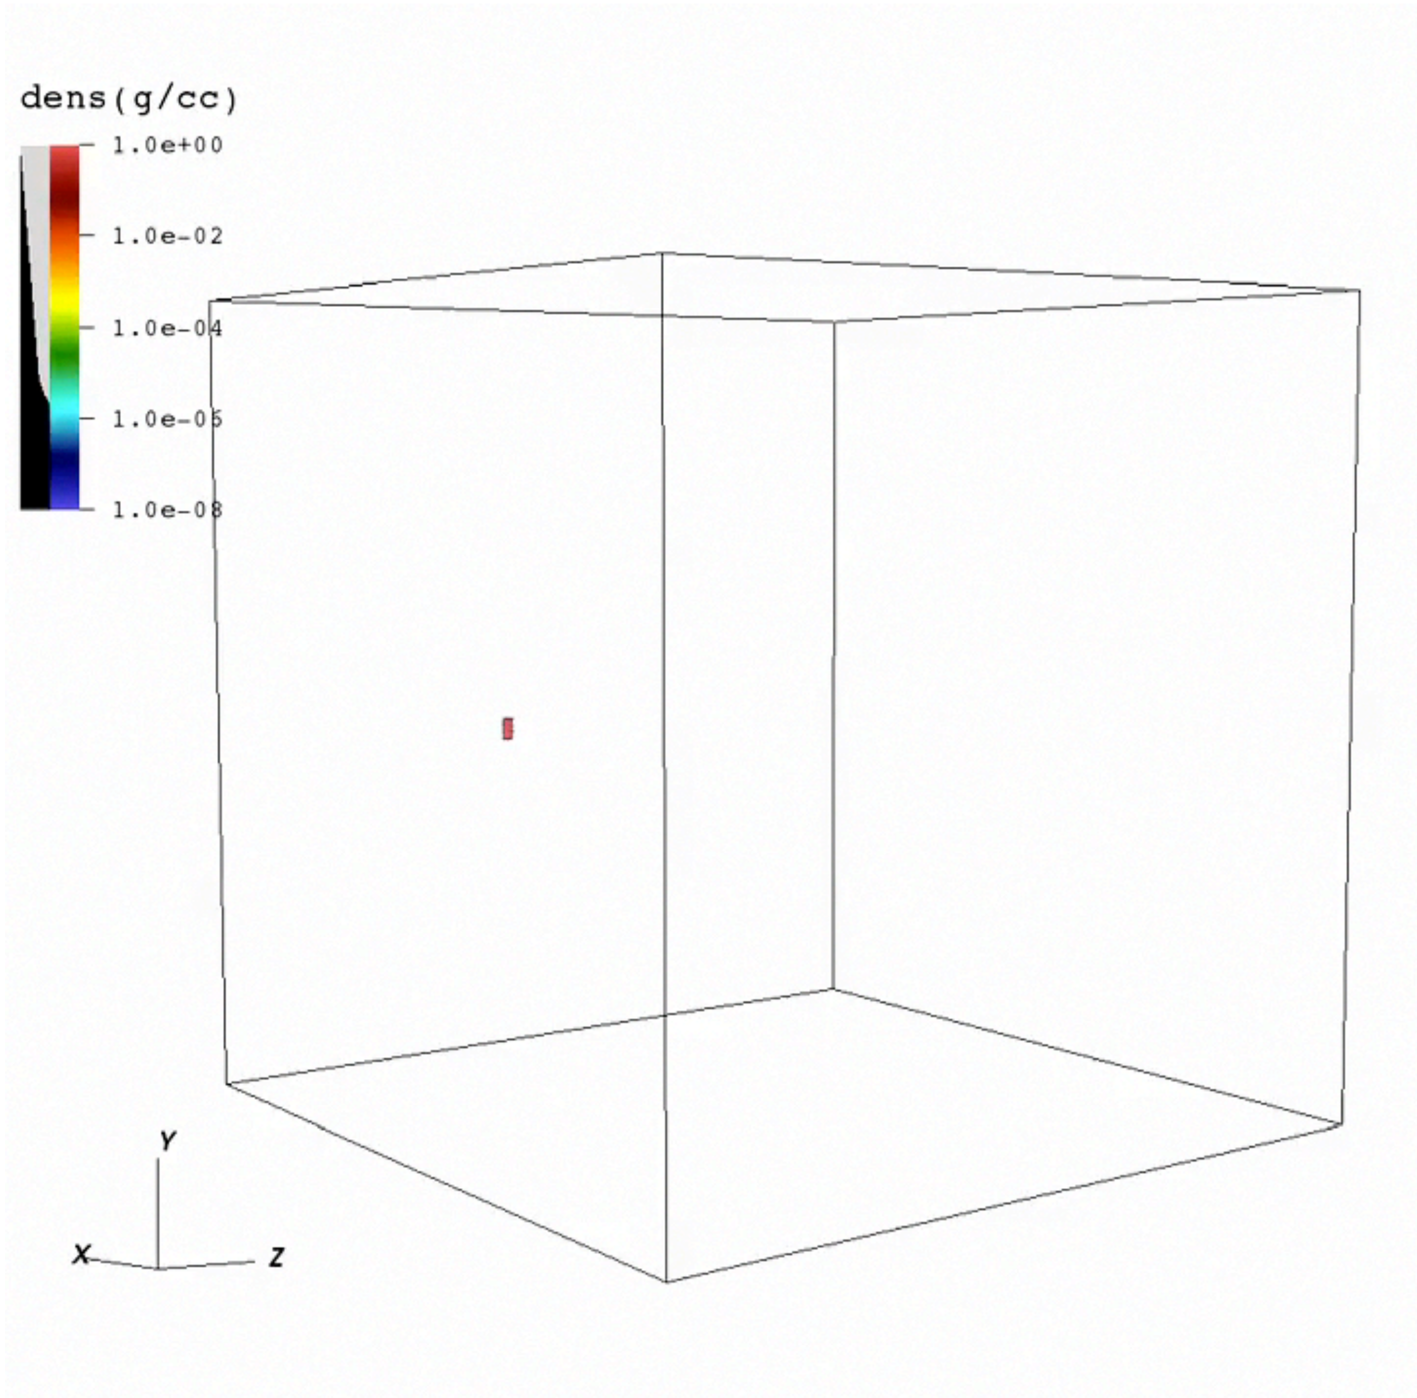

3D Simulation

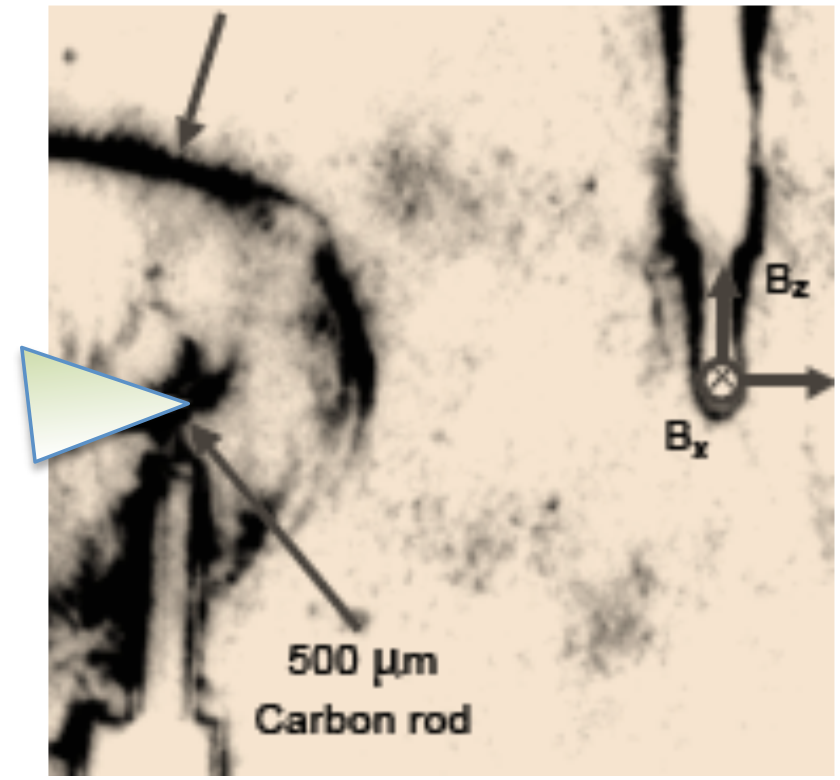

In collaboration with
the research teams in
U of Chicago & Oxford Univ.

