The Coming Revolution in Computational Astrophysics

Saul Teukolsky
Caltech & Cornell University

University of Toronto, October 19, 2017

Astrophysics Simulations

- Growing importance
- Many examples:
    - Cosmology
    - Globular clusters
    - Tidal disruption
    - Accretion disks
    - Planet formation
    - …

"It says it's sick of doing things like inventories and payrolls, and it wants to make some break-throughs in astrophysics."

GW150914: A Famous Example

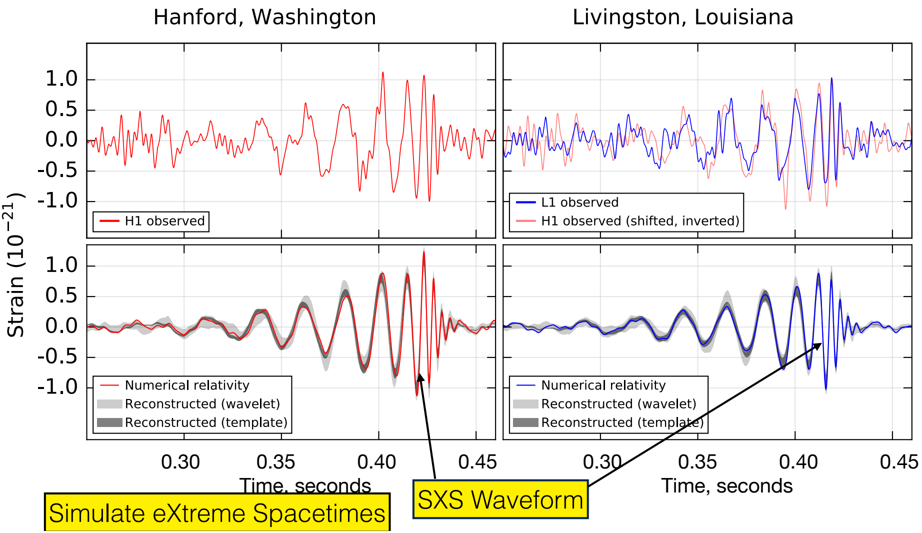

ANNALS OF PHYSICS: **29**, 304–331 (1964)

The Two-Body Problem in Geometrodynamics

SUSAN G. HAHN

International Business Machines Corporation, New York, New York

AND

RICHARD W. LINDQUIST

Adelphi University, Garden City, New York

The problem of two interacting masses is investigated within the framework of geometrodynamics. It is assumed that the space-time continuum is free of all real sources of mass or charge; particles are identified with multiply connected regions of empty space. Particular attention is focused on an asymptotically flat space containing a "handle" or "wormhole." When the two "mouths" of the wormhole are well separated, they seem to appear as two centers of gravitational attraction of equal mass. To simplify the problem, it is assumed that the metric is invariant under rotations about the axis of symmetry, and symmetric with respect to the time $t = 0$ of maximum separation

50 time steps
3 CPU hours (IBM 7090)
151×51 grid points
$t = 1.8M$

"In summary, the numerical solution of the Einstein field equations presents no insurmountable difficulties."

Focus of This Talk: PDEs

- Hydrodynamics
- MHD
- Gravity (Newton; Einstein)
- Radiation transport
- ...

Focus of This Talk: PDEs

- Hydrodynamics
- MHD
- Gravity (Newton; Einstein)
- Radiation transport
- …

Not N-body, Monte Carlo

The Dirty Secret

For the past 50 years, dominant algorithm essentially unchanged!

The Dirty Secret

For the past 50 years, dominant algorithm essentially unchanged!

Finite differencing (finite volume)

The First Stirrings . . .

- Solutions of Einstein's equations are smooth (away from singularities)
- Should use higher-order numerical methods

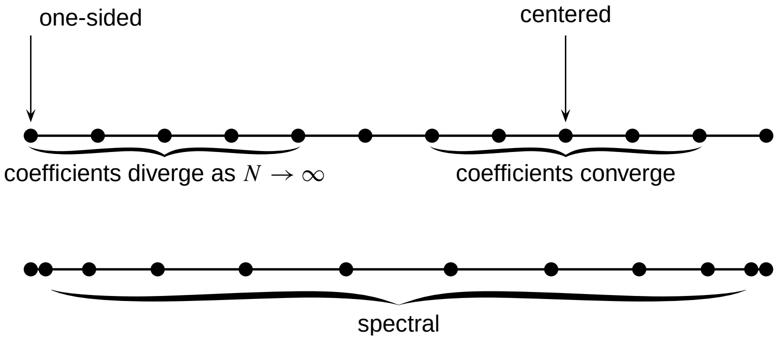

Why is SpEC So Good for BBHs?

- Approximate solution as sum of N basis functions

$$f(x,t) = \sum_{k=0}^{N-1} f_k(t)\phi_k(x)$$

- Spectral method:

$$f_k(t) = \int f(x,t)\phi_k(x)\,dx$$

- Pseudospectral method (Lanczos 1938):

$$f_k(t) = \sum_{n=0}^{N-1} w_n f(x_n,t)\phi_k(x_n)$$

- Uses N *collocation points* $\{x_n\} \to \{f(x_n,t)\}$ (momentum space vs. position space)
- Compute spatial derivatives analytically

Why is SpEC So Good for BBHs?

- Approximate solution as sum of N basis functions

$$f(x,t) = \sum_{k=0}^{N-1} f_k(t)\phi_k(x)$$

- Spectral method:

$$f_k(t) = \int f(x,t)\phi_k(x)\,dx$$

- Pseudospectral method (Lanczos 1938):

$$f_k(t) = \sum_{n=0}^{N-1} w_n f(x_n,t)\phi_k(x_n)$$

- Uses N *collocation points* $\{x_n\} \to \{f(x_n,t)\}$ (momentum space vs. position space)
- Compute spatial derivatives analytically
- Exponential convergence for smooth solutions

Why is SpEC So Good for BBHs?

- Approximate solution as sum of N basis functions

$$f(x,t) = \sum_{k=0}^{N-1} f_k(t)\phi_k(x)$$

- Spectral method:

$$f_k(t) = \int f(x,t)\phi_k(x)\,dx$$

- Pseudospectral method (Lanczos 1938):

$$f_k(t) = \sum_{n=0}^{N-1} w_n f(x_n,t)\phi_k(x_n)$$

- Uses N *collocation points* $\{x_n\} \to \{f(x_n,t)\}$ (momentum space vs. position space)
- Compute spatial derivatives analytically
- Exponential convergence for smooth solutions
- No good for shocks (Gibbs)