Aerodynamics: Supersonic Airflow

Theory

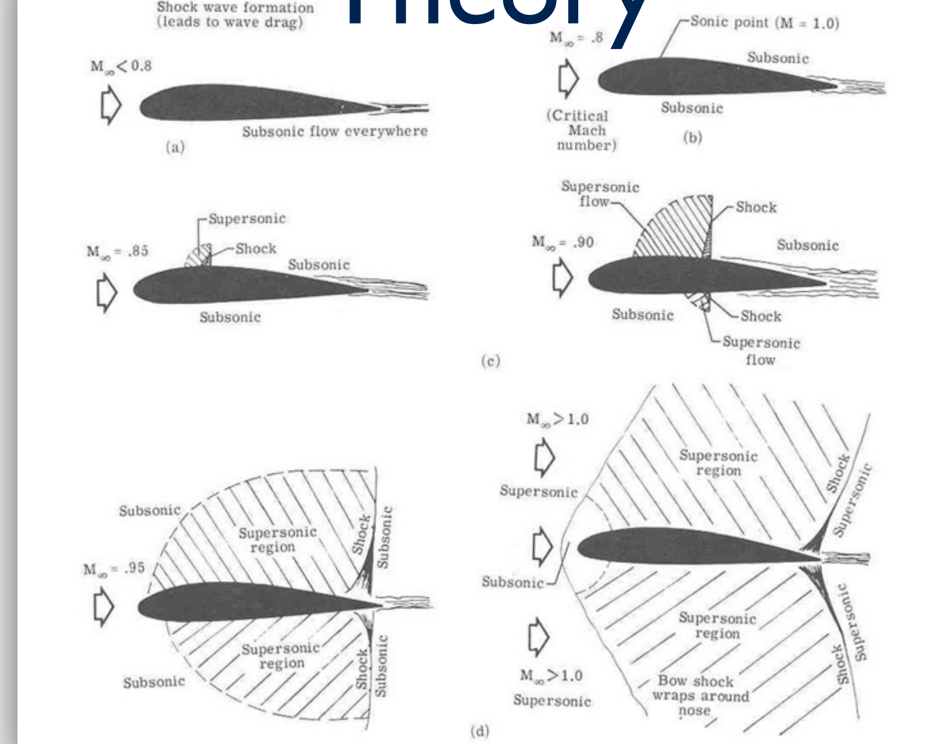

Experiment

Simulation

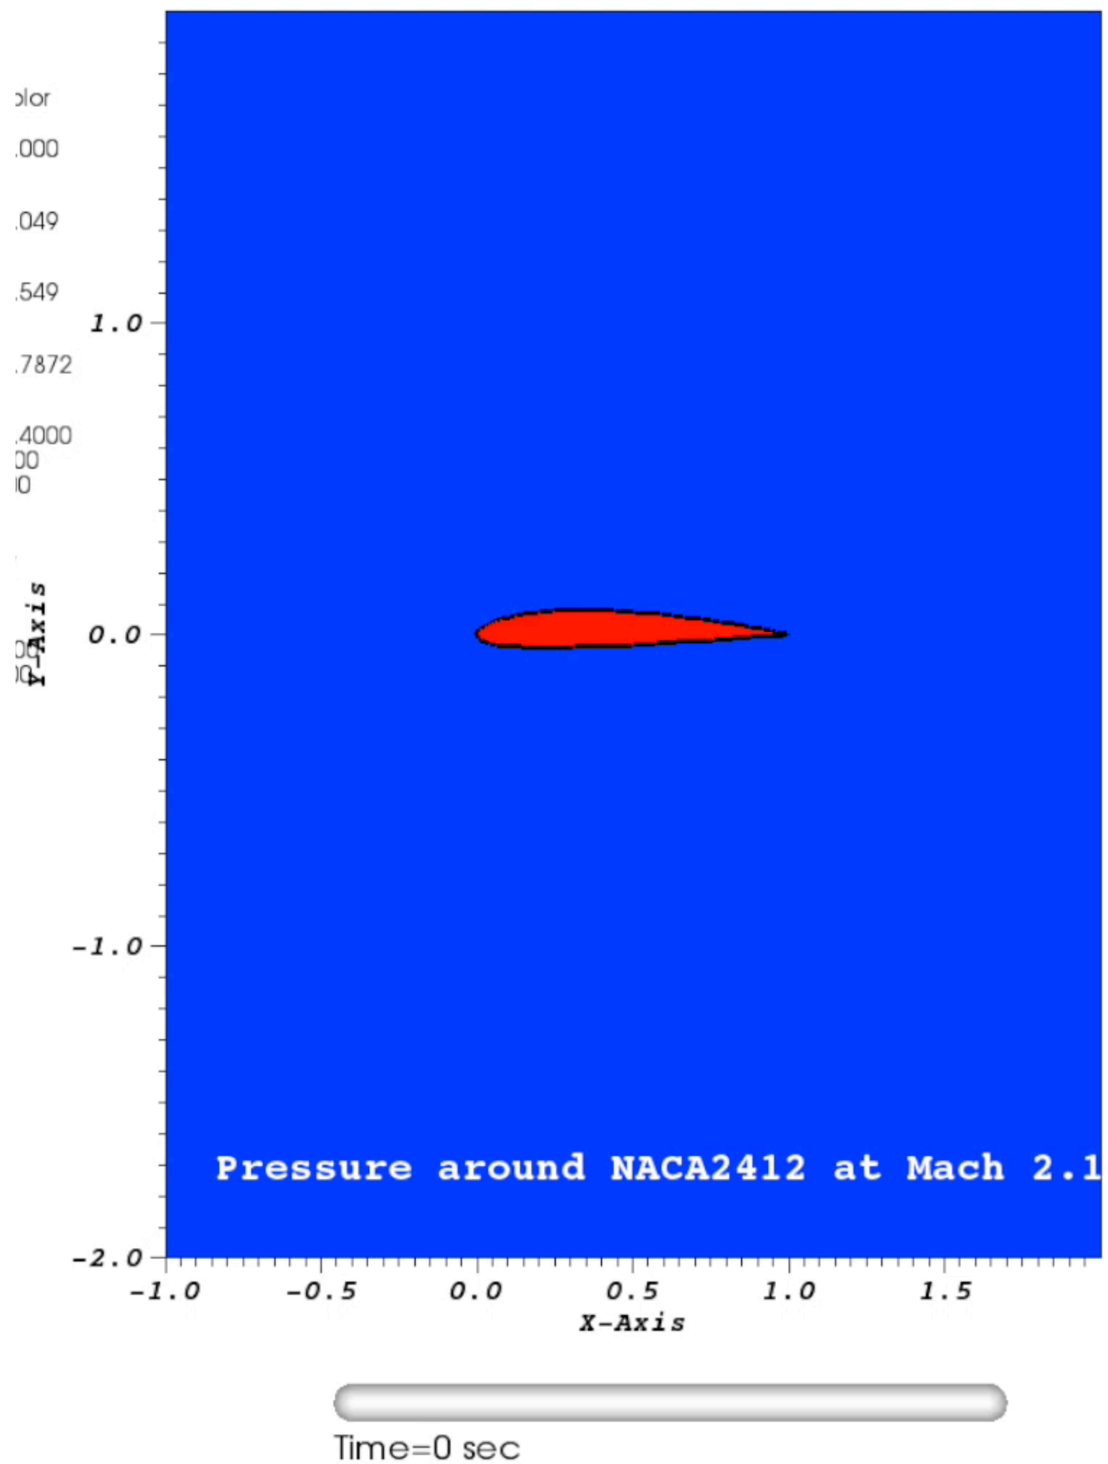

Cyclic Relationship

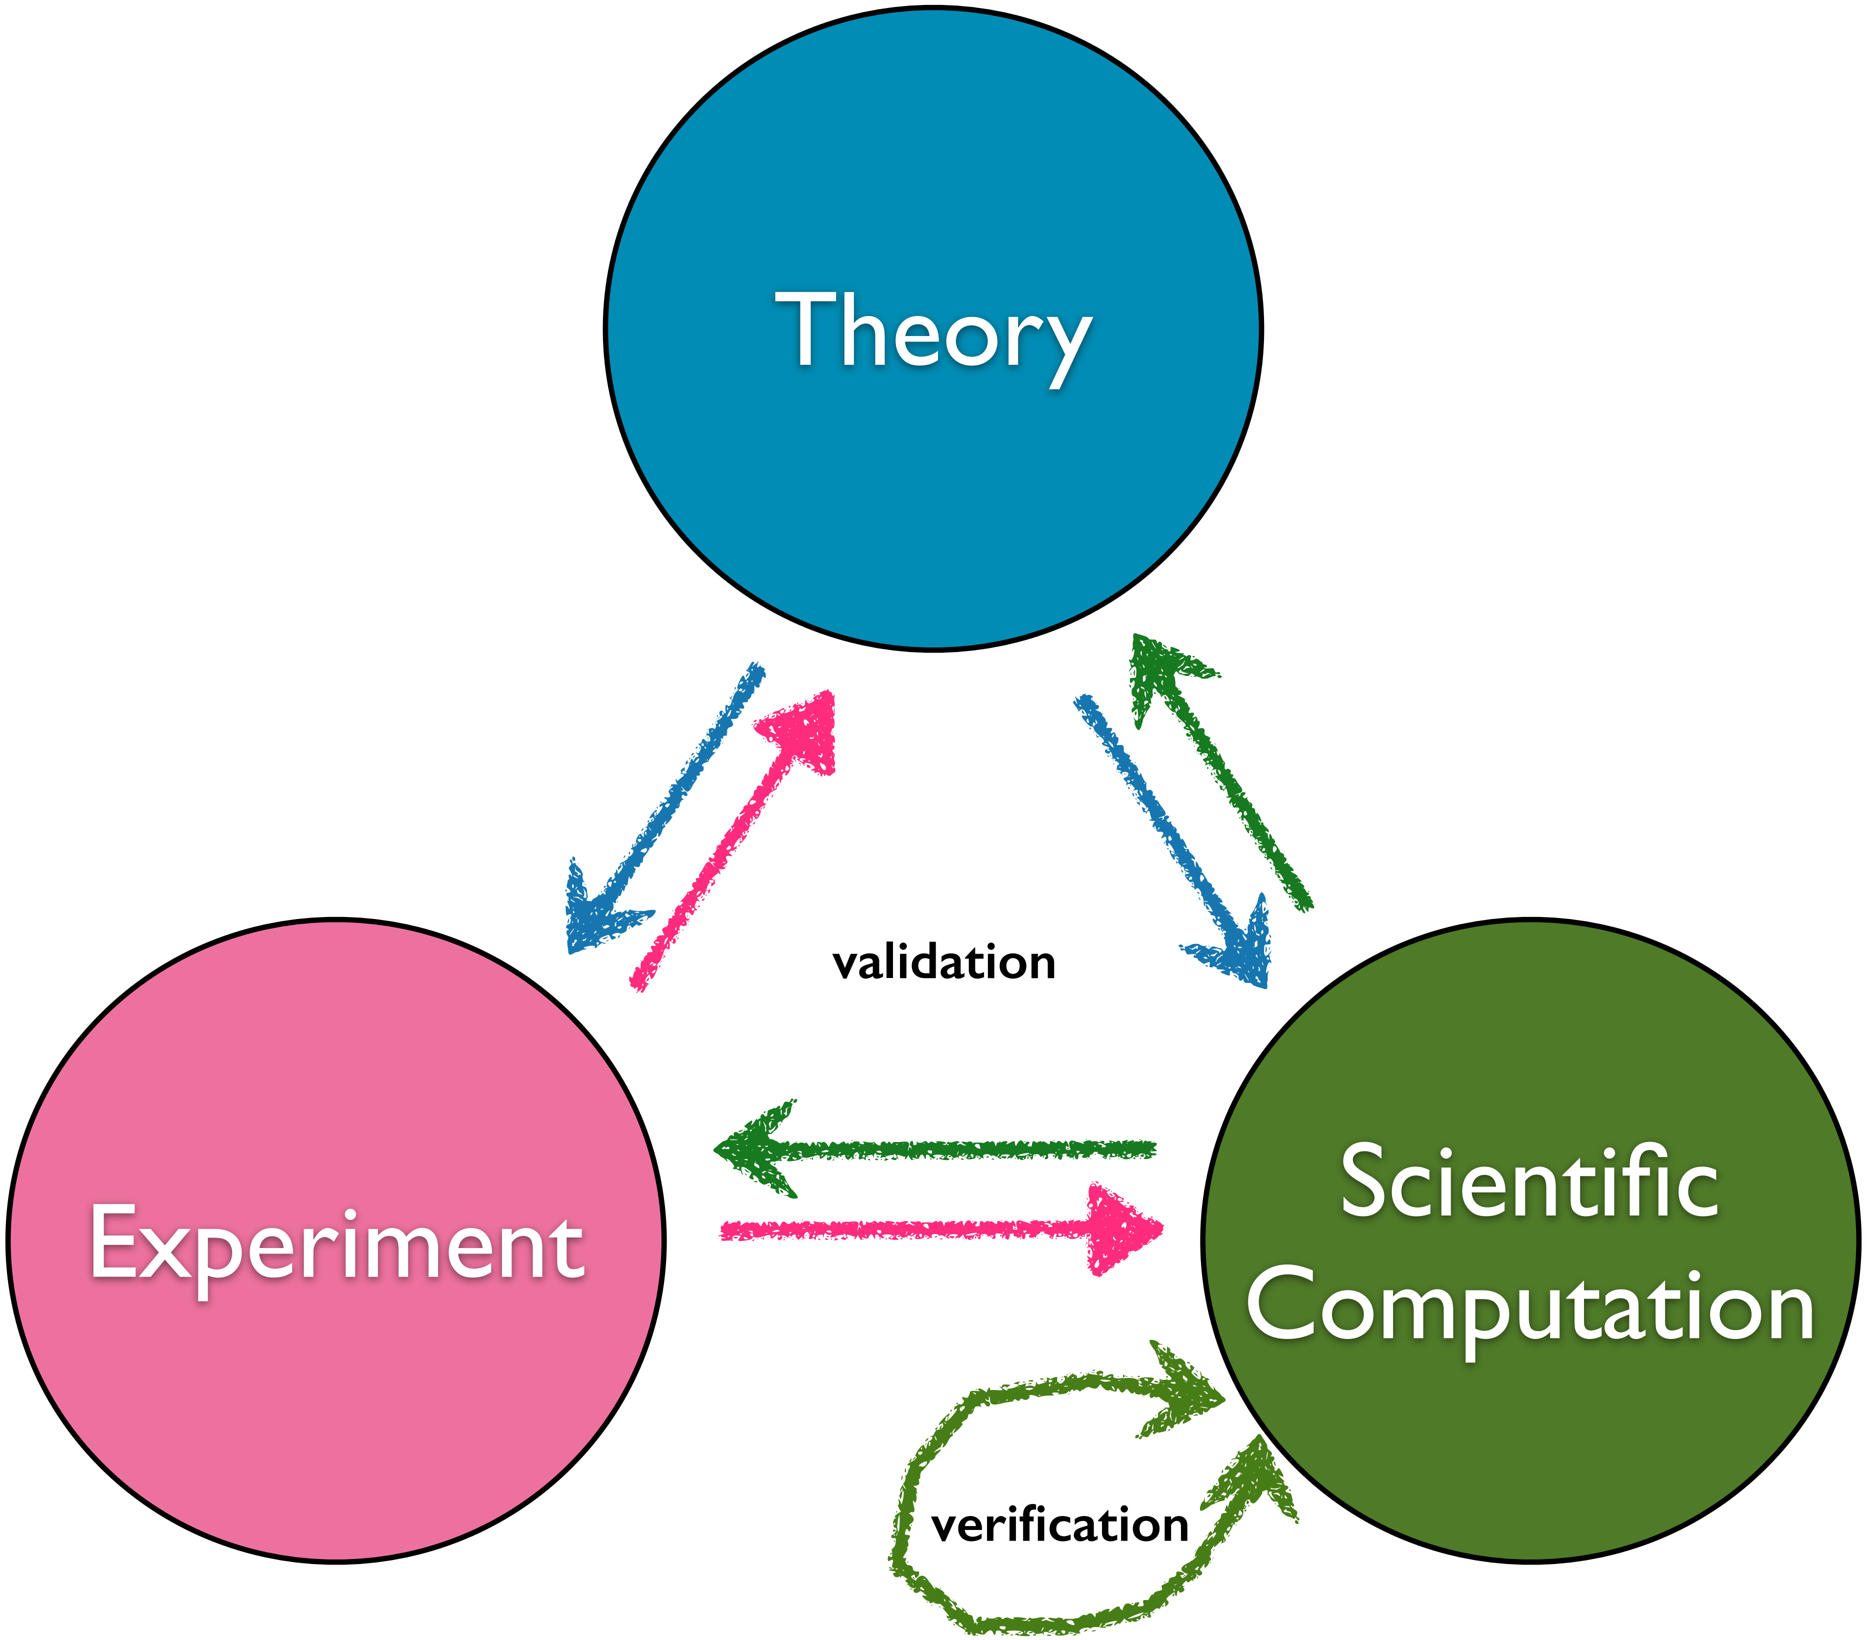

Scientific Computing Tasks

Science Problem
(IC, BC, ODE/PDE)

$$\frac{\partial \mathbf{B}}{\partial t} + \nabla \times \mathbf{E} = 0$$

Simulator
(code, computer)

Results
(Validation, verification, analysis)

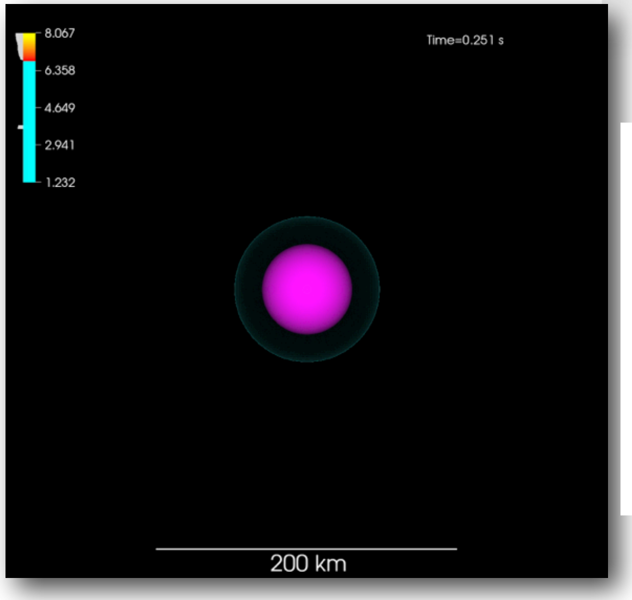

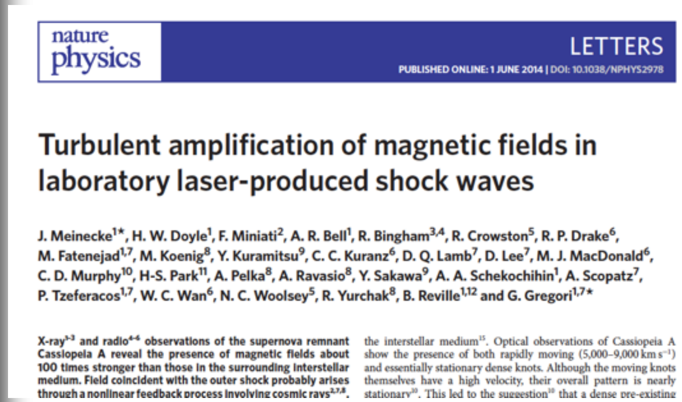
nature physics | LETTERS

PUBLISHED ONLINE: 1 JUNE 2014 | DOI: 10.1038/NPHYS2978

Turbulent amplification of magnetic fields in laboratory laser-produced shock waves

J. Meinecke[1*], H. W. Doyle[1], F. Miniati[2], A. R. Bell[1], R. Bingham[3,4], R. Crowston[5], R. P. Drake[6], M. Fatenejad[1,7], M. Koenig[8], Y. Kuramitsu[9], C. C. Kuranz[6], D. Q. Lamb[7], D. Lee[7], M. J. MacDonald[6], C. D. Murphy[10], H-S. Park[11], A. Pelka[8], A. Ravasio[8], Y. Sakawa[9], A. A. Schekochihin[1], A. Scopatz[7], P. Tzeferacos[1,7], W. C. Wan[6], N. C. Woolsey[5], R. Yurchak[8], B. Reville[1,12] and G. Gregori[1,7*]

X-ray[1-3] and radio[4-6] observations of the supernova remnant Cassiopeia A reveal the presence of magnetic fields about 100 times stronger than those in the surrounding interstellar medium. Field coincident with the outer shock probably arises through a nonlinear feedback process involving cosmic rays[2,7,8].

the interstellar medium[15]. Optical observations of Cassiopeia A show the presence of both rapidly moving (5,000–9,000 km s^{-1}) and essentially stationary dense knots. Although the moving knots themselves have a high velocity, their overall pattern is nearly stationary[10]. This led to the suggestion[10] that a dense pre-existing

Roles of Scientific Computing

- Scientific simulations of multi physics, multi scale phenomena have enabled us to enhance our scientific understanding via
  - advances in modeling and algorithms
  - growth of computing resources

Challenges in HPC

- Maximizing the scientific outcome of simulations, especially on high-performance computing (HPC) resources, requires many trade-offs
- Scientists, engineers, & software developers are often challenged to explore previously unexplored regimes in both physics and computing to maximally gain scientific utilization of large HPC

Today's Topics

- I am going to present a case study with the FLASH code on the following topics:
    - Code architecture
    - parallelization (MPI & threading)
    - optimization (algorithm tweaking & improvements)
    - FLASH performance on HPC

FLASH Code

- FLASH is is free, open source code for astrophysics and HEDP (http://flash.uchicago.edu)

  - modular, multi-physics, adaptive mesh refinement (AMR), parallel (MPI & OpenMP), finite-volume Eulerian compressible code for solving hydrodynamics and MHD

  - professionally software engineered and maintained (daily regression test suite, code verification/validation), inline/online documentation

  - FLASH can run on various platforms from laptops to supercomputing (peta-scale) systems such as IBM BG/P and BG/Q with great scaling over a hundred thousands processors

FLASH Code

- FLASH is is free, open source code for astrophysics and HEDP (http://flash.uchicago.edu)

  - modular, multi-physics, adaptive mesh refin… …llel (MPI & OpenMP), finite-volume Eu… solving hydrodynamics an…

  - prof… …), inline/online

  - …on various platforms from laptops to s…puting (peta-scale) systems such as IBM BG/P and BG/Q with great scaling over a hundred thousands processors

- Over 1.2 million lines of code (Fortran, C, pyton)
- 25% are comments
- Extensive code docs in user's manual (~ 500 pages)
- Great scaling ~ 100,000 procs

FLASH around the World

Baskin Engineering UC SANTA CRUZ

Munich University Observatory
CAST
Computational Astrophysics Group

Germany

http://www.usm.uni-muenchen.de/CAST/people.html
Pawel Ciecielag: Planets formation: planetesimals accretion, gas drag, disk-planet interaction; Cosmology: large scale velocity fields, density-velocity comparisons; Numerical methods: direct n-body (Nbody4++ code, special purpose hardware - GRAPE), hydro (PPM, AMR, FLASH code).
Steffi Walch: Star formation with emphasis on the formation of protostellar disks, evolution and fragmentation of protoplanetary

Practical FLASH Notes 10/14/2005 04:14 PM

Learning and Testing FLASH: Practical Notes

This page is intended for members of the Computational Astrophysics group at McMaster University who whish to learn and be involved with the testing of the FLASH code. The code was developed at the ASCI/Alliances Center for Astrophysical Thermonuclear Flashes at the University of Chicago.

Getting Started:

Canada

-step spoon-feed for ıs along the way.

THE FLASH CODE AT OAPA 10/14/2005 01:21 PM

OSSERVATORIO ASTRONOMICO DI PALERMO GIUSEPPE S. VAIANA

The FLASH code at OAPA

Italy

The Osservatorio Astronomico di Palermo is one of the test sites for the very accurate FLASH code, a 3-dimensional astrophysical hydrodynamic code for supercomputers mainly developed at the "Flash Center", the University of Chicago.

Palermo is the test site in which FLASH has been ported to Compaq architectures. The FLASH code solves the compressible Euler equations on a block-structured adaptive mesh, and its modular design permits the introduction of additional physics and of different solvers.

The Palermo team collaborates with the Flash Center to upgrade and to apply extensively FLASH to astrophysical systems.

The group in Palermo also develops new modules for FLASH which extend the field of applicability of the code to other problems in astrophysics, from solar and stellar coronae, to supernova remnants, and to galaxy clusters halos. In particular, the new modules so far developed and tested include:

1. the non-equilibrium ionization effects of the most abundant elements in astrophysical plasmas
2. the thermal conduction according to the formulation of Spitzer (1962)
3. the radiative losses from an optically thin plasma according to the Raymond spectral code and Peres et al. (1982) for the chromosphere.
4. the viscosity according to the formulation of Spitzer (1962)

In this project, the Palermo team takes advantage of its long experience in developing hydrodynamic codes for modeling astrophysical plasma and in optimizing the codes for efficient parallel execution on high performance computers. The group has recently acquired and uses, for the FLASH development, a high performance computing (HPC) cluster of 16 powerful alpha EV67 processors distributed in 4 compaq ES40 (interconnected with a highly efficient Memory Channel II), entirely dedicated to HPC projects (for more information see the SCAN facility homepage).

http://www.astropa.unipa.it/FLASH/ Page 1 of 1

THE ASTROPHYSICAL JOURNAL, 628:205–209, 2005 July 20

SUPERNOVA BLAST WAVES IN LOW-DENSITY HOT MEDIA: A MECHANISM FOR SPATIALLY DISTRIBUTED HEATING

SHIKUI TANG AND Q. DANIEL WANG
Department of Astronomy, University of Massachusetts, LGRT-B 619E, 710 North Pleasant Street, Amherst, MA 01003; tangsk@astro.umass.edu, wqd@astro.umass.edu
Received 2004 December 17; accepted 2005 April 11

ABSTRACT

Most supernovae are expected to explode in low-density hot media, particularly in galactic bulges and elliptical galaxies. The remnants of such supernovae, although difficult to detect individually, can be profoundly important in heating the media on large scales. We characterize the evolution of this kind of supernova remnant, based on analytical approximations and hydrodynamic simulations. We generalize the standard Sedov solution to account for both temperature and density effects of the ambient media. Although cooling can be neglected, the expansion of such a remnant deviates quickly from the standard Sedov solution and asymptotically approaches the ambient sound speed as the swept-up thermal energy becomes important. The relatively steady and fast expansion of the remnants over large volumes provides an ideal mechanism for spatially distributed heating, which may help to alleviate the overcooling problem of hot gas in groups and clusters of galaxies, as well as in galaxies themselves. The simulations were performed with the FLASH code.

Subject headings: cooling flows — galaxies: clusters: general — galaxies: ISM — ISM: structure — supernova remnants

Online material: color figures

CITA–ICAT
ABOUT THE INSTITUTE
RESEARCH @ CITA
WORKING @ CITA
EVENTS & CALENDAR
Archive
Seminars
Lunch Talks
Meetings &

10/14/2005 01:30 PM

FLASH Code Development

ım Members:

ı Boss (DTM)
abeth Myhill, Justin Domes (Marymount University)
ri Vanhala (Challenger Center)

US

: FLASH Adaptive Mesh Refinement Hydrodynamics Code

seek to study the dynamics of mixing and transport processes in the presolar cloud and in the solar ıla, in the context of isotopic heterogeneity introduced either by shock-triggered collapse of the presolar d or by infall from an x-wind outflow, the two leading explanations for the widespread evidence of t-lived radioactivities in chondritic refractory inclusions and, much more rarely, in chondrules. Myhill Domes are leading our effort to develop a new hydrodynamical code for studying these problems, the SH adaptive mesh refinement code. FLASH will allow the problem of shock-wave triggering and tion to be studied with an unprecedented degree of high spatial resolution, which is likely to be critical e question of simultaneous triggering and injection when nonisothermal shock front thermodynamics is loyed. We will work with Vanhala to define a nonisothermal shock test case to be calculated with both FLASH code and Vanhala's EVH-1 code. In addition, the FLASH code will permit extending these stigations down to the scale of the solar nebula, where nebular transport and mixing processes can be ied in general terms, applicable to isotopically heterogeneous grains falling onto the nebular surface as a lt of either shock-triggered collapse or x-wind outflows. We will seek in part to learn whether spatial temporal heterogeneity inherited from such sources can survive subsequent nebular mixing processes can help to explain certain isotopic abundance patterns seen in the inner Solar System and the asteroid The development of the FLASH code should also prove useful for Boss's studies of the formation of t planets by the disk instability mechanism.

EVENTS & CALENDAR

FLASH Workshop

Canada

Workshop
Wed, Mar 23, 2005, 9:00 AM
Location: MP1318A

Workshop on the FLASH code: a general-purpose, parallel, adaptive mesh, reactive astrophysical hydrodynamics code that is available to the research community

Tentative Schedule

Abstract:

On March 23 and 24th, CITA will host a mini-workshop on the FLASH code: a general-purpose, parallel, adaptive mesh, reactive astrophysical hydrodynamics code that is available to the research community. The FLASH code has been applied to problems from burning in or on compact objects, to the efficiency of cooling flows in galaxy clusters, to cosmological simulations.

Four people will be coming from the Chicago FLASH centre to give one day of talks on the capabilities of the code, and will be available to schedule meetings with on the second day for more focused, less formal discussions on applying the FLASH code to problems of interest to local projects.

3D AMR Simulations of Point-Symmetric Nebulae

Erik-Jan Rijkhorst, Vincent Icke, and Garrelt Mellema, Sterrewacht Leiden, The Netherlands
http://www.strw.LeidenUniv.nl/AstroHydro3D/

Numerical Implementation

We used the three-dimensional hydrocode *Flash* [6] to model the interaction between a spherical wind and a warped disk. This parallelized code implements block-structured adaptive mesh refinement (AMR) [7] (see images below) and a PPM type hydrosolver [8].

We added to the code the proper initial conditions for the wind-di… warped disk, Eq. (1) was combined with a constant 'wedge angle… was given a constant density. The spherical wind was implemen… $1/r^2$ density profile and a constant velocity. The pressure was cal… constant Poisson index γ.

Netherlands

THE ASTROPHYSICAL JOURNAL, 630:740–749, 2005 September 10

Germany

ACTIVE GALACTIC NUCLEI HEATING AND DISSIPATIVE PROCESSES IN GALAXY CLUSTERS

M. BRÜGGEN
International University Bremen, Campus Ring 1, 28759 Bremen, Germany; m.brueggen@iu-bremen.de
M. RUSZKOWSKI[1]
Joint Institute for Laboratory Astrophysics, Campus Box 440, University of Colorado, Boulder, CO 80309-0440; mr@quixote.colorado.edu
AND
E. HALLMAN
Center for Astrophysics and Space Astronomy, University of Colorado, Boulder, CO 80309; hallman@origins.colorado.edu
Received 2005 January 10; accepted 2005 May 22

ABSTRACT

Recent X-ray observations reveal growing evidence for heating by active galactic nuclei (AGNs) in clusters and groups of galaxies. AGN outflows play a crucial role in explaining the riddle of cooling flows and the entropy problem in clusters. Here we study the effect of AGNs on the intracluster medium in a cosmological simulation using the adaptive mesh refinement FLASH code. We pay particular attention to the effects of conductivity and viscosity on the dissipation of weak shocks generated by the AGN activity in a realistic galaxy cluster. Our three-dimensional simulations demonstrate that both viscous and conductive dissipation play an important role in distributing the mechanical energy injected by the AGNs, offsetting radiative cooling and injecting entropy to the gas. These processes are important even when the transport coefficients are at a level of 10% of the Spitzer value. Provided that both conductivity and …ity are suppressed by a comparable amount, conductive dissipation is likely to dominate over viscous dissipation. …heless, viscous effects may still affect the dynamics of the gas and contribute a significant amount of dissipation …red to radiative cooling. We also present synthetic *Chandra* observations. We show that the simulated buoyant …s inflated by the AGN, and weak shocks associated with them, are detectable with the *Chandra* observatory.