Including Matter: BH-NS and NS-NS Collisions

- GW sources
- Short-duration GRBs

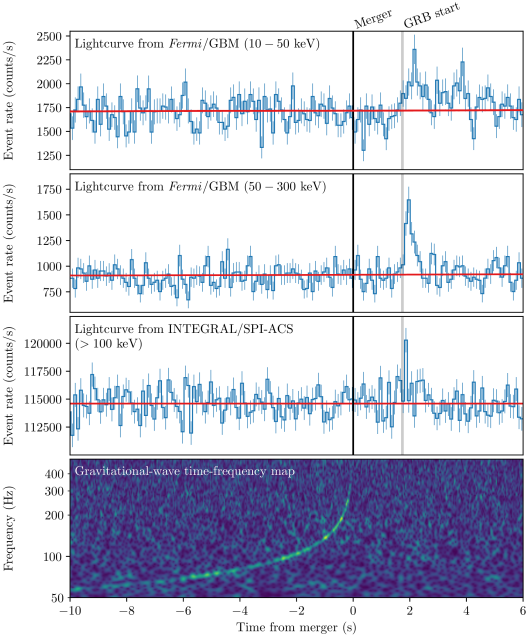

Including Matter: BH-NS and NS-NS Collisions

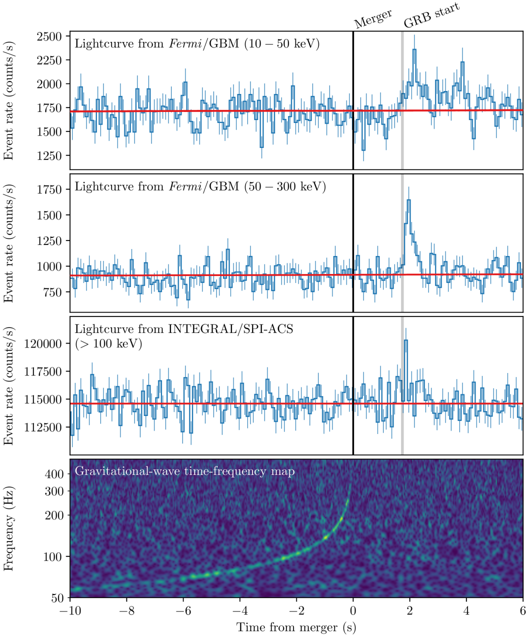

- GW sources
- Short-duration GRBs

Need full GR!

Multimessenger Astronomy

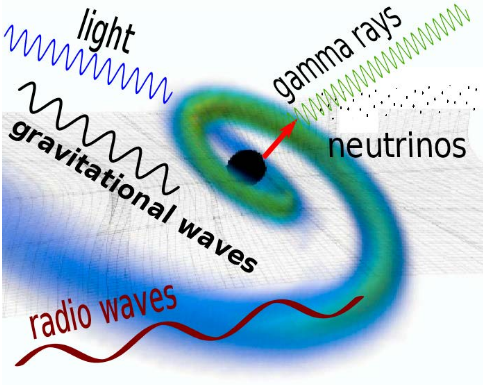

Multimessenger Astronomy Is Here!

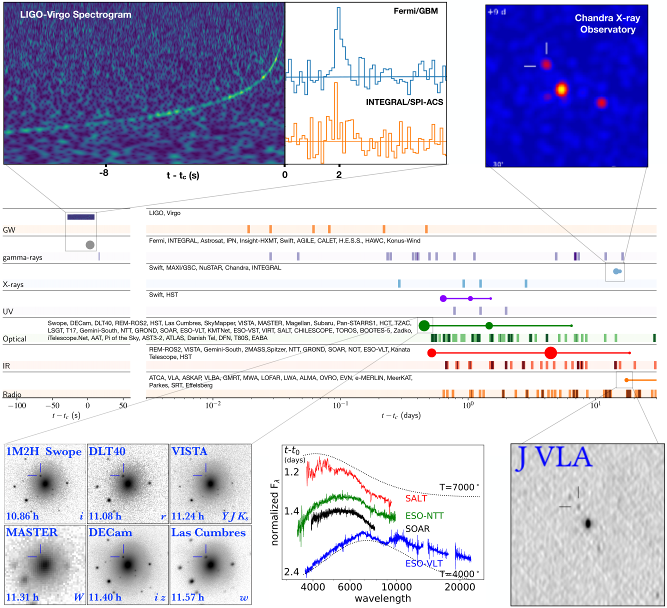

NSNS and BHNS: GR Hydro in SpEC

- Standard finite volume HRSC
  - WENO5 + HLL
  - MP5 + Roe
- FMR
- GR $\Leftrightarrow$ hydro grid via interpolation
- Accuracy
  - GW phase: ~ 1 radian (10 orbits)
  - BH: $\sim 1\%$
  - Matter: 10 – 50%

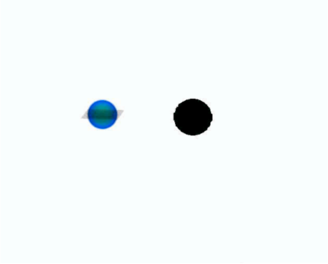

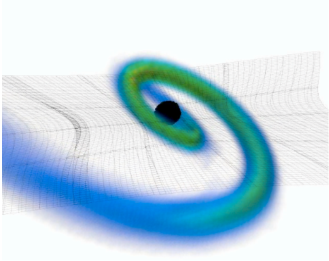

Challenges for BBH Codes

- LIGO SN will improve by ~ 3 in next 3 years
  - Event rate ~ 1 per day
  - Some events with SN ~ 100. Need $\delta\phi \lesssim .01$ at merger
- Another factor of 4 in 10 years (Voyager)
- LISA: SN $\sim 10,000$ in 2030