…t *headings:* cooling flows — galaxies: active — galaxies: clusters: general — X-rays: galaxies

Monthly Notices of the Royal Astronomical Society
Volume 355 Issue 3 Page 995 - December 2004
doi:10.1111/j.1365-2966.2004.08381.x

UK

Quenching cluster cooling flows with recurrent hot plasma bubbles
Claudio Dalla Vecchia[1], Richard G. Bower[1], Tom Theuns[1,2], Michael L. Balogh[1], Pasquale Mazzotta[3] and Carlos S. Frenk[1]

Papers using FLASH

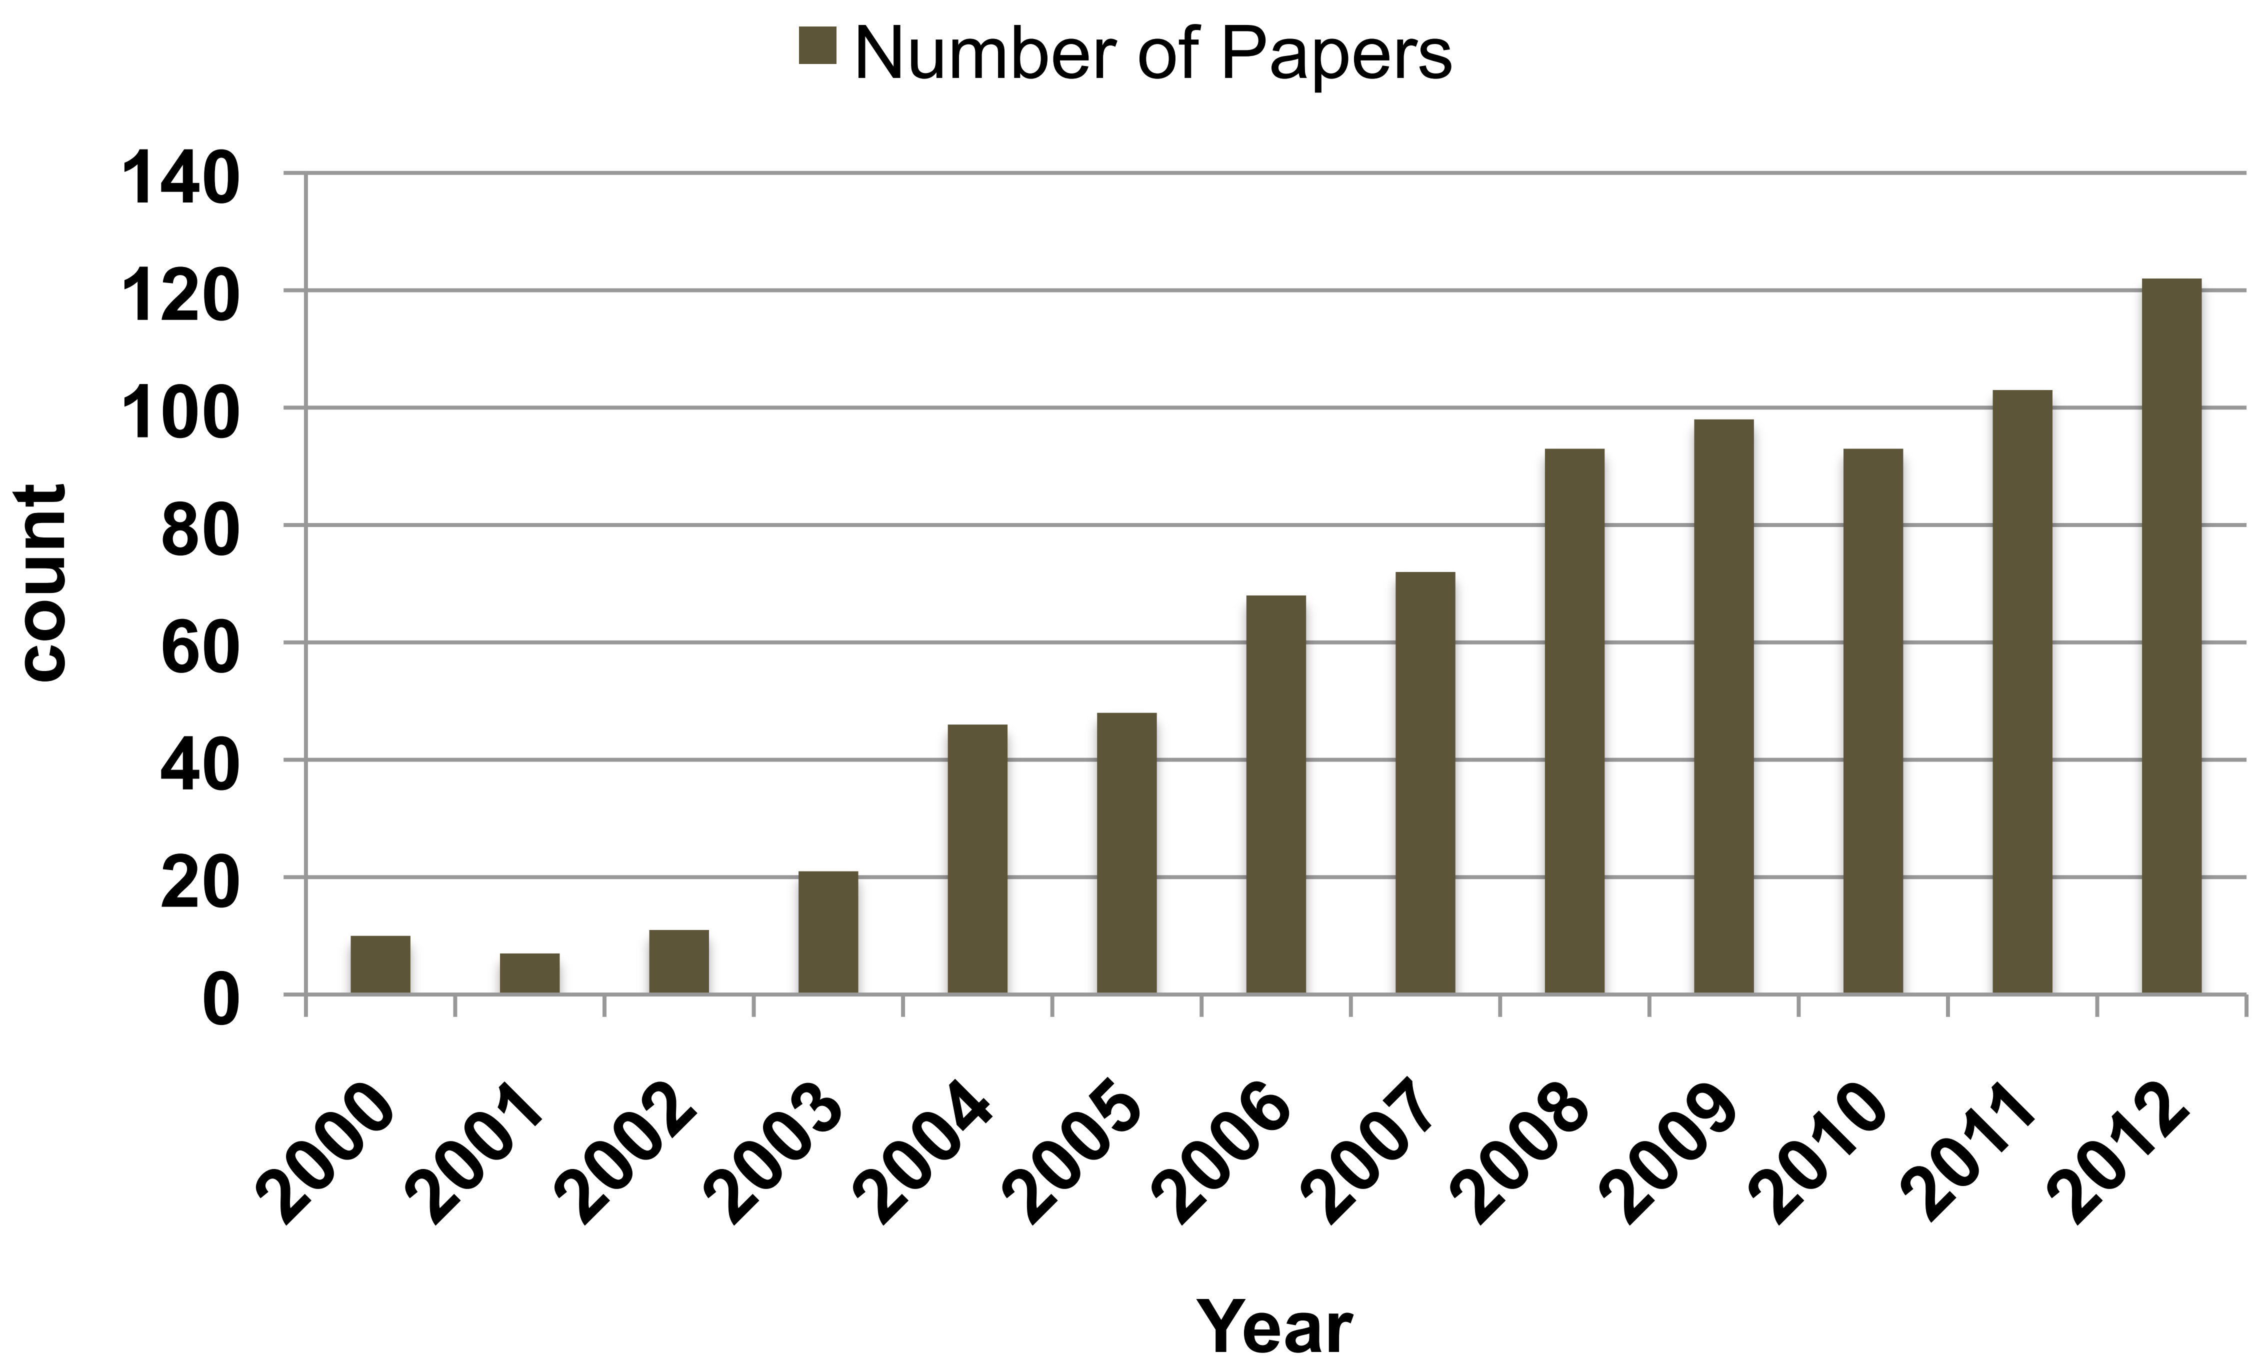

downloads > 8500; authors > 1500; papers > 1000

Research Applications

- Major Research Applications
  - thermonuclear flashes
  - high energy density physics (HEDP) - B field amplification
  - fluid-structure interaction
  - star formation
  - star-star & star-planets interactions
  - cosmology
  - galaxy & galaxy cluster simulations
  - turbulence
  - …

Scientific Simulations using FLASH

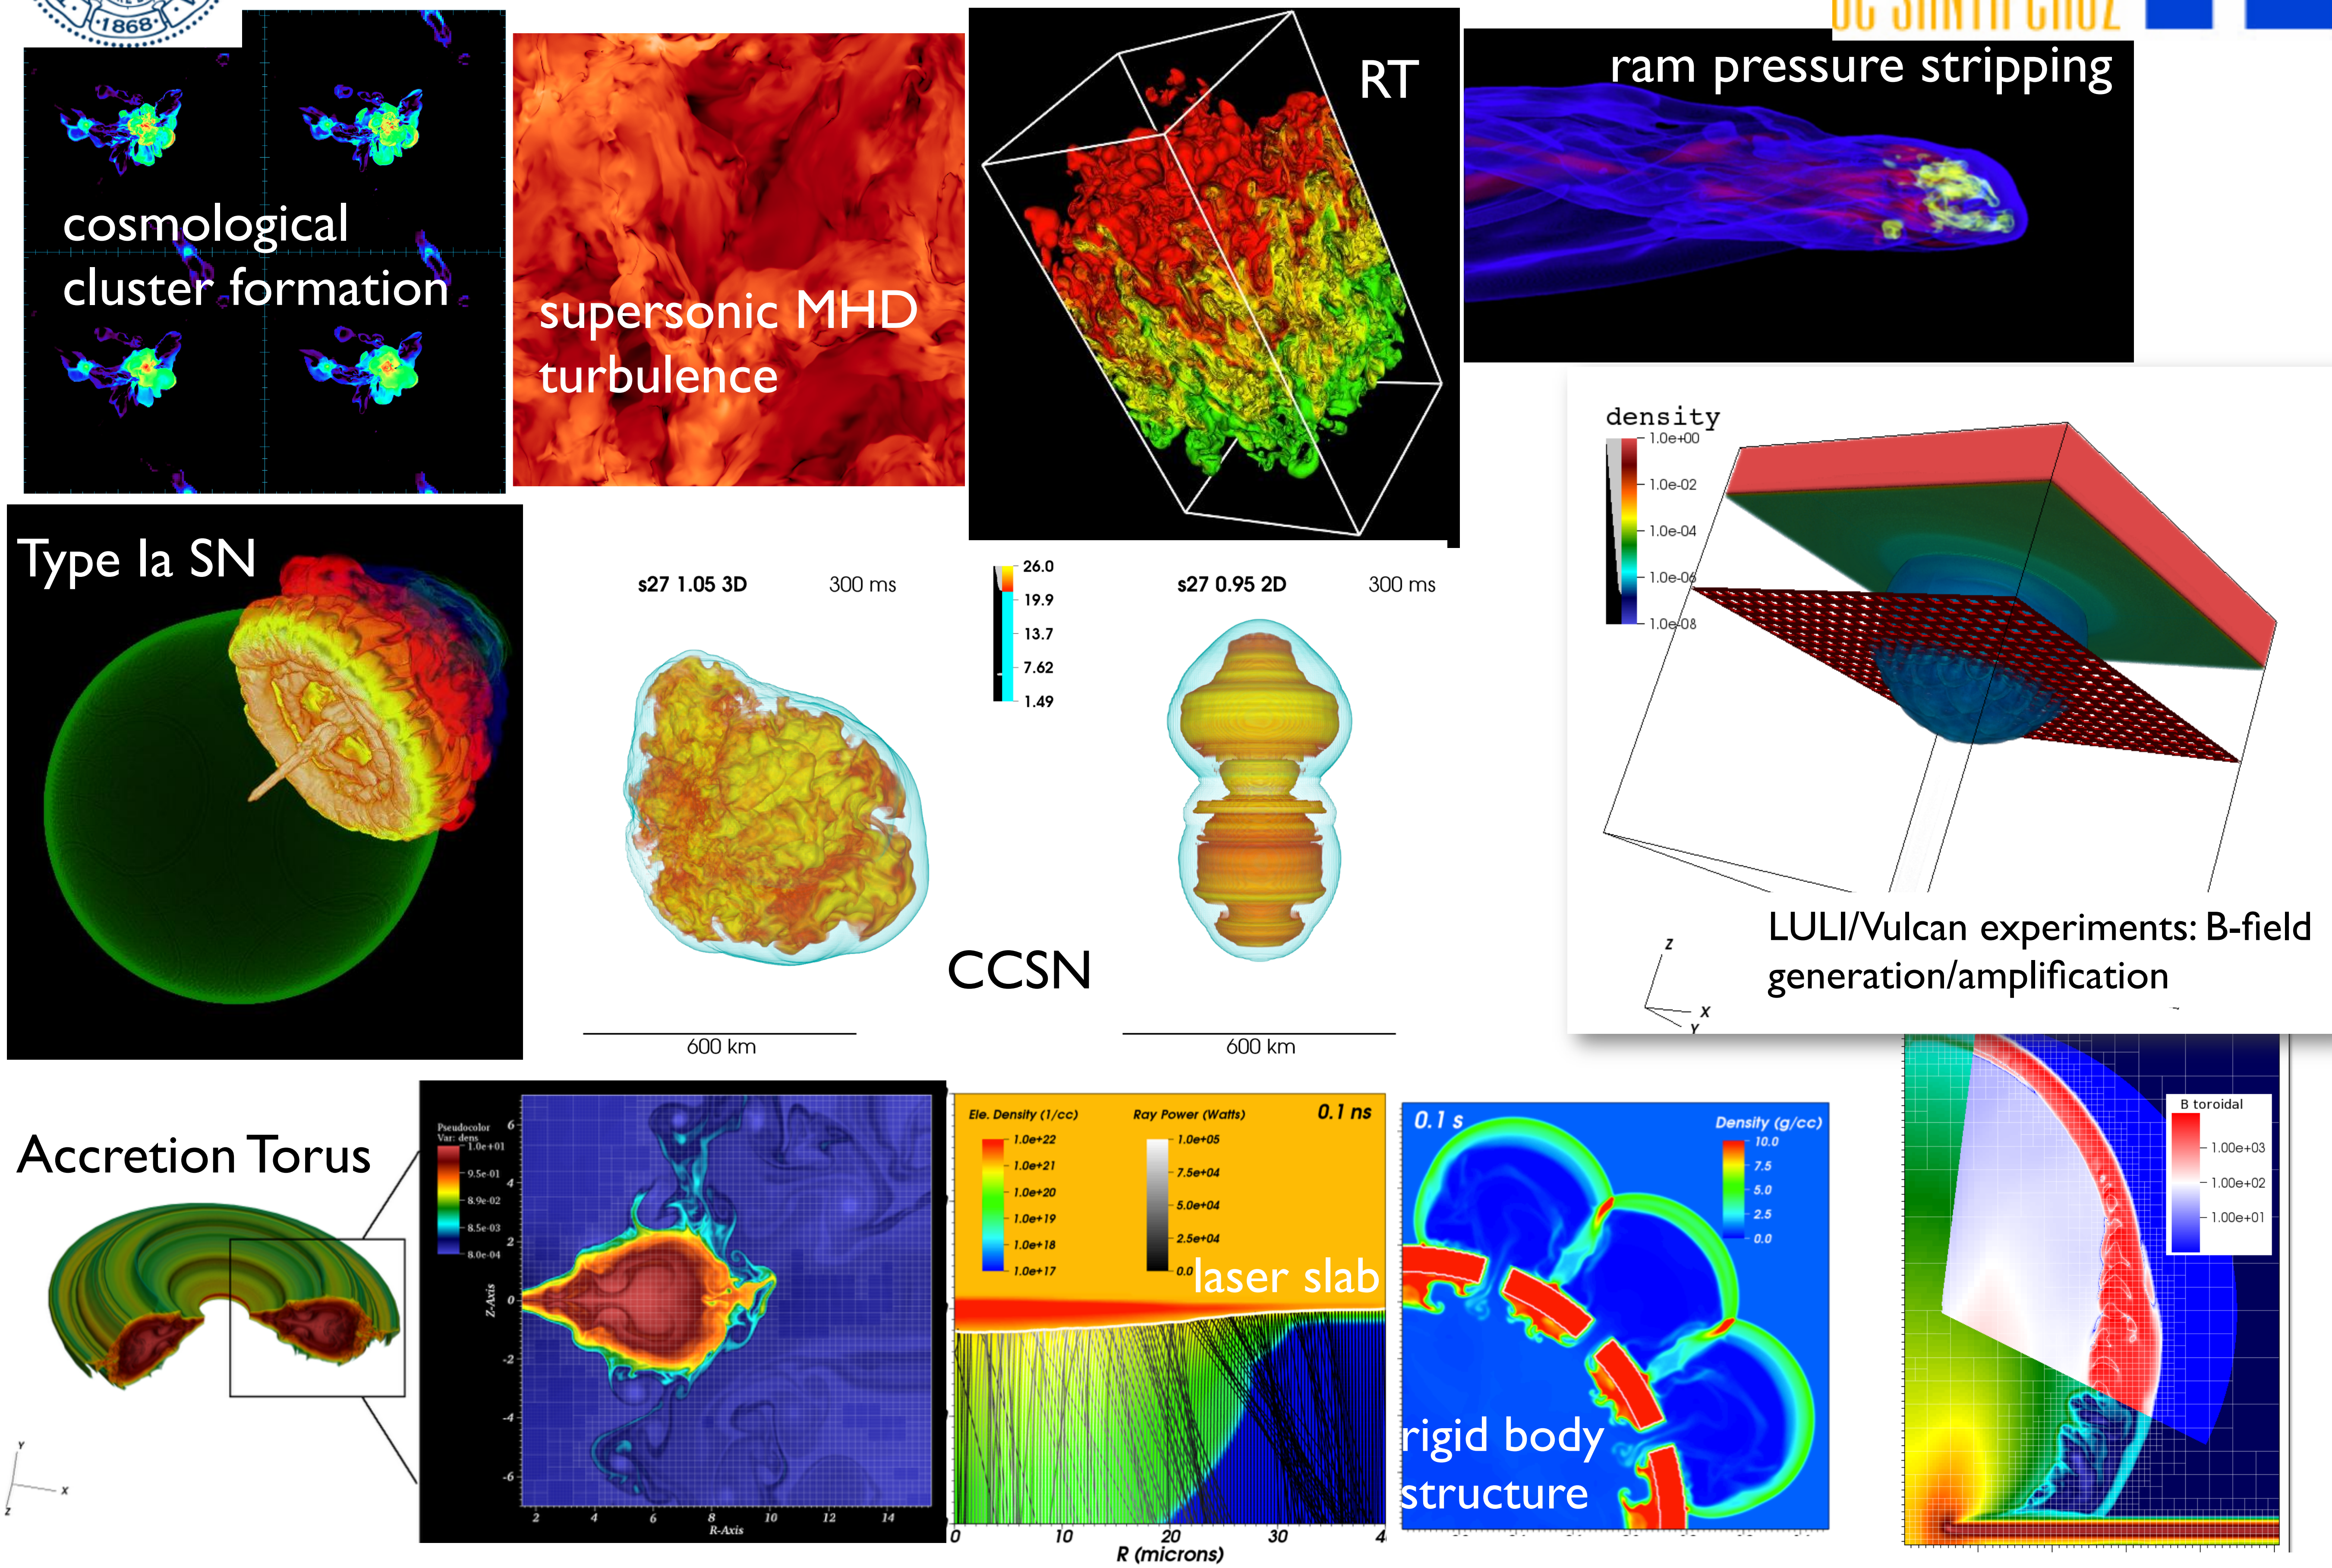

History

- 1997: Founding of ASCI (Accelerated Strategy Computing Initiative) FLASH Center
  - thermonuclear flashes at neutron stars and WD surfaces

- 2000: first FLASH code release
  - Fryxell, Olson, Ricker, Timmes, Zingale, Lamb, MacNeice, Rosner, Tururan, Tufo (The FLASH Paper, ApJS)
  - PPM hydro FVM code, no MHD
  - no-self gravity
  - AMR using PARAMESH

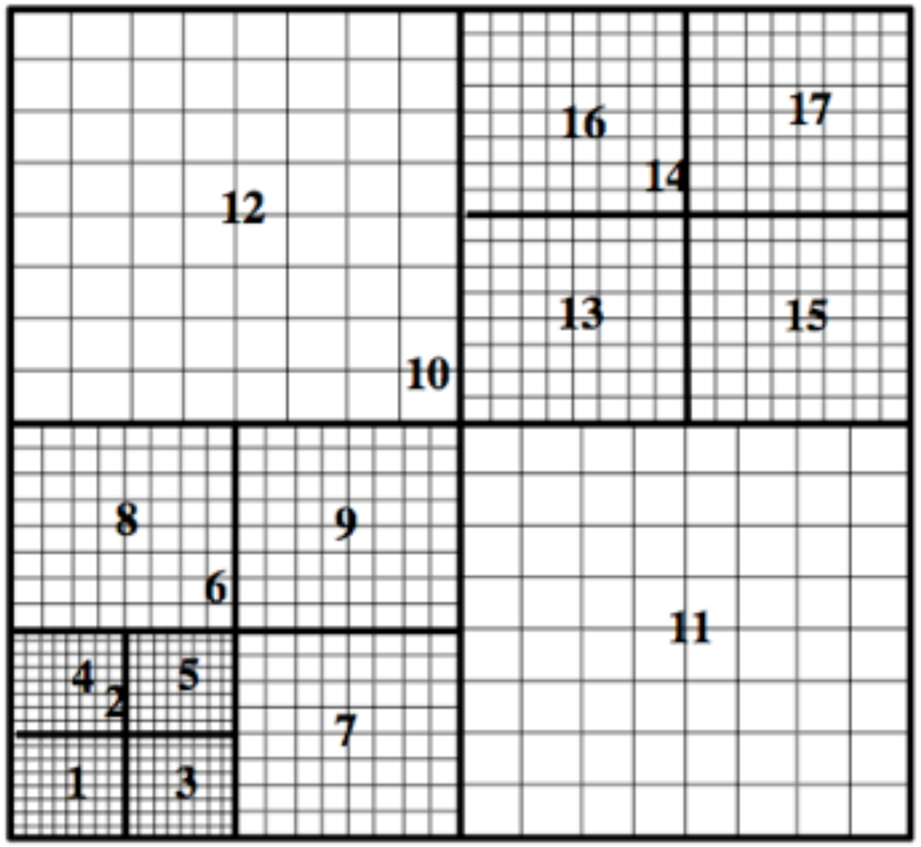

History

- 2002: Major revision updated to 2.3
  - MHD solver based on the 8-wave scheme: directionally split, FVM (Powell, Roe, Linde, Gombosi, De Zeeuw, 1999)
  - self-gravity/Poisson solver (Ricker)
  - tracer particle
  - cosmology
  - HDF5 I/O