Challenges for Neutron Star Codes

- Computational errors are too large, 1 – 10%
- Can't even quantify the errors
- Simulations take too long

Challenges for Neutron Star Codes

- Computational errors are too large, 1 – 10%
- Can't even quantify the errors
- Simulations take too long
- Methods do not scale to extreme-scale machines

What's Wrong With 1% Accuracy?

What's Wrong With 1% Accuracy?

Examples:

- $M_{\rm disk} \sim 1\% M_{\rm tot}$
- Core-collapse supernovae
- EOS from tidal effects in NSNS or BHNS
- Wrong physics from unresolved scales, e.g. MRI

What's So Hard?

- NS surface + shocks $\Longrightarrow$ solution not smooth
- Multiple time scales
- Multiple spatial scales (adaptivity)
- Geometry changes (disruption, merger, black hole formation)
- Multiphysics (GR, hydro, MHD, neutrinos, photons, nuclear reactions,. . .)

The answer:

What's So Hard?

- NS surface + shocks $\Longrightarrow$ solution not smooth
- Multiple time scales
- Multiple spatial scales (adaptivity)
- Geometry changes (disruption, merger, black hole formation)
- Multiphysics (GR, hydro, MHD, neutrinos, photons, nuclear reactions,. . .)

The answer: (?)

Discontinuous Galerkin

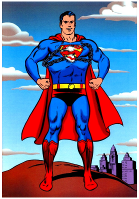

Finite Volume Methods

- solution represented by cell averages

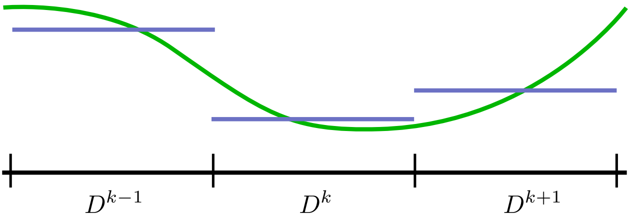

Figure: Francois Hebert

Finite Volume Methods

- solution represented by cell averages

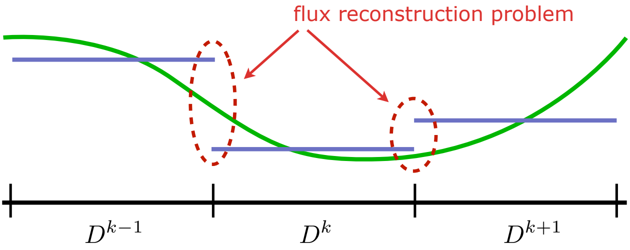

Figure: Francois Hebert

Finite Volume Methods

- solution represented by cell averages
- flux reconstruction can handle shocks
- but high order requires wide stencils

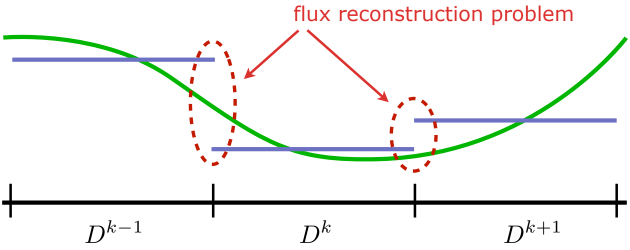

Figure: Francois Hebert

Spectral Code

- solution expanded on a local basis

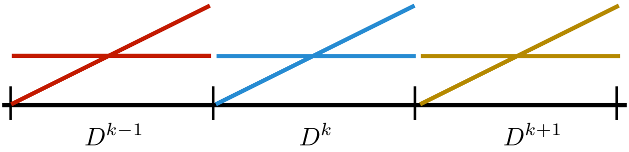

Spectral Code

- solution expanded on a local basis

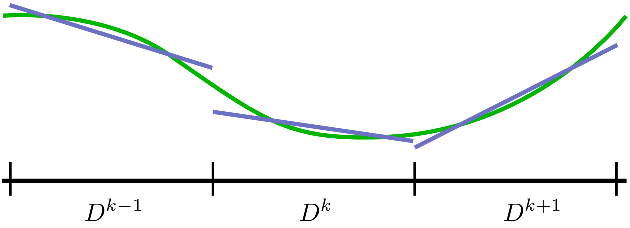

Spectral Code

- solution expanded on a local basis
- exponential convergence in smooth regions
- but flux can't do shocks