History

- 2008: completely restructured version 3
  - unsplit HD & USM-MHD solvers (Dongwook, JCP 2009; 2013)
  - more flexible grid structures (UG & AMR)
  - decentralized "database" structure
  - reorganized directory structure

- 2011: Version 4 (Last released ver. 4.2.2, 2014)
  - full 3D AMR for MHD
  - many new physical modules
    - e.g., implicit diffusion solver for thermal conduction, energy deposition via laser, 3-temperature for HEDP, solid-boundary interface

Out-of-Box Examples

Domain Decomposition

- **Adaptive Mesh Refinement** (w/ Paramesh)
  - conventional parallelism via MPI (Message Passing Interface)
  - domain decomposition distributed over multiple processor units
  - distributed memory (cf. shared memory)

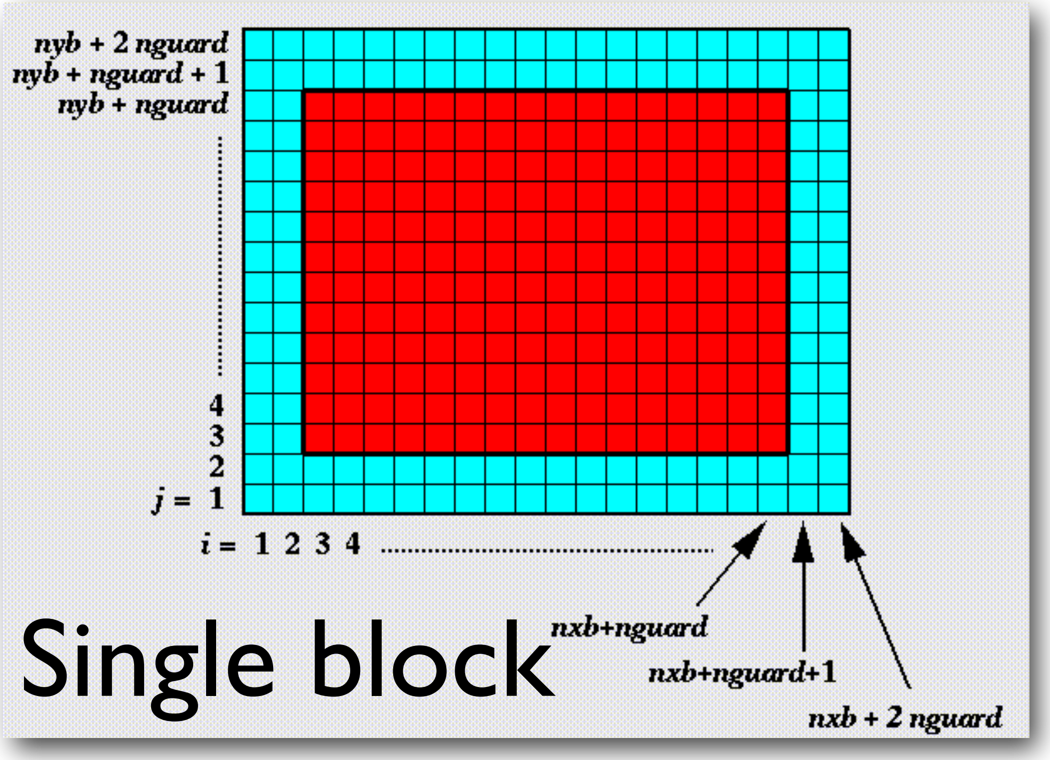

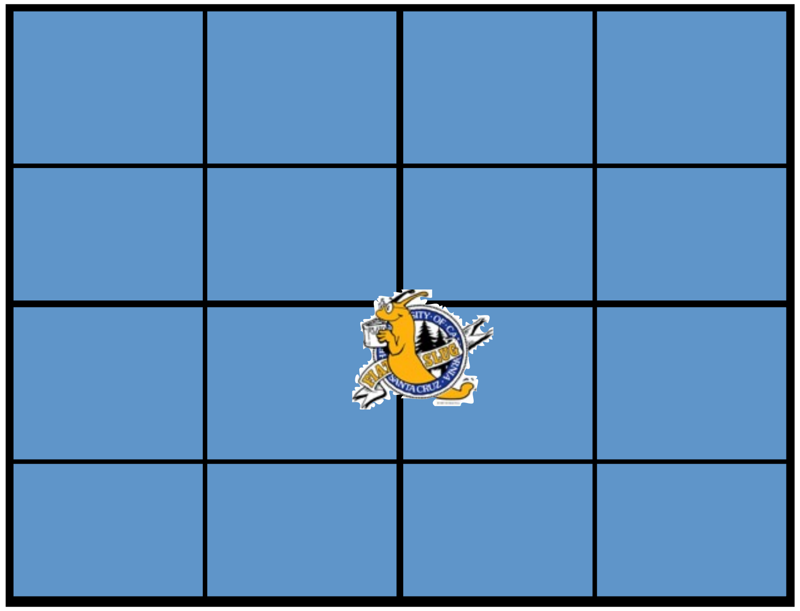

uniform grid

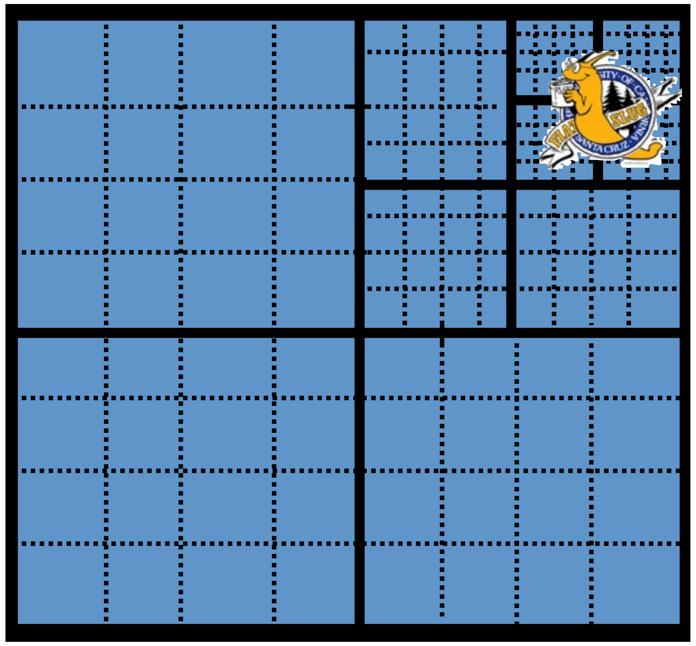

oct-tree-based block AMR

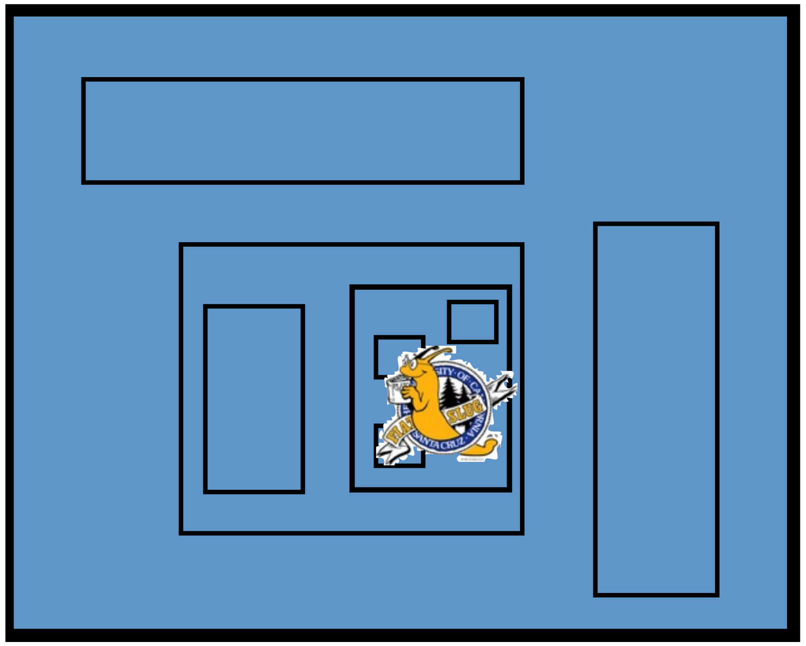

patch-based AMR

FLASH Grid Structures

- Two ways to setup
  - AMR using PARAMESH (MacNiece et al, 2000)
  - block structured, oct-tree based AMR
    - alternative CHOMBO (Colella et al.) library is under development
  - UG without AMR overhead

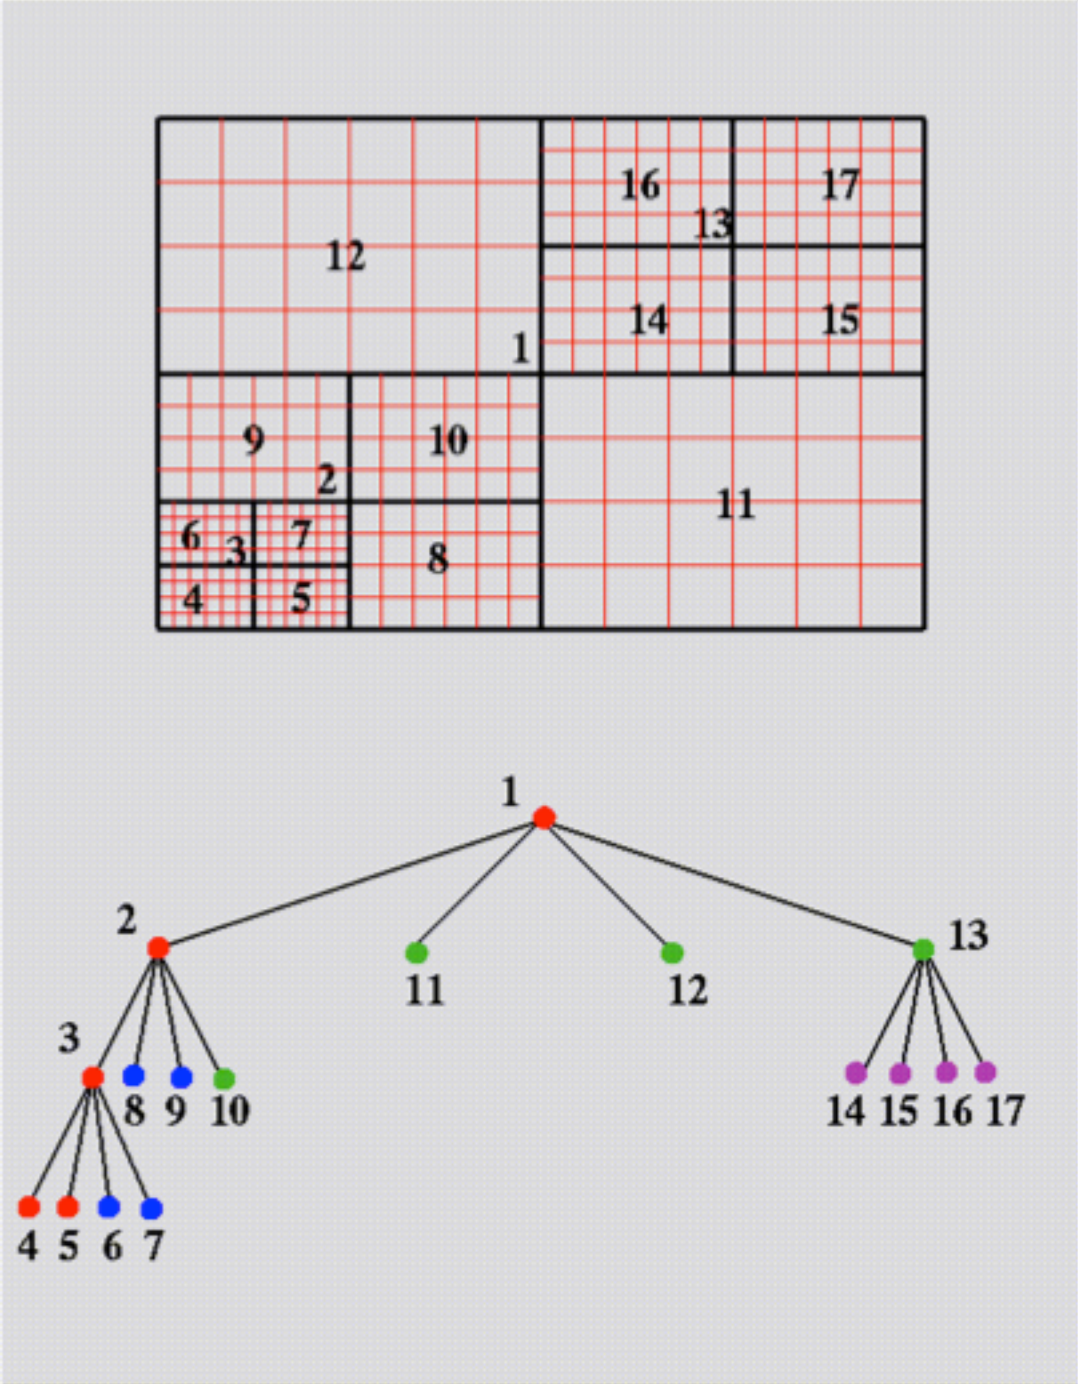

FLASH AMR Grid Structures

- Oct-tree structure where each branch of tree produces 2^d leaf blocks
- FLASH evolves only on the leaf blocks
- Ref ratio is a factor of 2
- Mesh can be refined or de-refined adaptively (customizable)
- Morton space-filling curve (Warren and Salmon, 1993) for load balancing

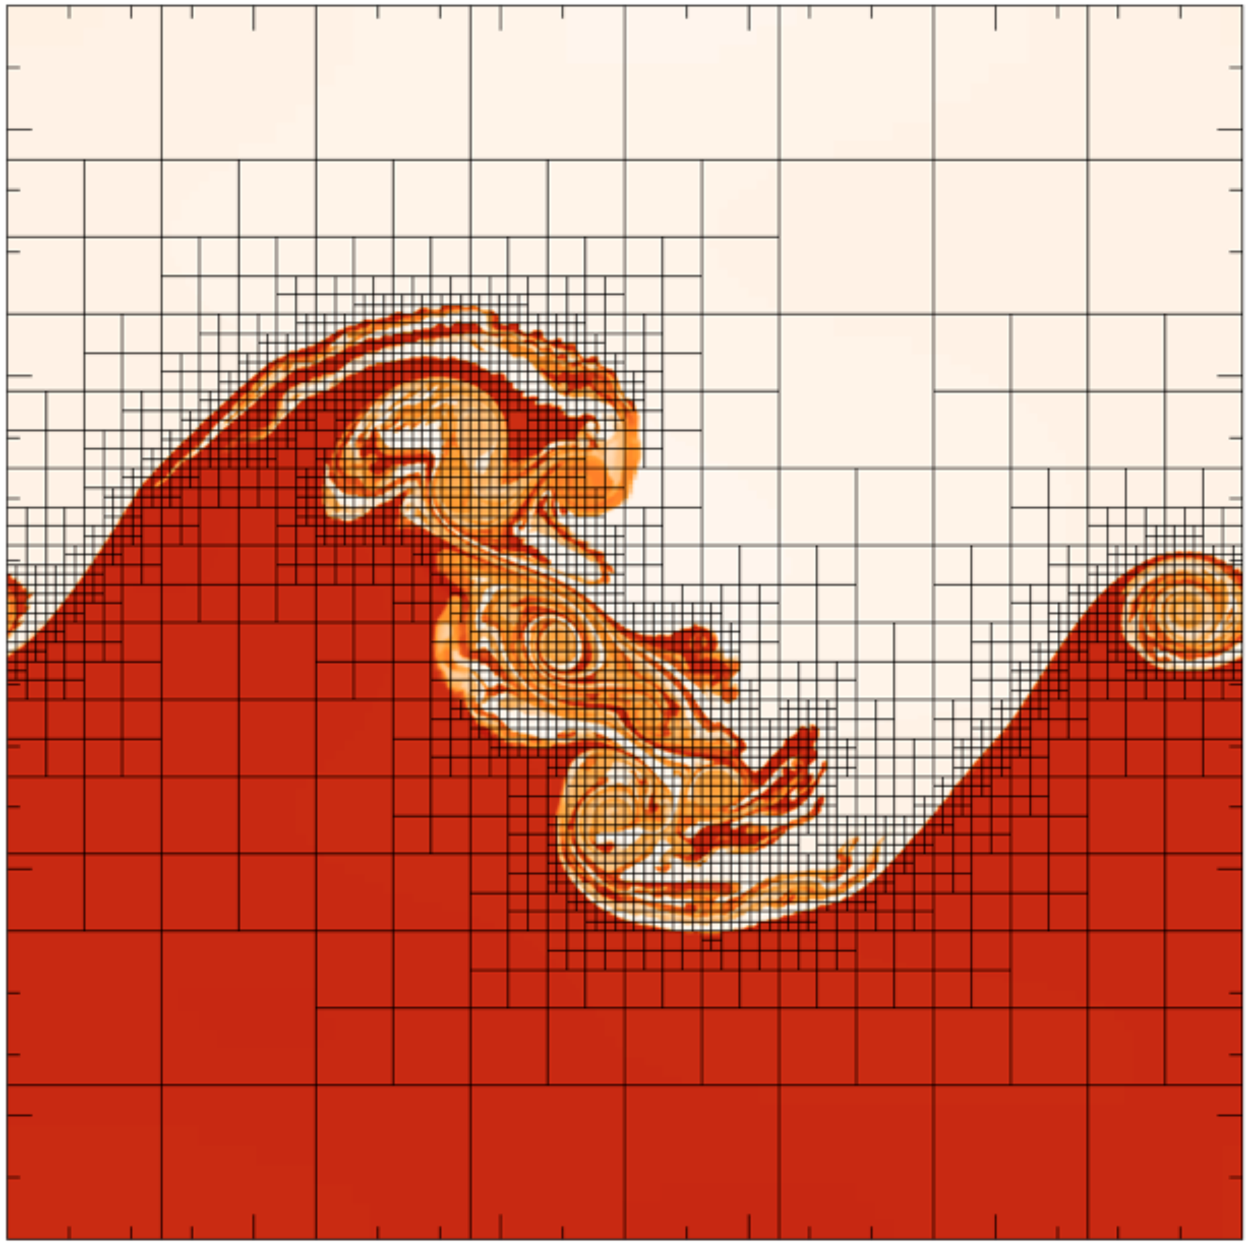

FLASH AMR Grid Structures

- Oct-tree structure where each branch of tree produces 2^d leaf blocks
- FLASH evolves only on the leaf blocks
- Ref ratio is a factor of 2
- Mesh can be refined or de-refined adaptively (customizable)
- Morton space-filling curve (Warren and Salmon, 1993) for load balancing

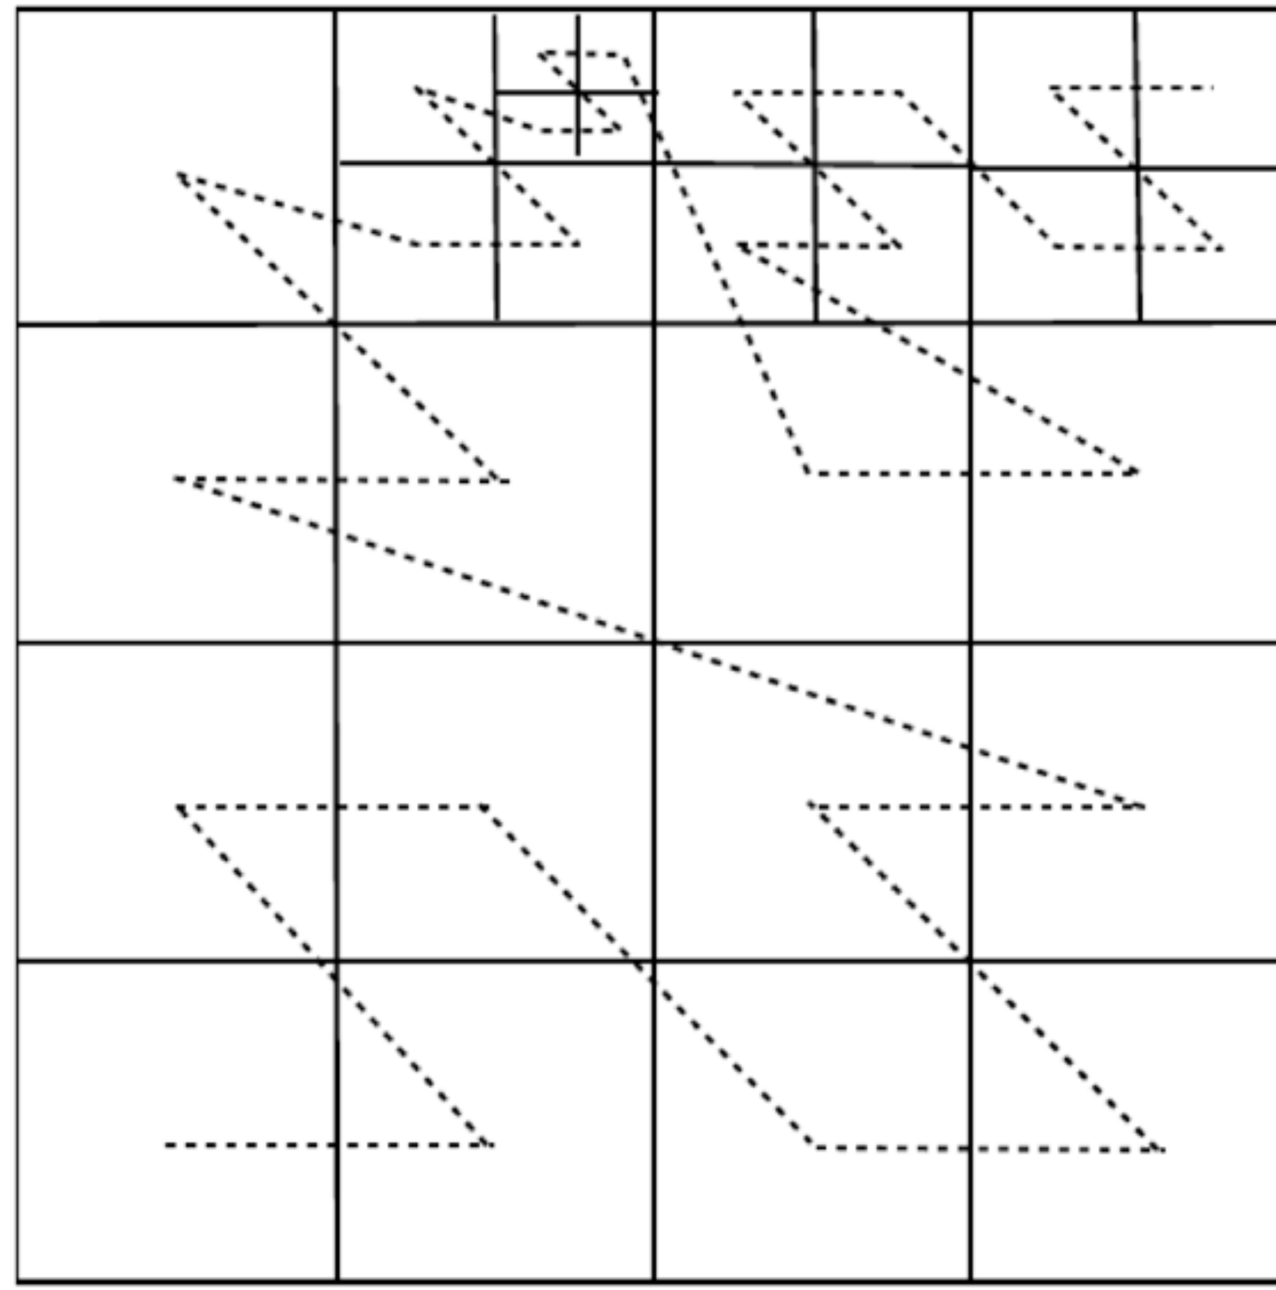

Various Tested Supports

- FLASH has been tested on various machines, platforms, compilers, OS, etc.