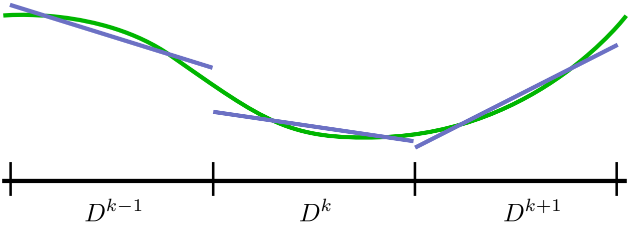

DG Code

- solution expanded on a local basis (local high order)

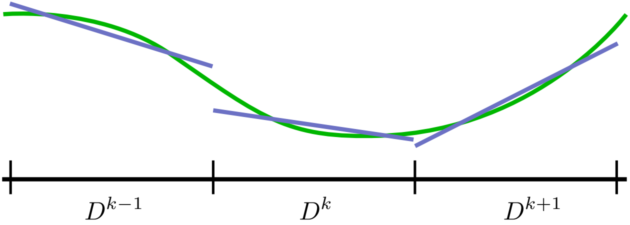

DG Code

- solution expanded on a local basis (local high order)

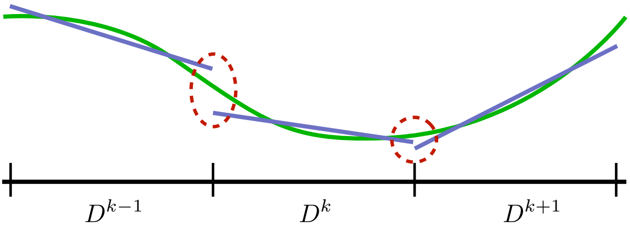

DG Code

- solution expanded on a local basis (local high order)
- exponential convergence in smooth regions
- formulation allows “arbitrary” fluxes—shocks OK

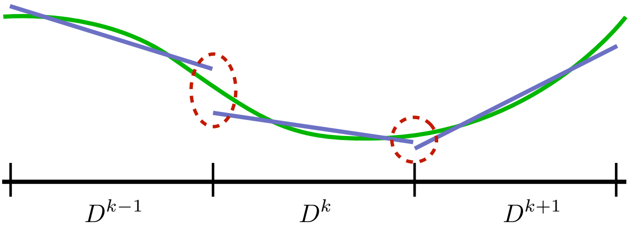

How Does DG Work?

$$\frac{\partial u}{\partial t} + \partial_a F^a(u) = s$$

Expand in basis functions:

$$u = \sum_i u_i \phi_i(\mathbf{x}), \qquad F^a = \sum_i F_i^a \phi_i(\mathbf{x}), \quad \ldots$$

N eqns. for u_i by projecting residual on space of test functions:

$$\int \left(\frac{\partial u}{\partial t} + \partial_a F^a - s \right) \phi_i(\mathbf{x}) \, d^3x = 0 \qquad \text{(Galerkin)}$$

The Standard Manipulation:

In each subdomain:

$$\int \partial_a F^a \, \phi_i(\mathbf{x}) \, d^3x =$$

The Standard Manipulation:

In each subdomain:

$$\int \partial_a F^a \, \phi_i(\mathbf{x}) \, d^3x = \int \partial_a (F^a \phi_i) \, d^3x - \int F^a \partial_a \phi_i \, d^3x$$

The Standard Manipulation:

In each subdomain:

$$\int \partial_a F^a \, \phi_i(\mathbf{x}) \, d^3x = \int \partial_a (F^a \phi_i) \, d^3x - \int F^a \partial_a \phi_i \, d^3x$$

$$= \oint F^a n_a \phi_i \, d^2S - \int F^a \partial_a \phi_i \, d^3x$$

The Standard Manipulation:

In each subdomain:

$$\begin{aligned}\int \partial_a F^a \, \phi_i(\mathbf{x}) \, d^3x &= \int \partial_a (F^a \phi_i) \, d^3x - \int F^a \partial_a \phi_i \, d^3x \\ &= \oint F^a n_a \phi_i \, d^2S - \int F^a \partial_a \phi_i \, d^3x \\ &\to \oint (F^a)^* n_a \phi_i \, d^2S - \int F^a \partial_a \phi_i \, d^3x\end{aligned}$$

The Standard Manipulation:

In each subdomain:

$$
\begin{aligned}
\int \partial_a F^a \, \phi_i(\mathbf{x}) \, d^3x &= \int \partial_a (F^a \phi_i) \, d^3x - \int F^a \partial_a \phi_i \, d^3x \\
&= \oint F^a n_a \phi_i \, d^2S - \int F^a \partial_a \phi_i \, d^3x \\
&\to \oint (F^a)^* n_a \phi_i \, d^2S - \int F^a \partial_a \phi_i \, d^3x
\end{aligned}
$$

- $(F^a)^* =$ numerical flux (art!)
- Generalize to curved spacetime (Teukolsky 2016)

Relativistic Hydro Test Example

Relativistic inflow

High-density gas

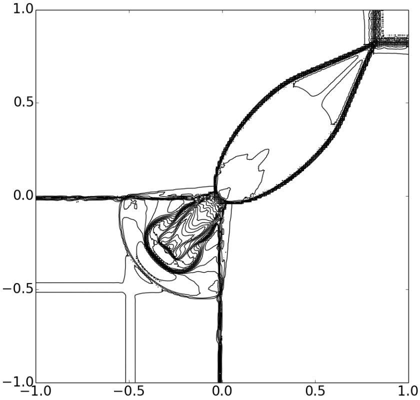

Relativistic inflow

Interacting shocks and contact discontinuities

Figure: Francois Hebert

SpECTRE: A Radically New Computer Code

- How will we accomplish our goal?
- Moore's Law is broken
- Next-generation machines will have millions of processors

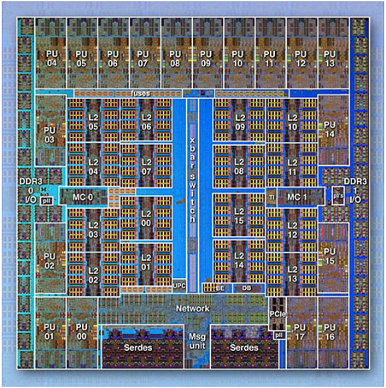

IBM Blue Gene Q chip
3/4" square
1.5 billion transistors

Why Not Run Current Codes on Millions of Processors?

- Currently, cells distributed across processors, MPI to communicate data
- Processors often idle during communication
- Load balancing: processors doing different amounts of computing
  - inside turbulent NS vs in vacuum
  - apparent horizon finding
  - trace light rays or neutrino paths (radiative transfer)

Solution: Task-based parallelism (in principle!)

A New Way to Parallelize

Conventional Parallelization (e.g. SpEC)

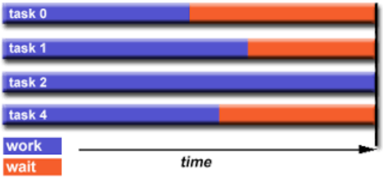

Task-based Parallelization (SpECTRE)

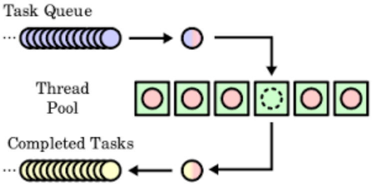

Implementing Task-Based Parallelism

- No standard packages
- MPI + OpenMP
- HPX
- Charm++
- ...