computer alias names

prototype makefiles:
linux, Mac OS X, …

```
Aliases
Prototypes/
SEAS10927.gwu.edu/
alc.llnl.gov/
animal5/
archimedes.uchicago.edu/
bassi.nersc.gov/
bgl.llnl.gov/
bgl.mcs.anl.gov/
bgl.sdsc.edu/
bonsai.cfa.harvard.edu/
brassica.asci.uchicago.edu/
```

…

```
hyades.ucsc.edu/
hydra.si.edu/
icc-9.0_fornax.uchicago.edu/
ignition/
```

FLASH Directory Structures

- Most of the code in Fortran 90
- highly modular with some set of coding rules
  - e.g., API, name conventions, de-centralized database, etc.
- configuration via setup python script

```
Driver/
Grid/
IO/
Multispecies/
Particles/
PhysicalConstants/
RuntimeParameters/
Simulation/
flashUtilities/
monitors/
physics/
```

```
Cosmology/
Diffuse/
Eos/
Flame/
Gravity/
Hydro/
RadTrans/
materialProperties/
sourceTerms/
```

```
Burn/
Cool/
Deleptonize/
EnergyDeposition/
Heat/
Heatexchange/
Ionize/
Polytrope/
PrimordialChemistry/
Stir/
```

FLASH Directory Structures

```
Driver/
Grid/
IO/
Multispecies/
Particles/
PhysicalConstants/
RuntimeParameters/
Simulation/
flashUtilities/
monitors/
physics/
```

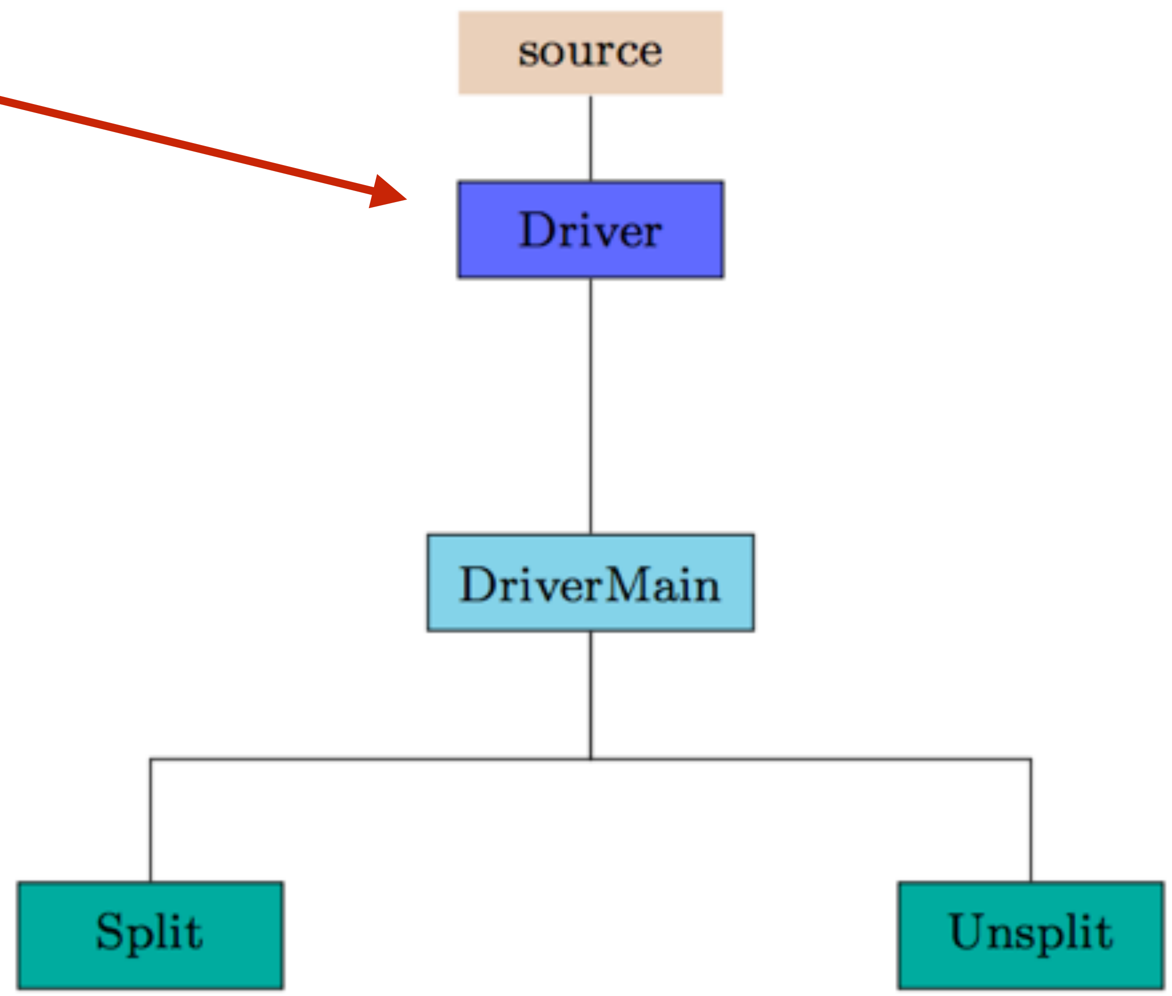

FLASH Directory Structures

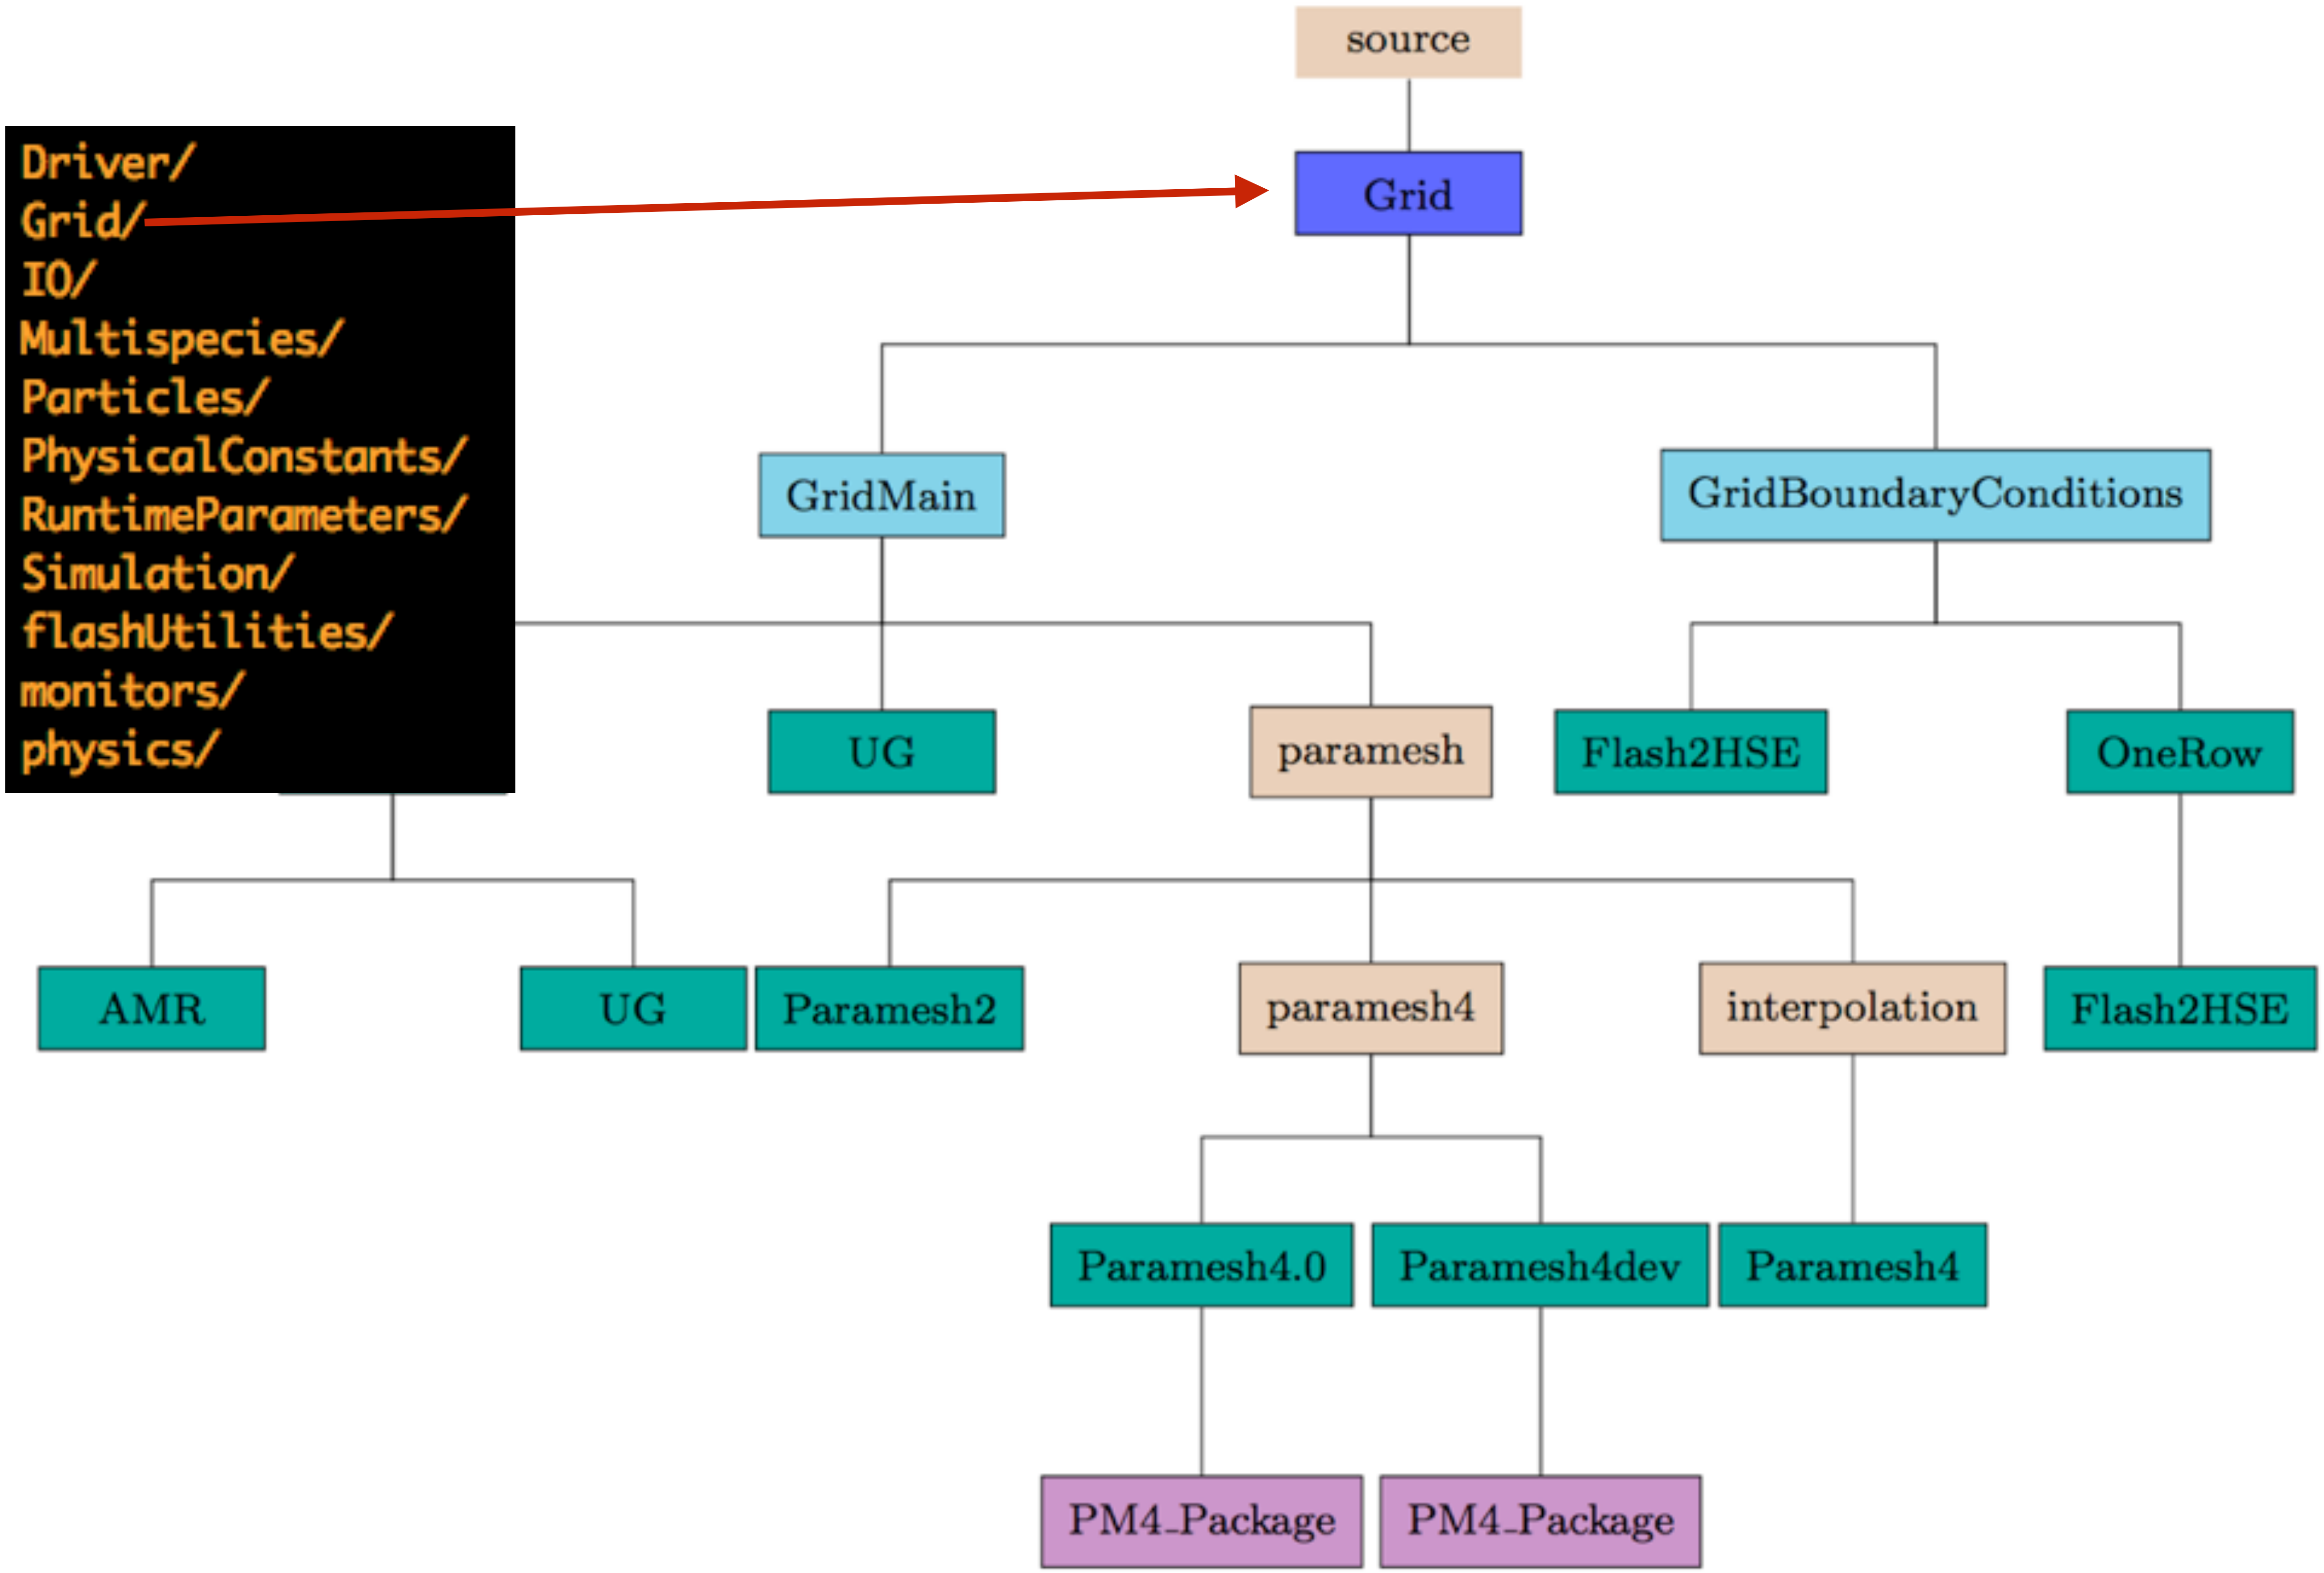

FLASH Directory Structures

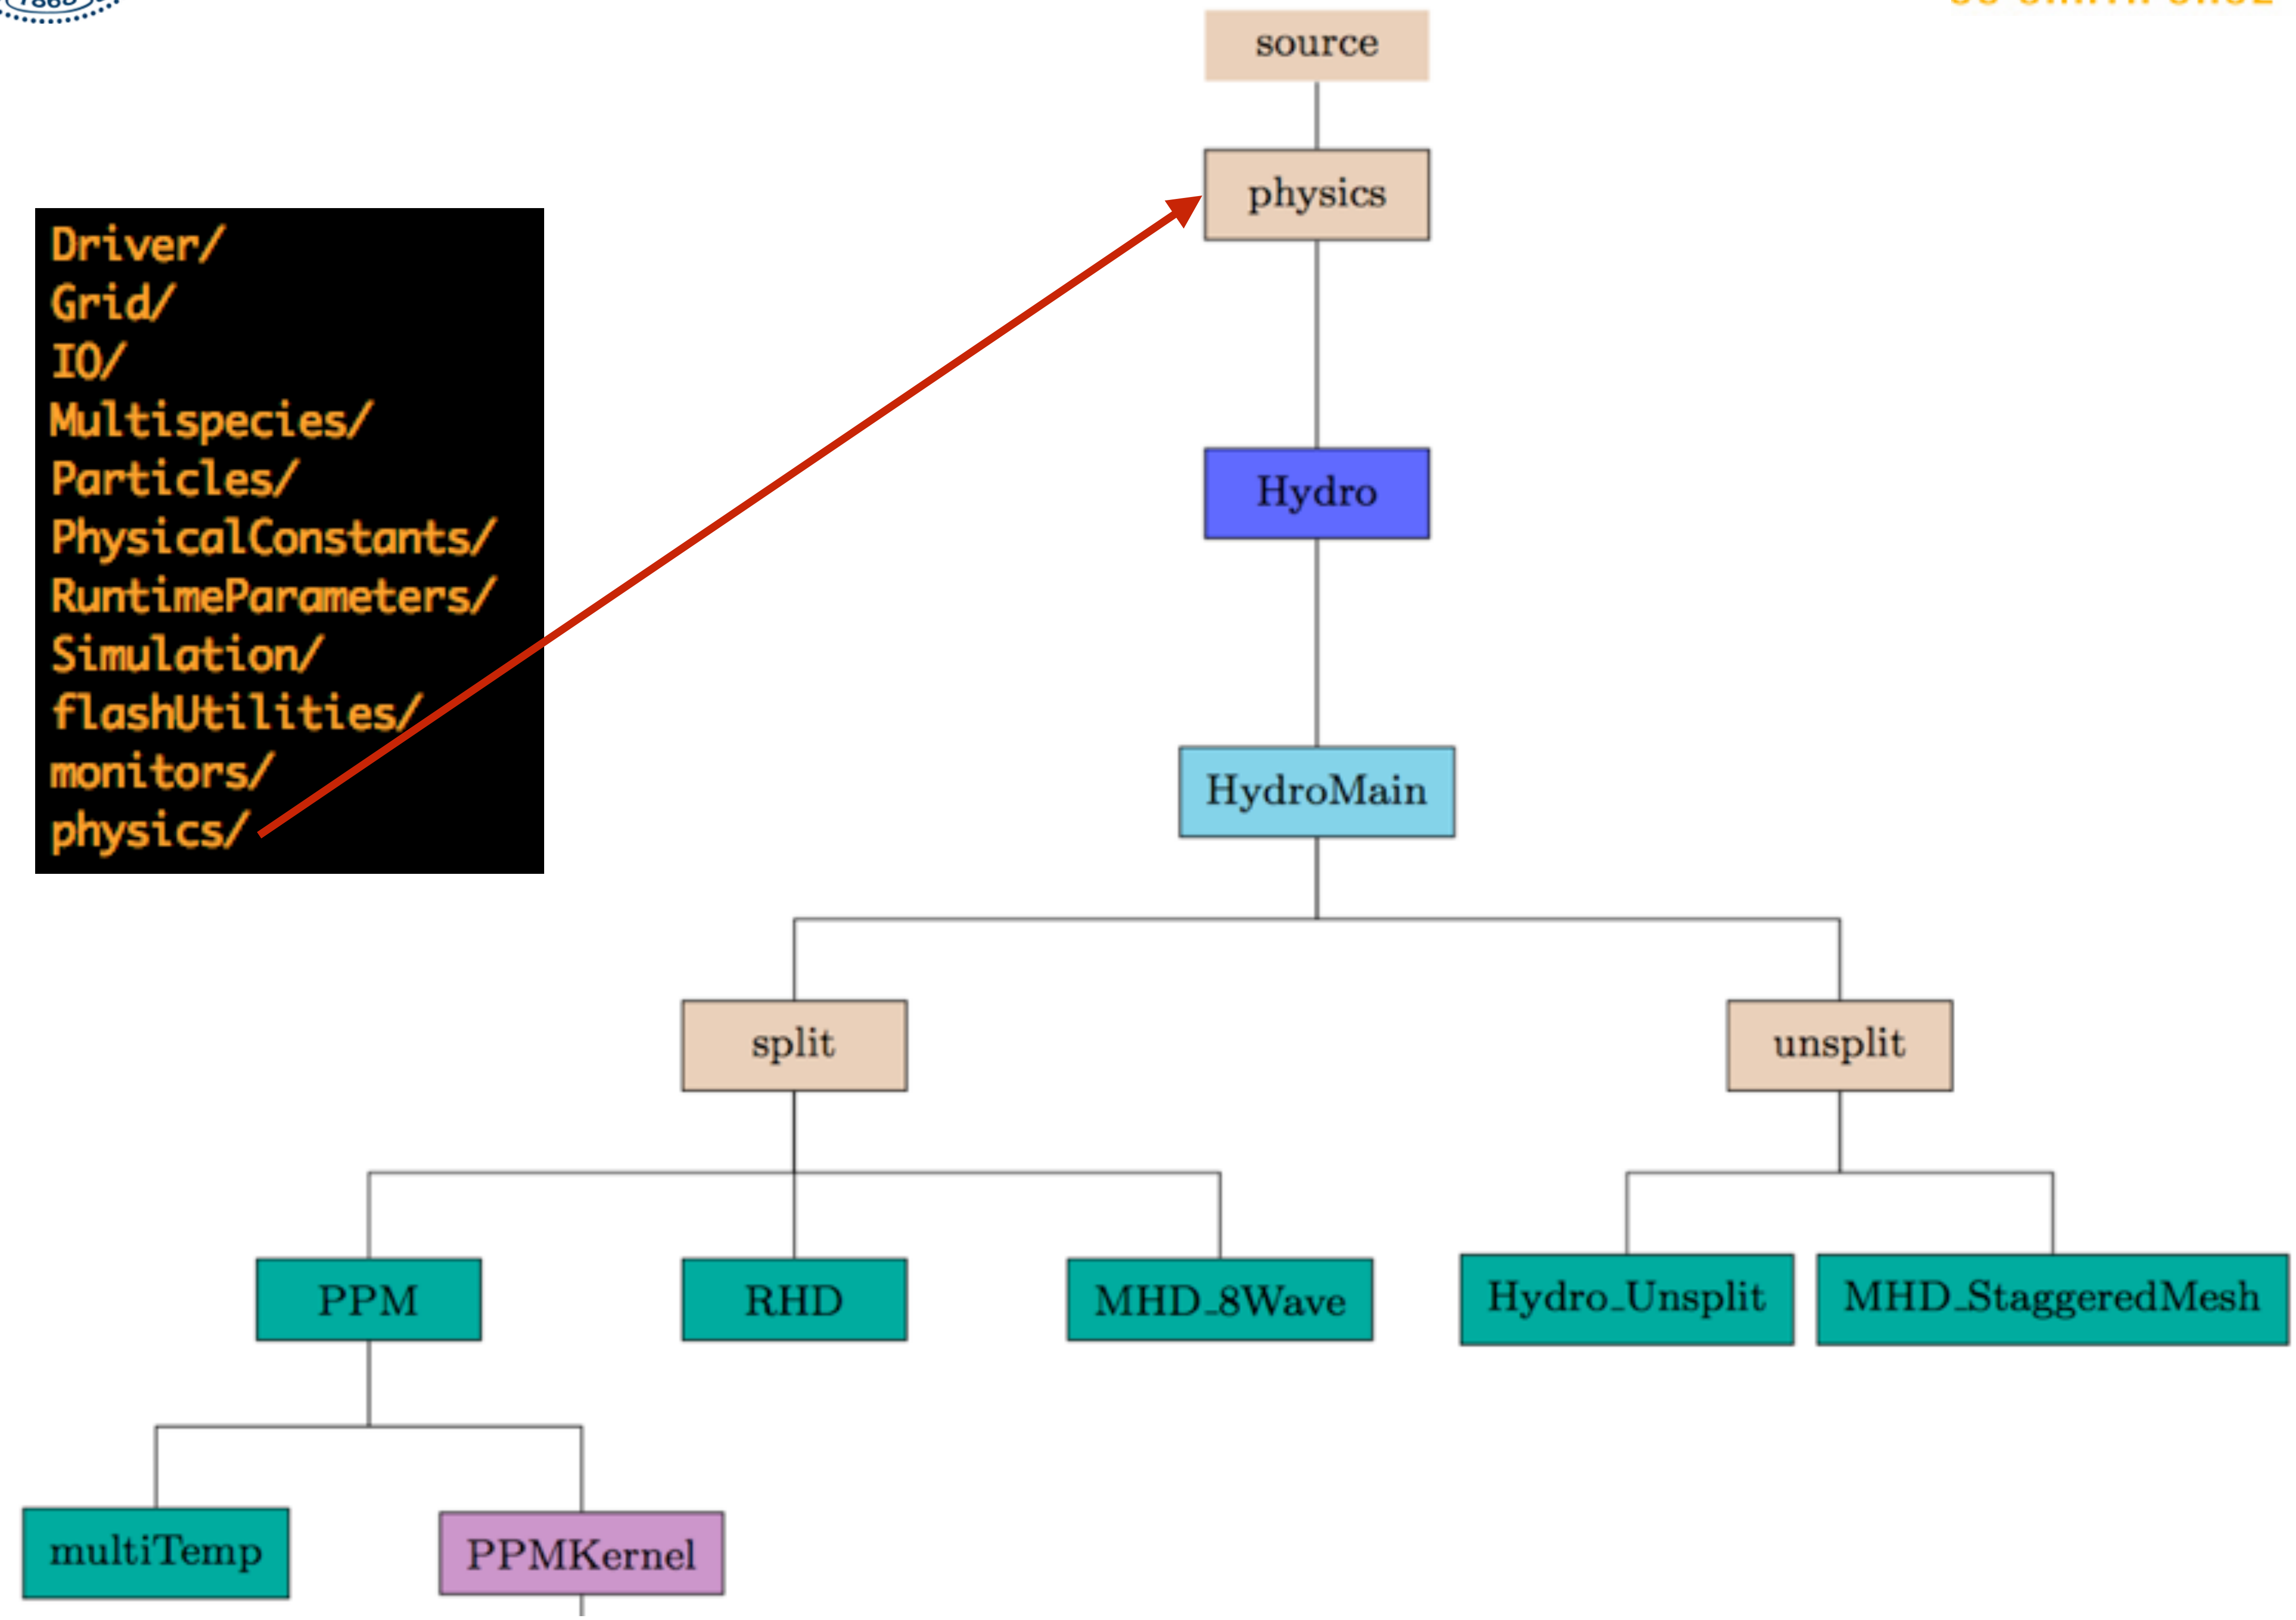

Libraries/Softwares

- Additional libraries/softwares required for FLASH:
  - fortran, C compilers with openMP
  - MPI
  - HYPRE for implicit diffusion
  - HDF5

FLASH Parallelizations

- Two types of parallelizations:
  - inter-node parallelism: domain decomposition with MPI (distributed memory)
    - Paramesh or Chombo
  - intra-node parallelism with OpenMP (shared memory)
    - thread block list
    - thread within block
- More parallelization…
  - Parallel I/O with HDF5

FLASH Strong vs. Weak Scalings

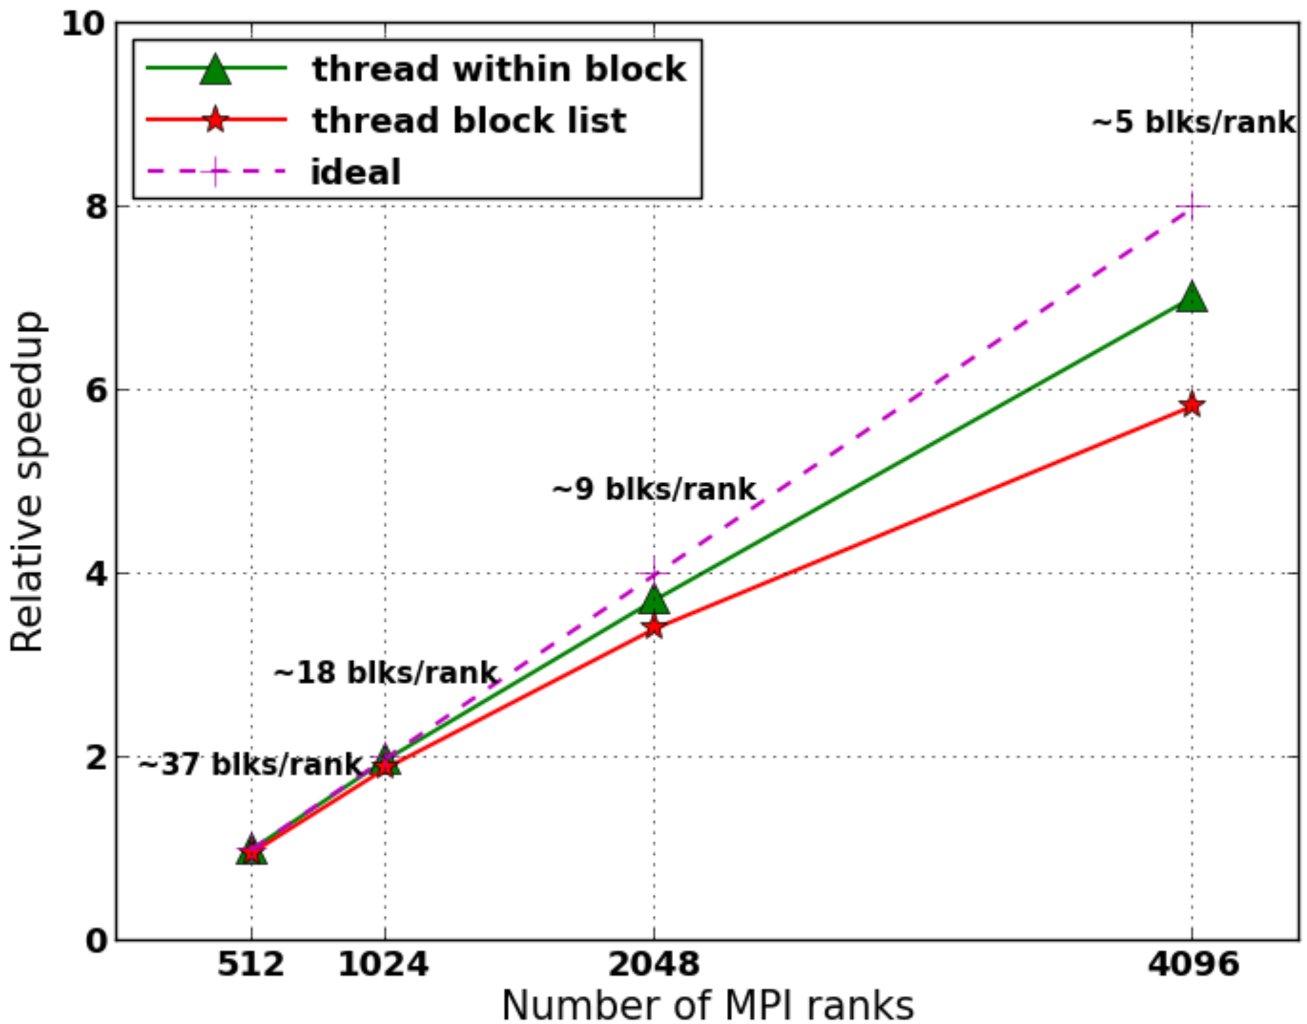

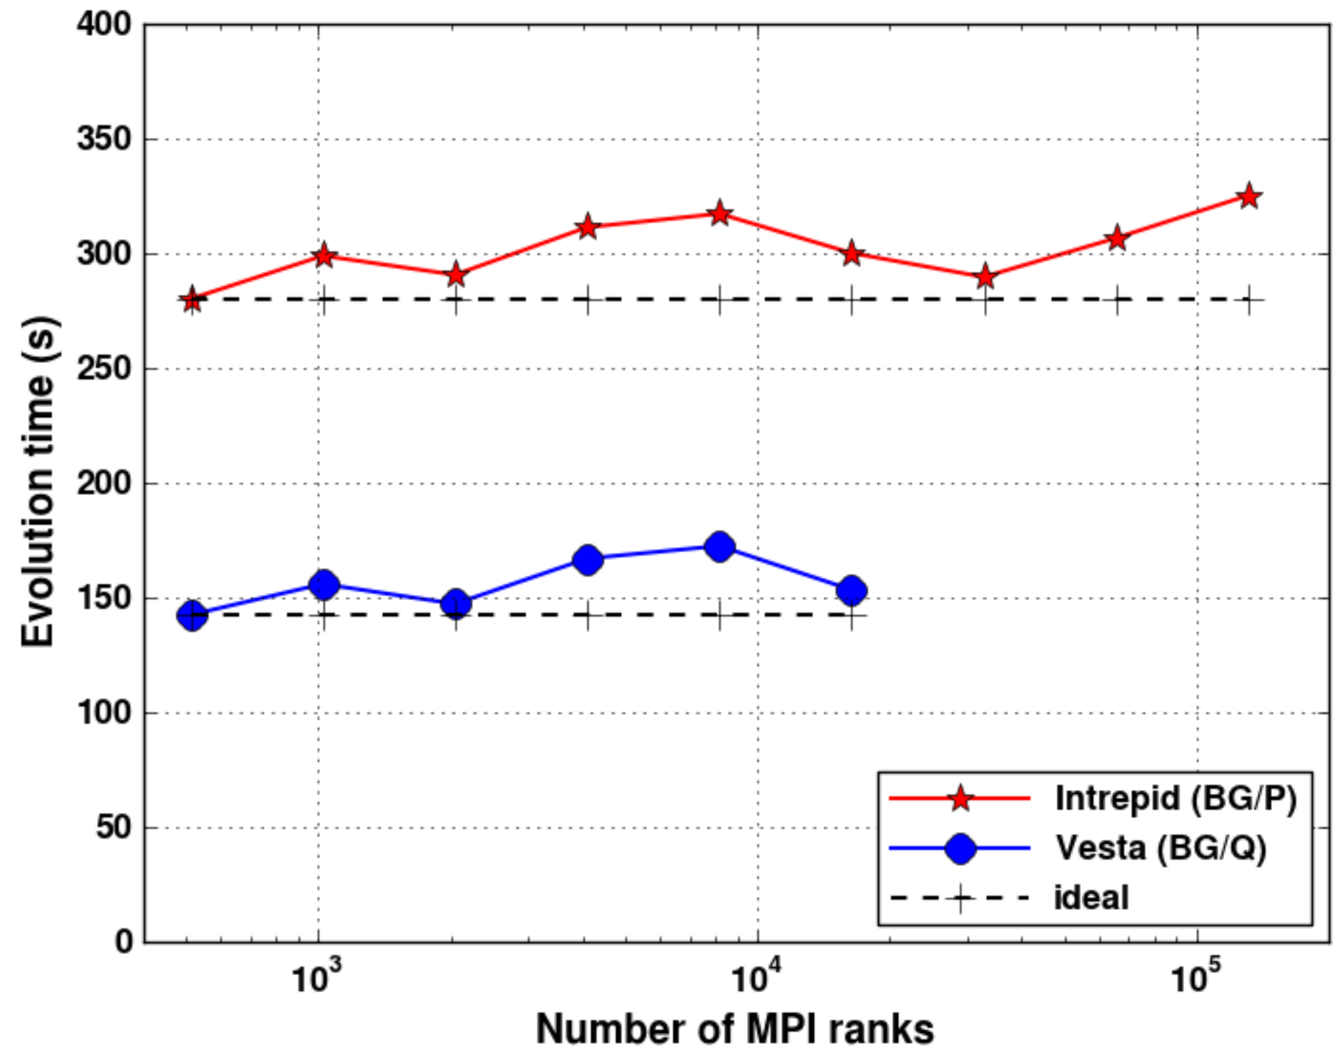

BG/P to BG/Q transition

- Intrepid BG/P
  - 4 cores/node, 2GB/node, 40,960 nodes
  - FLASH has been run on Intrepid for the last several years
    - scales to the whole machine
    - MPI-only is sufficient

- Mira BG/Q
  - 4 hw threads/core, 16 cores/node, 16GB/node, 152 nodes
    - MPI-only approach not suitable for BG/Q
    - OpenMP directives have been recently added to FLASH to take advantages of the additional intra-node parallelism