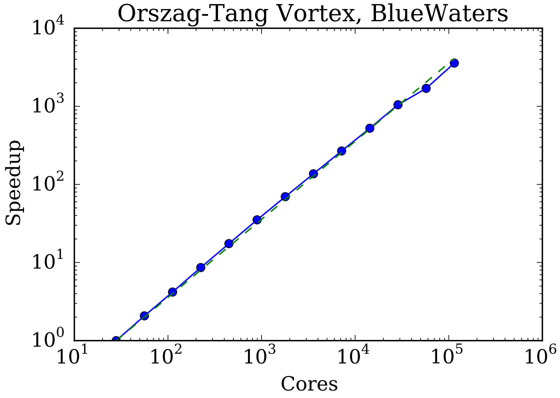

$330 \times 330 \times 2$ cells
2^3 points per cell

Figure: Scott Field

Time Profile

10 steps of relativistic MHD test

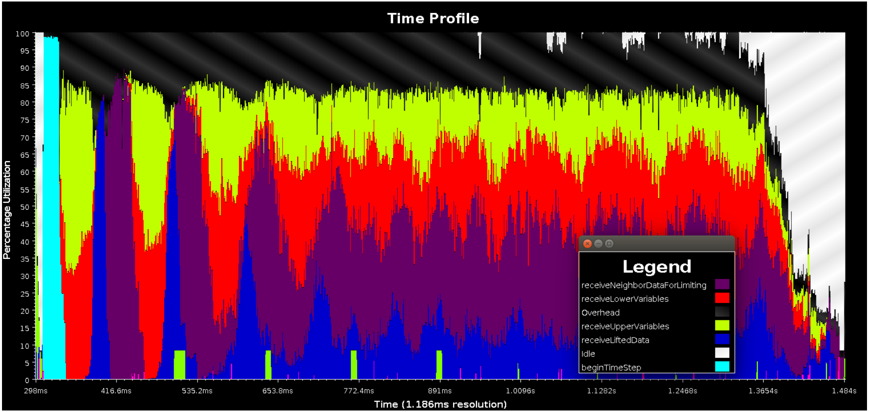

Red/Yellow: data to interfaces (hides RHS vol.)
Blue: fluxes to elements
Cyan: setup

Purple: slope limiting
Black: Charm++
White: idle

Challenges — Local Time Stepping

- AMR $\rightarrow$ large range of Courant conditions
- Advance each element with its own timestep (task-based!)

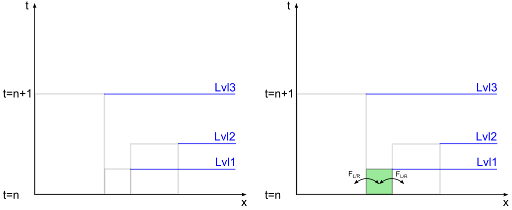

Calculate Fluxes for first time level edges and update 1st level cells

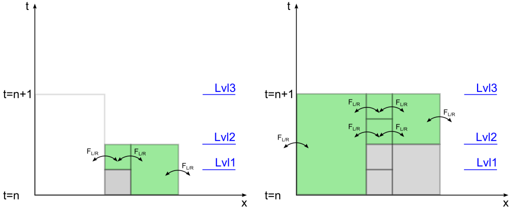

Calculate Fluxes for first and second time level edges and update 1st and 2nd level cells

Calculate Fluxes for first, second and third time level edges and update all cells to t=n+1

How to Fool a Computer Allocation Committee:

Advance all elements in lockstep! Perfect scaling, but only 10% of machine doing useful work

How to Fool a Computer Allocation Committee:

Advance all elements in lockstep! Perfect scaling, but only 10% of machine doing useful work

Dubey et al. (2014) survey of AMR packages:

- Scaling bad if local time stepping turned on
- Exception: Uintah (task-based parallelism)

Summary

- After 50 years of finite differencing, it’s time for us to move on if we want to tackle complex problems
- Algorithms like DG are high order, robust for shocks, local (good scaling)
- Task-based parallelism will enable exascale computing