BG/Q

16 GB/node

16 cores/node

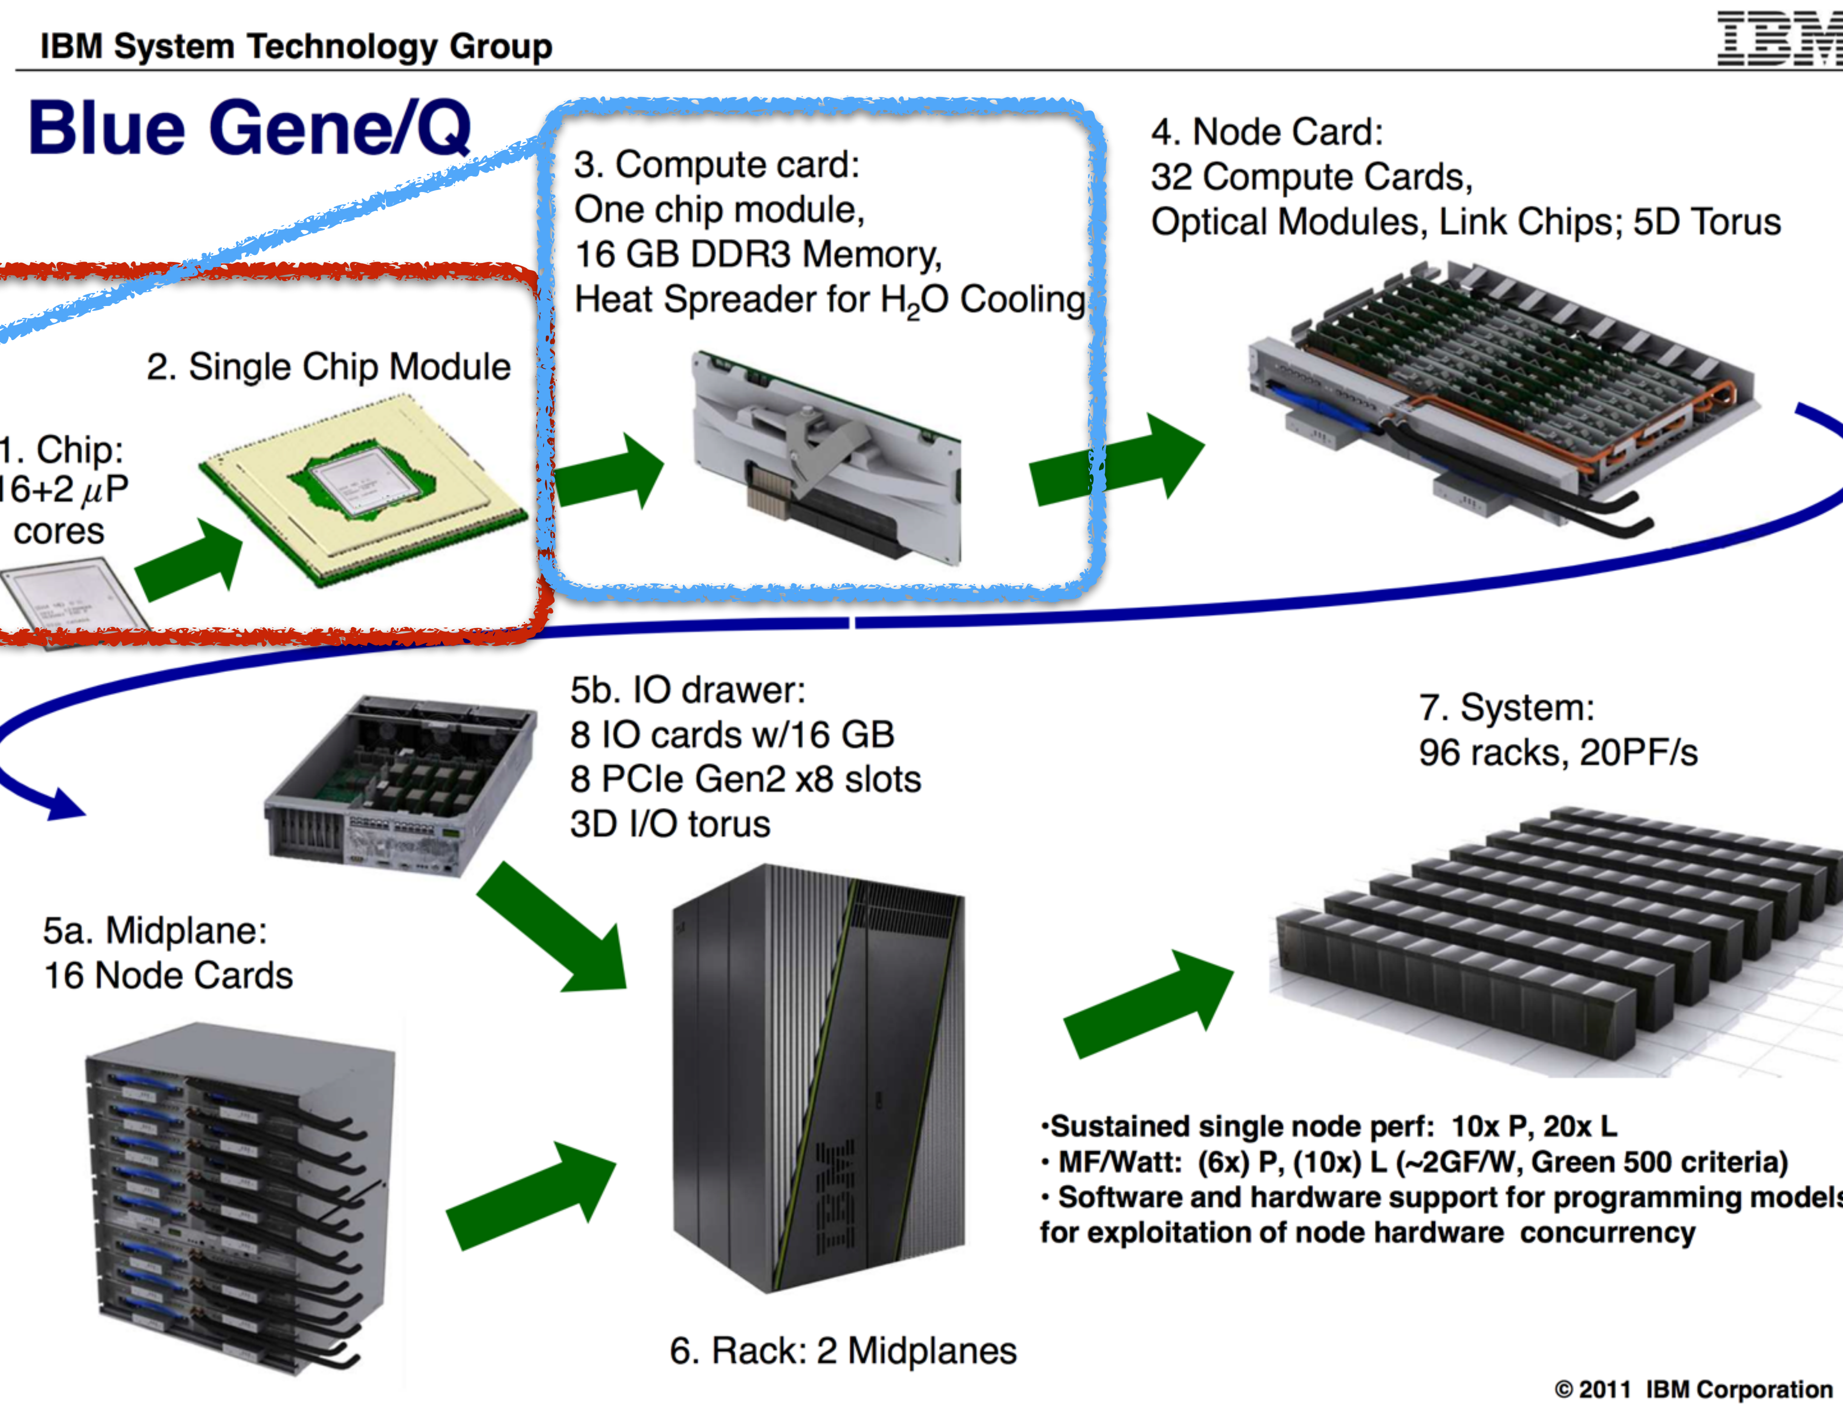

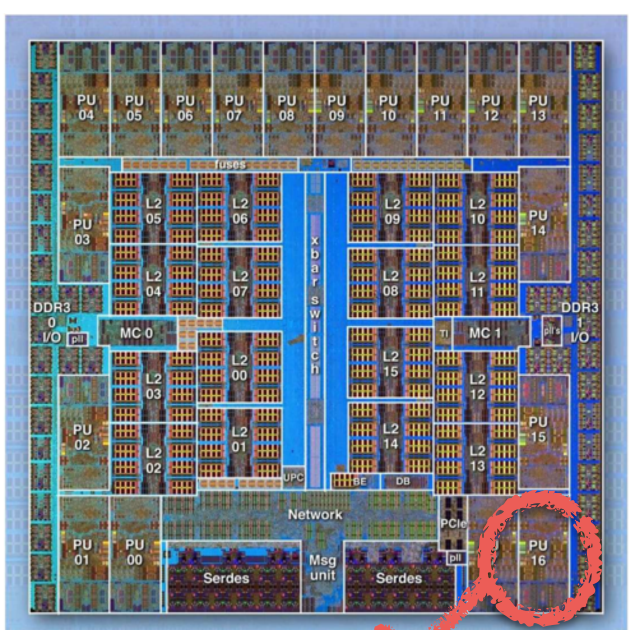

4 threads/core

FLASH Threading

- Blocks 12~17 being assigned to a single MPI rank
  - 6 total blocks
    - 5 leaf & 1 parent (#13)
- Mesh is divided into blocks of fixed size
- Oct-tree hierarchy
- Blocks are assigned to MPI ranks

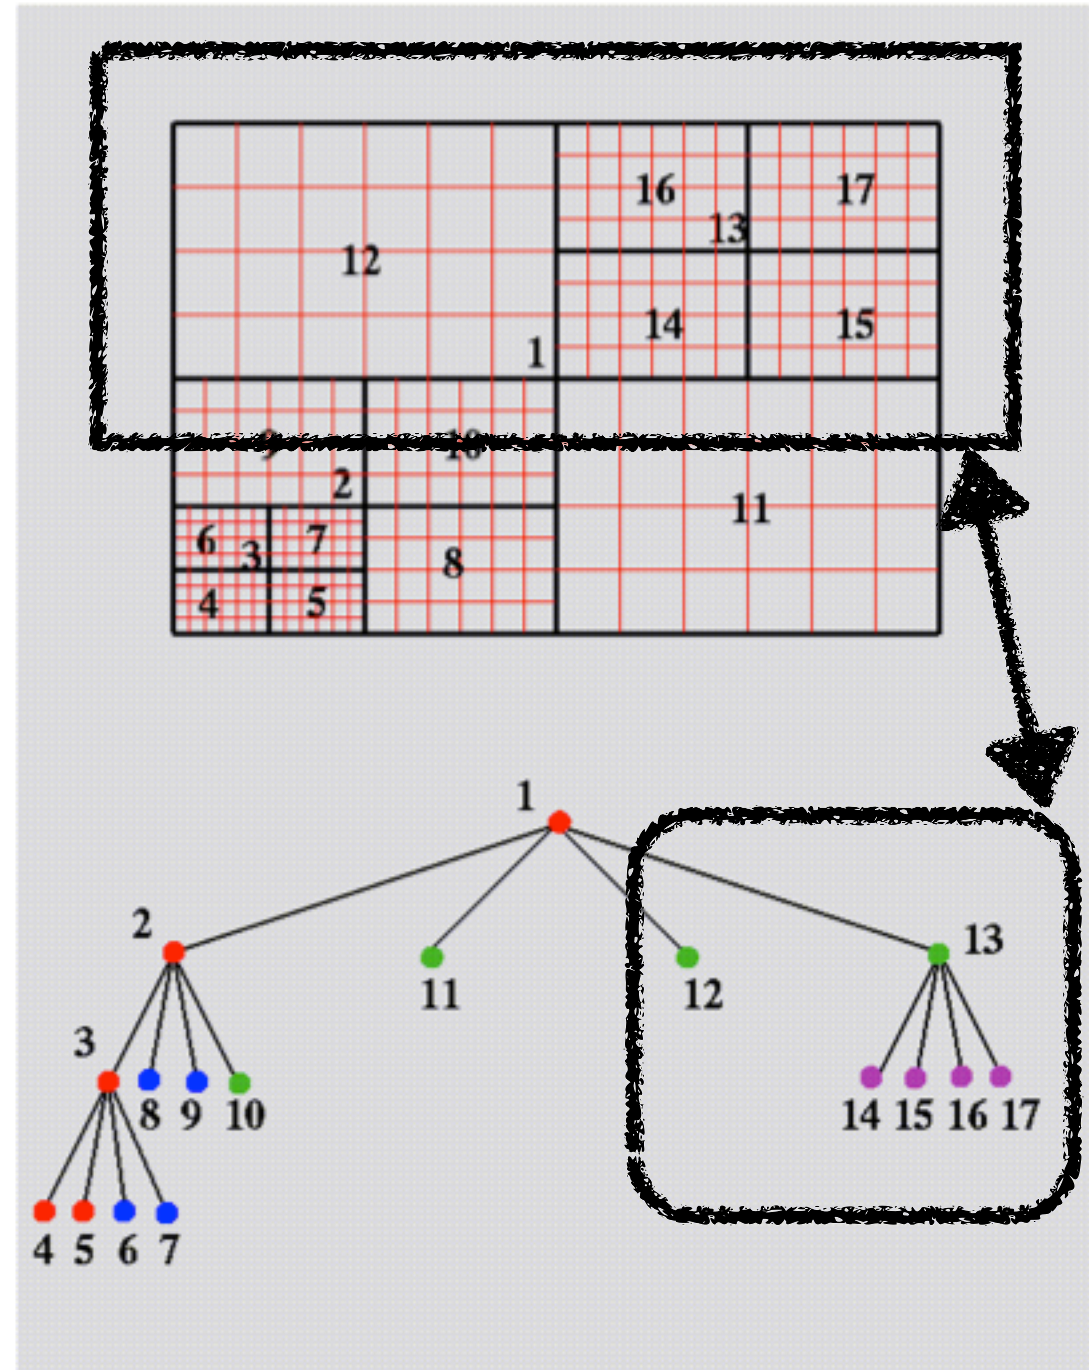

FLASH Threading

- FLASH solvers update the solution only on local leaf blocks
  - FLASH uses multiple threads to speed up the solution update

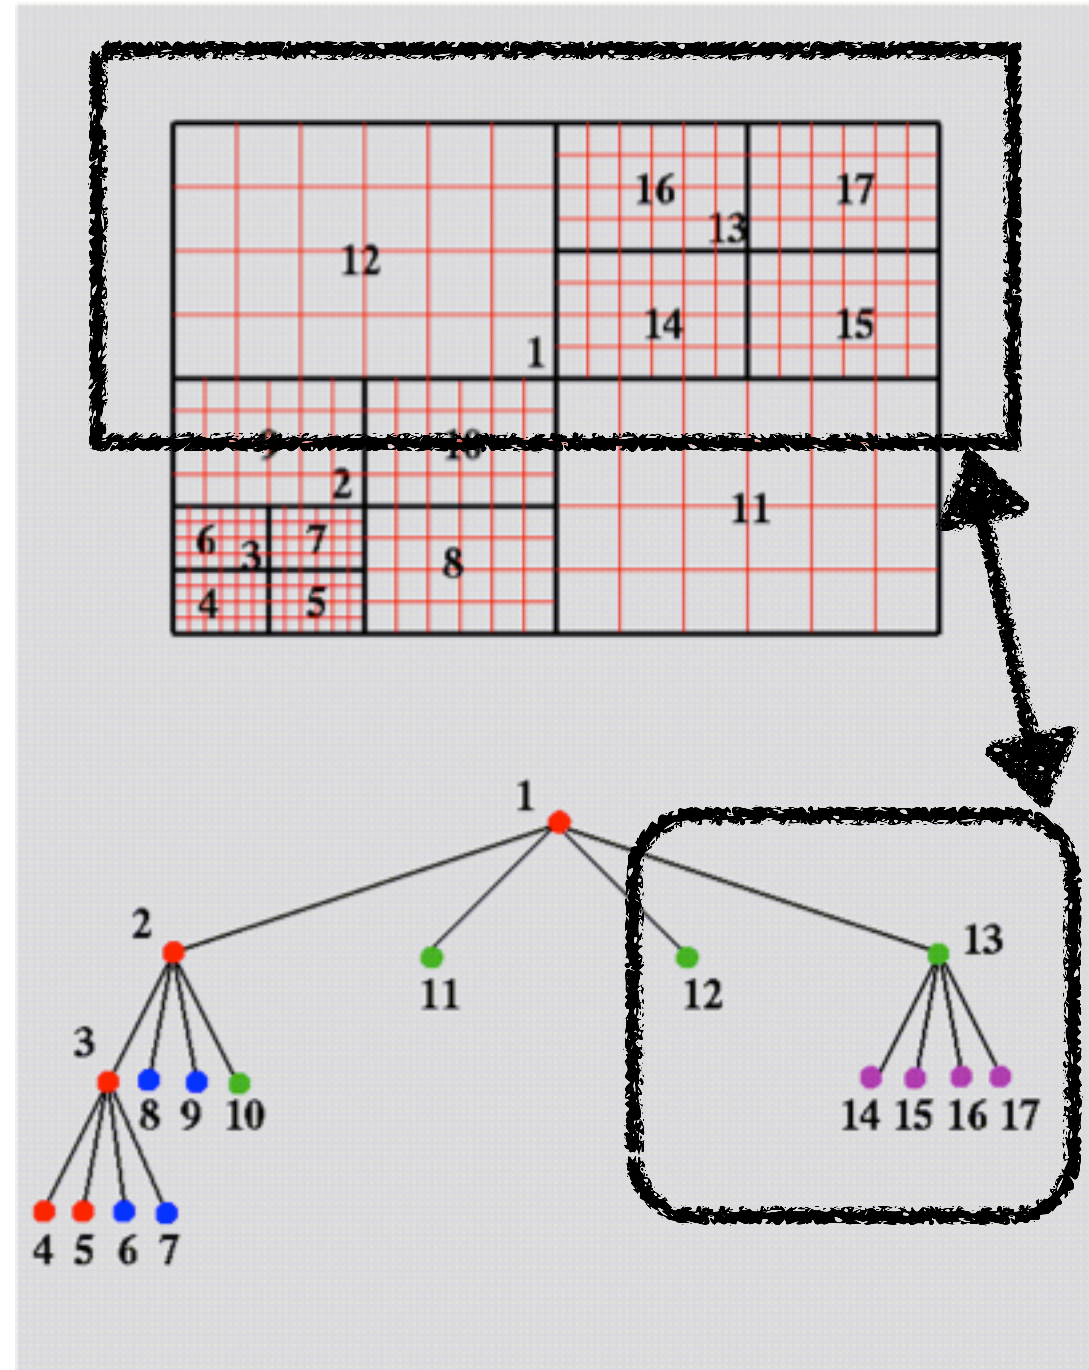

Threading Strategy 1

- Assign different blocks to different threads
- Assuming 2 threads per MPI rank
  - thread 0 (blue) updates 3 full blocks — 72 cells
  - thread 1 (yellow) updates 2 full blocks — 48 cells
- ‘thread block list’ — coarse grained threading

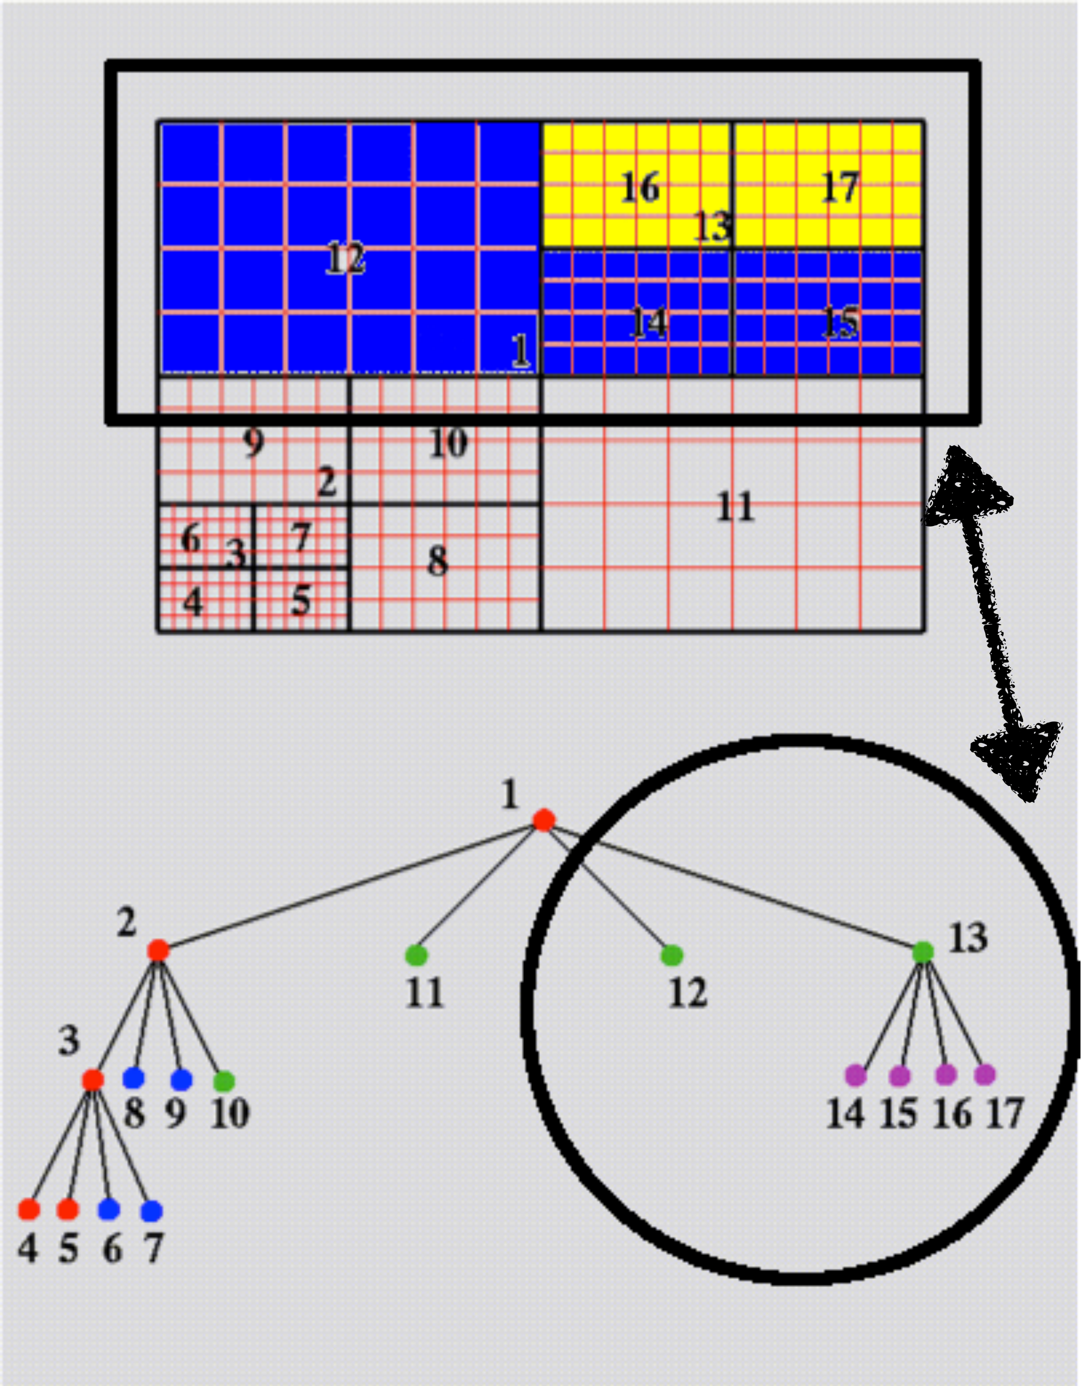

Threading Strategy 2

- Assign different cells from the same block to different threads
- Assuming 2 threads per MPI rank
  - thread 0 (blue) updates 5 partial blocks — 60 cells
  - thread 1 (yellow) updates 5 partial blocks — 60 cells
- 'thread within block' - fine grained threading

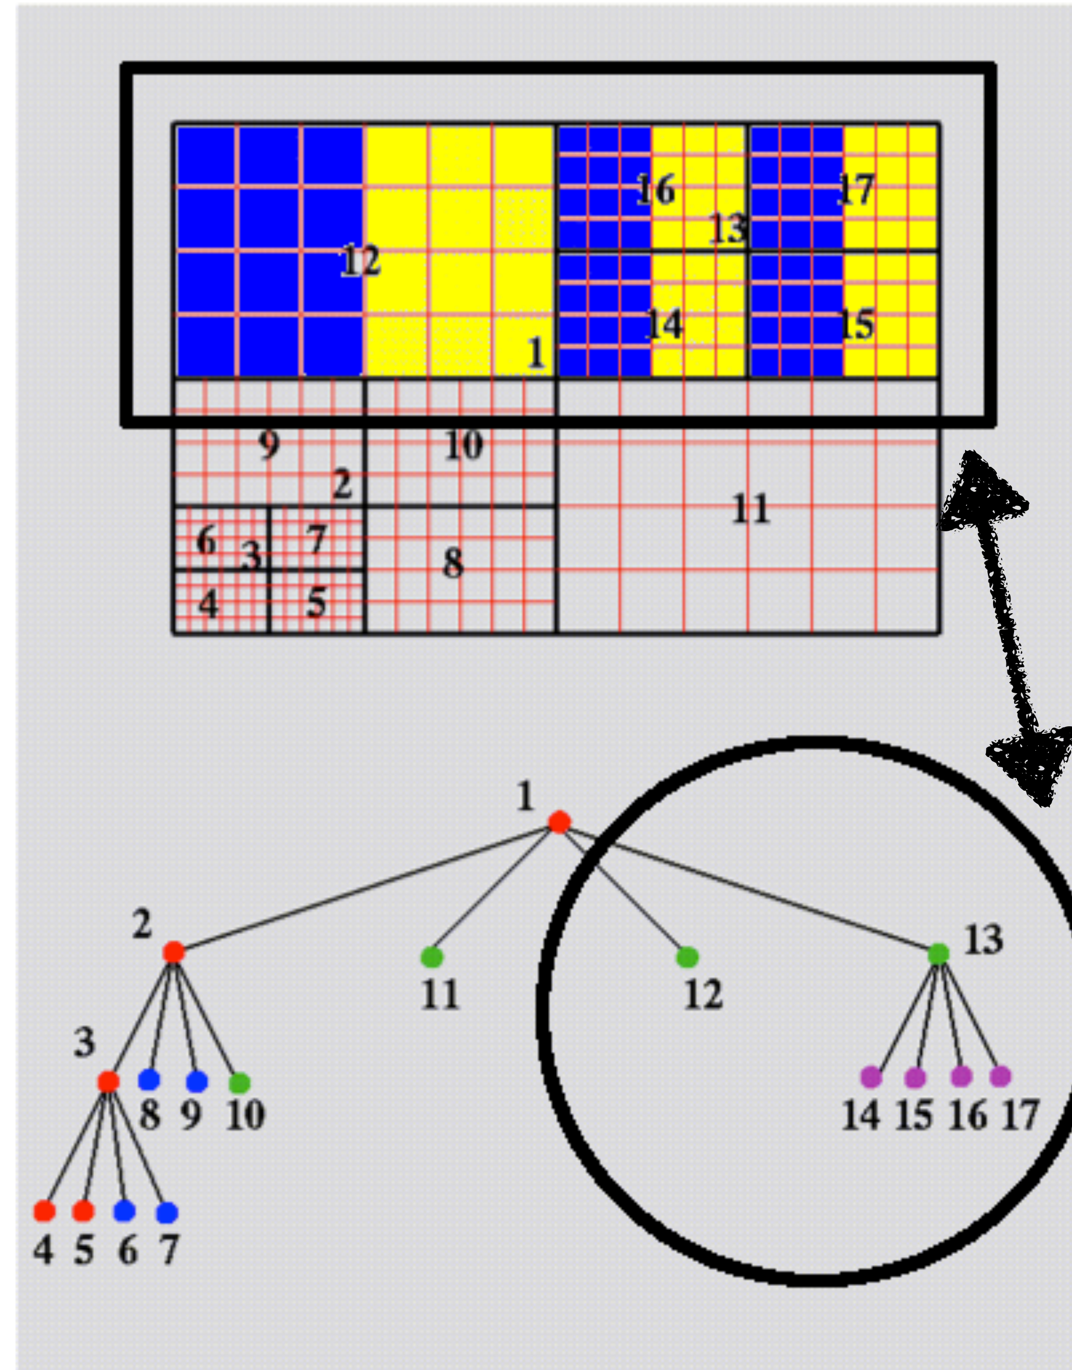

Strong Scaling of RT Flame

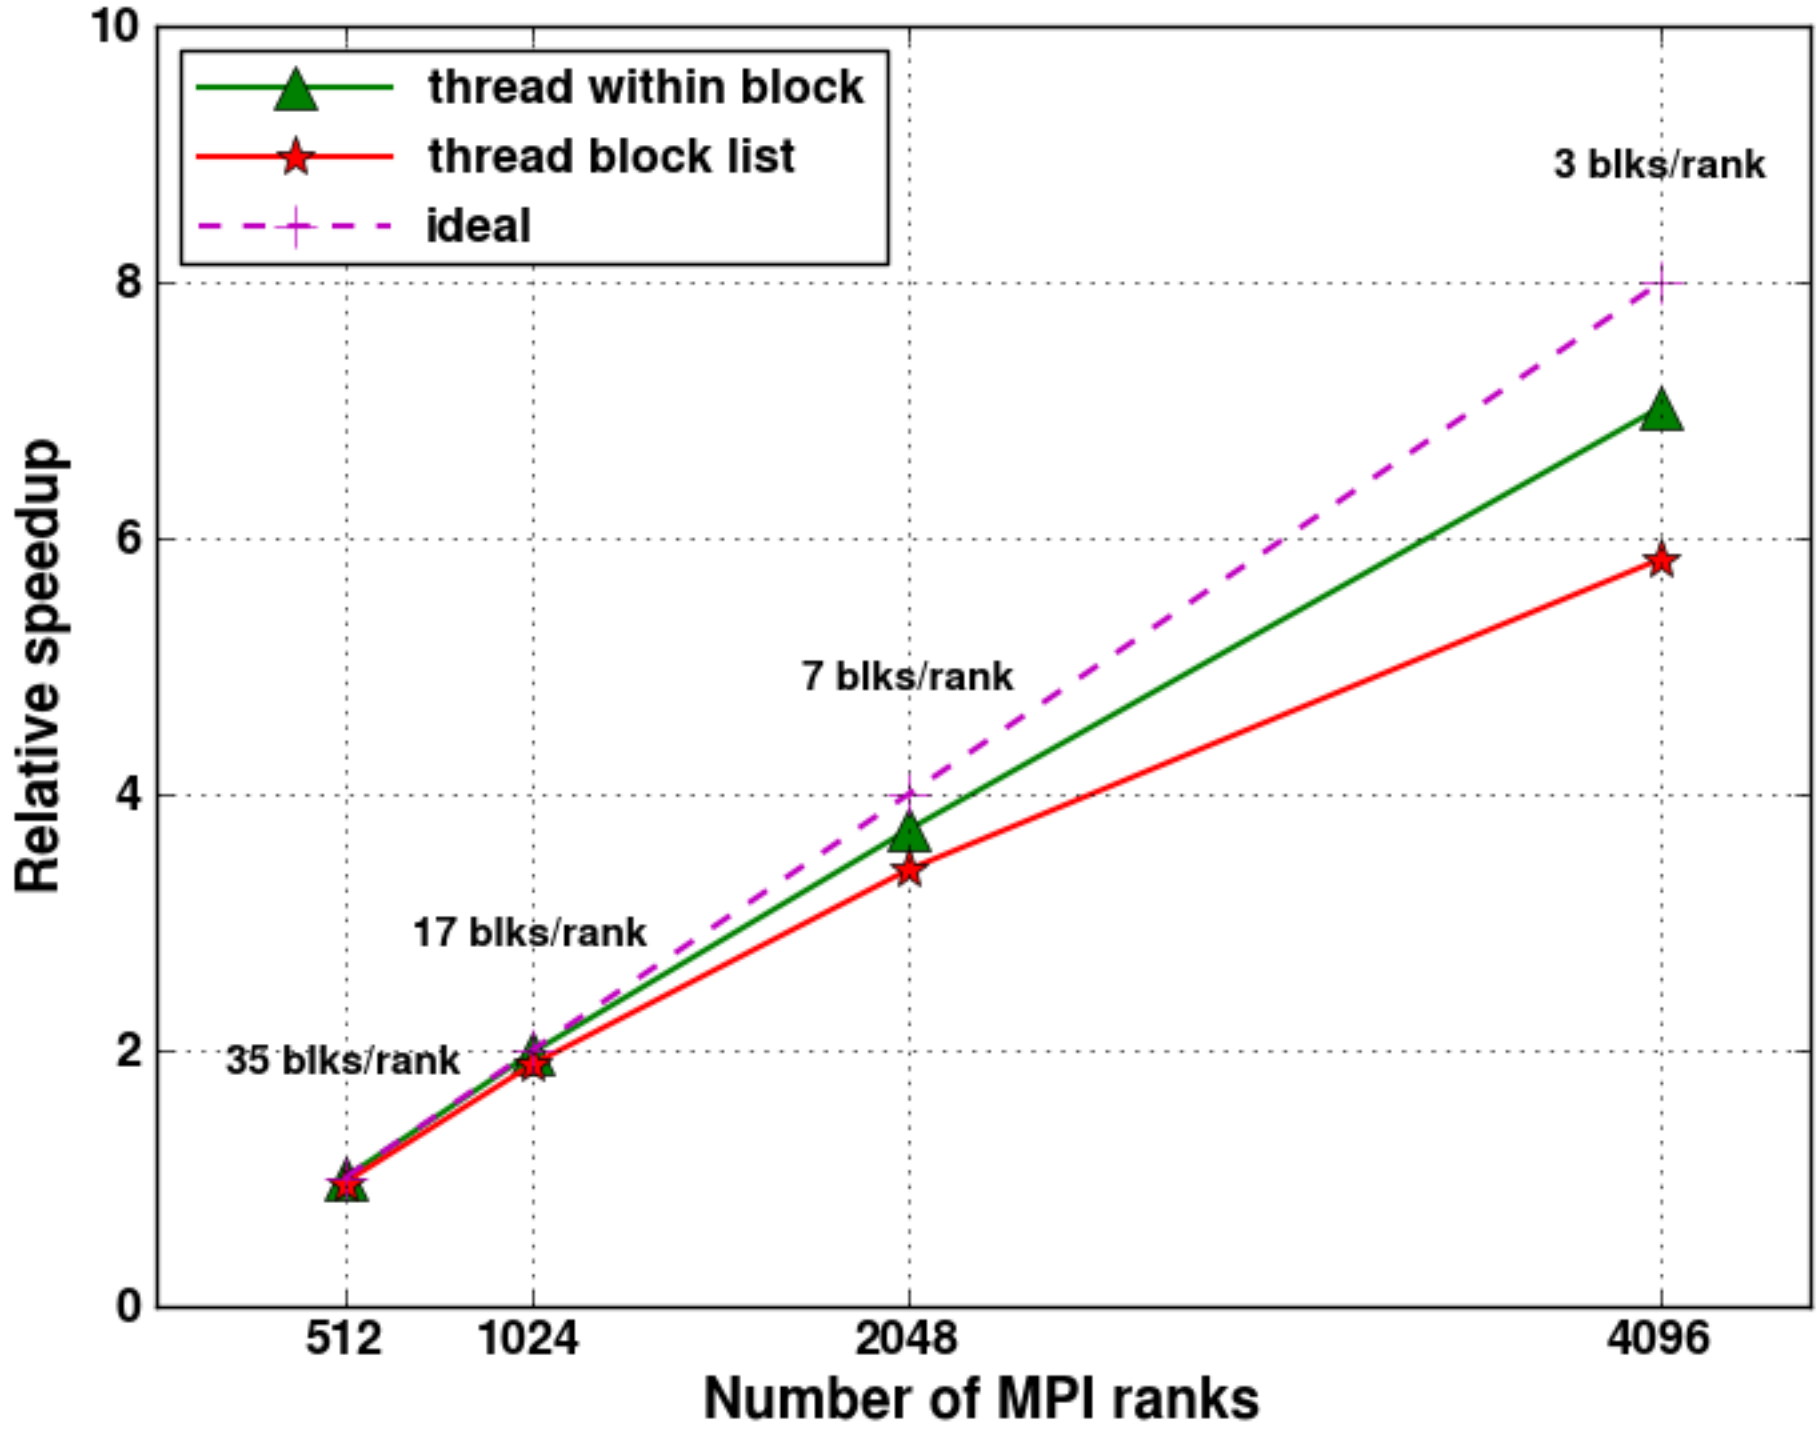

Strong Scaling Results

- Good strong scaling for both multithreading strategies
  - Only 3 blocks of 16^3 cells on some MPI ranks for the 4096 MPI rank calculation (256 nodes)
  - Coarse-grained threading: list of blocks >> # of threads to avoid load imbalance

- Finer-grained threading performs slightly better
  - Better load balancing within an MPI rank
  - Performance advantage increases as work becomes more finely distributed

Further Optimizations

- For instance, reordering arrays in kernels reduced to time to solution from 184 sec to 156 sec

Before: 4.79 sec

```
689    479    call  hy_uhd_dataReconstOnestep &
...
711             leig(1:NDIM,i,j,k,1:HY_WAVENUM,1:HY_VARINUM),&
712             reig(1:NDIM,i,j,k,1:HY_VARINUM,1:HY_WAVENUM))
```

After: 0.37 sec

```
689     37    call  hy_uhd_dataReconstOnestep &
...
711             leig(1,1,1,i,j,k),&
712             reig(1,1,1,i,j,k))
```

Further Optimizations

- FLASH re-gridding optimization in Paramesh

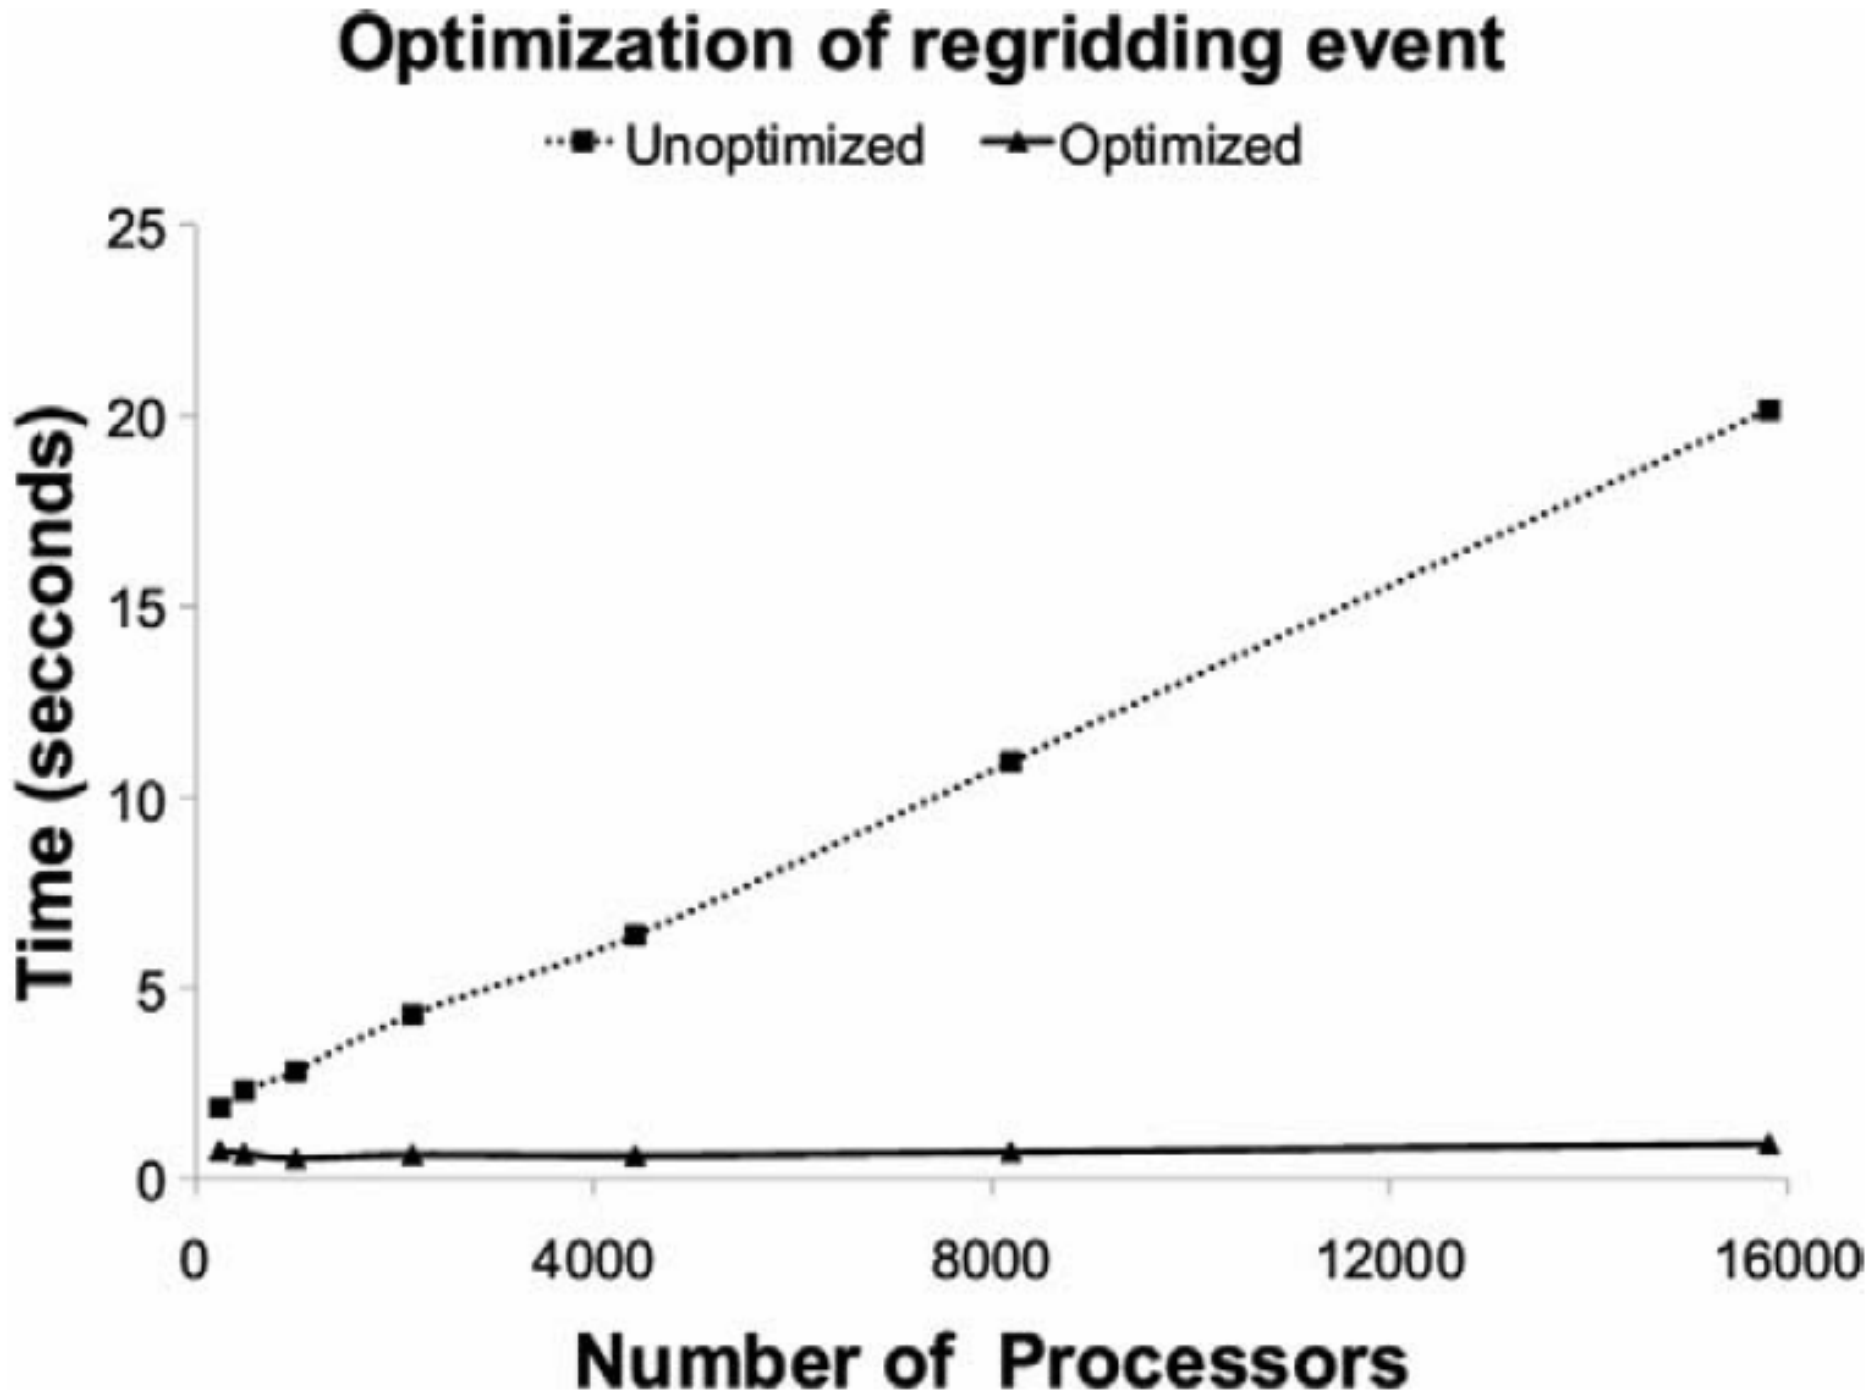

Conclusion

- It is challenging to modify large software like FLASH to contain AMR and multiple physics modules for the next generation of many-core architecture
- FLASH's successful approach is to add OpenMP directives to deliver a large, highly-capable piece of software to run efficiently on the BG/Q platforms
  - coarse grained and fine grained threading strategies have been explored and tested
  - fine grained threading performs better
- Extra optimizations can speedup the code performance