Global Address Space Applications

Kathy Yelick

NERSC/LBNL and U.C. Berkeley

NERSC

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Advancing Computational Science of Scale—Producing Real Results

Algorithm Space

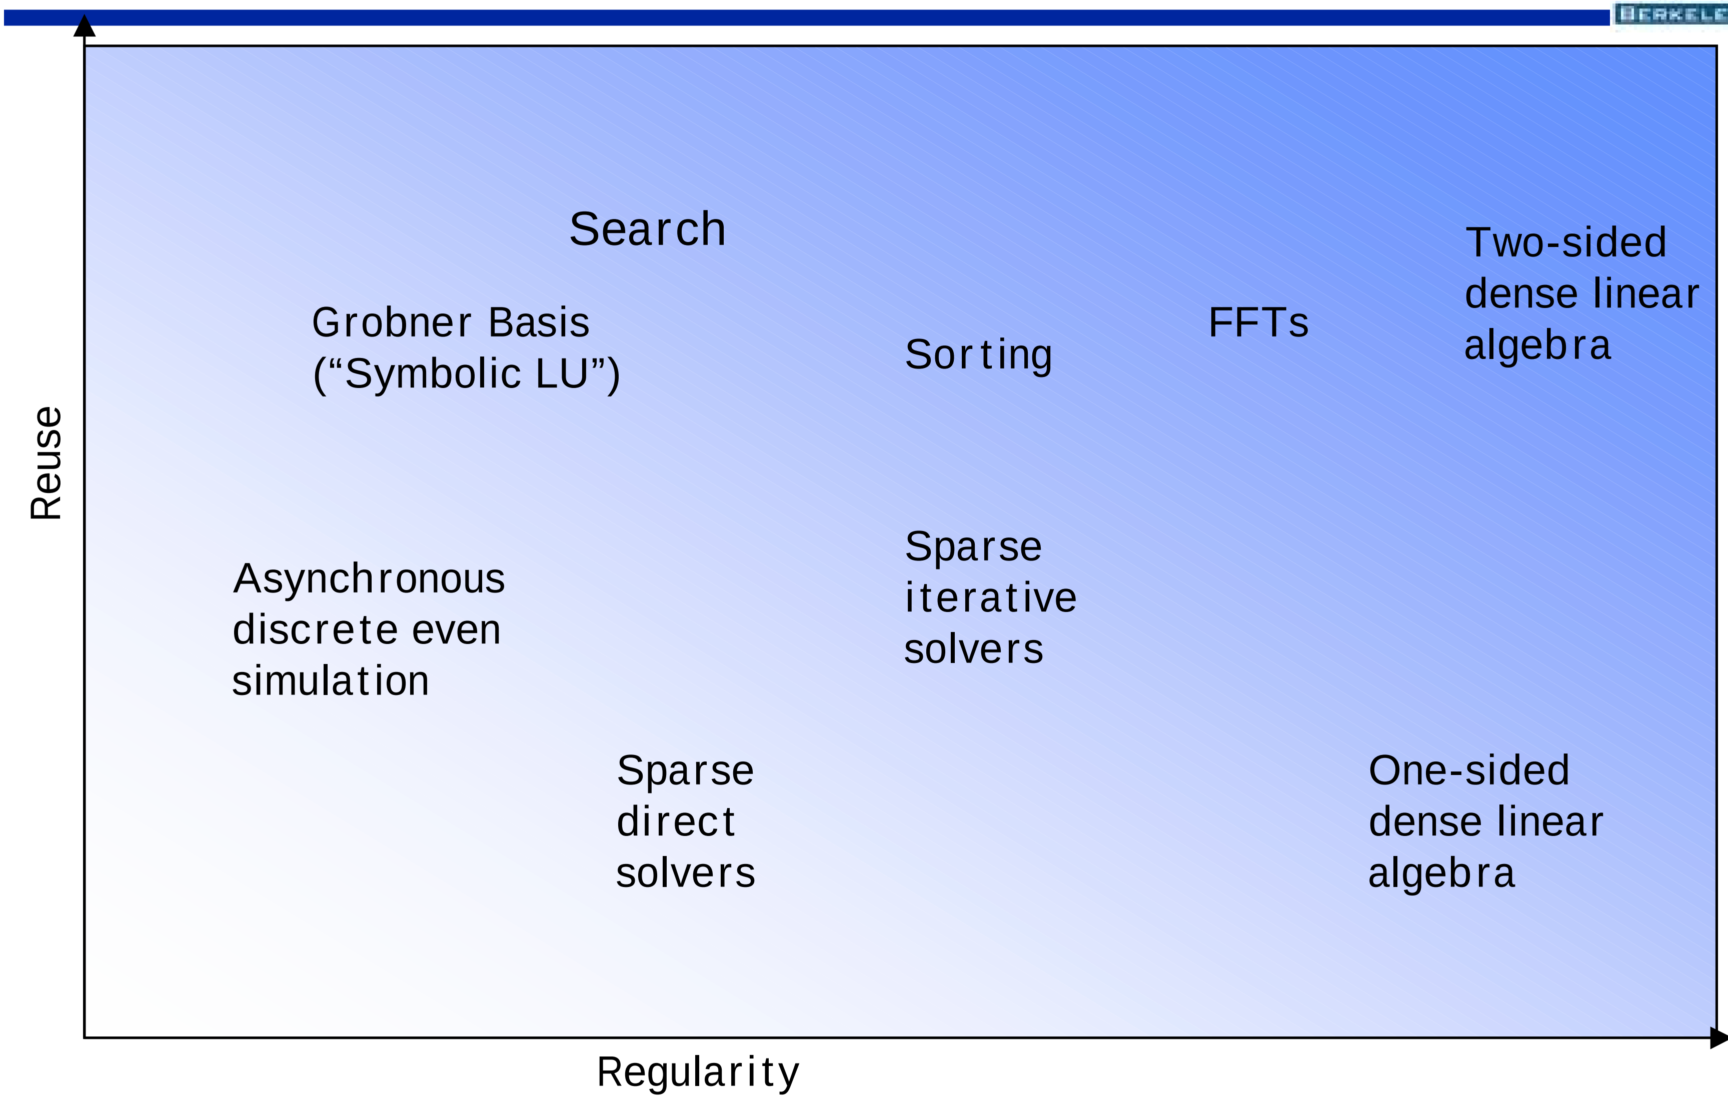

Scaling Applications

- Machine Parameters
  - Floating point performance
    - Application dependent, not theoretical peak
  - Amount of memory per processor
    - Use 1/10th for algorithm data
  - Communication Overhead
    - Time processor is busy sending a message
    - Cannot be overlapped
  - Communication Latency
    - Time across the network (can be overlapped)
  - Communication Bandwidth
    - Single node and bisection
- Back-of-the envelope calculations !

Running Sparse MVM on a Pflop

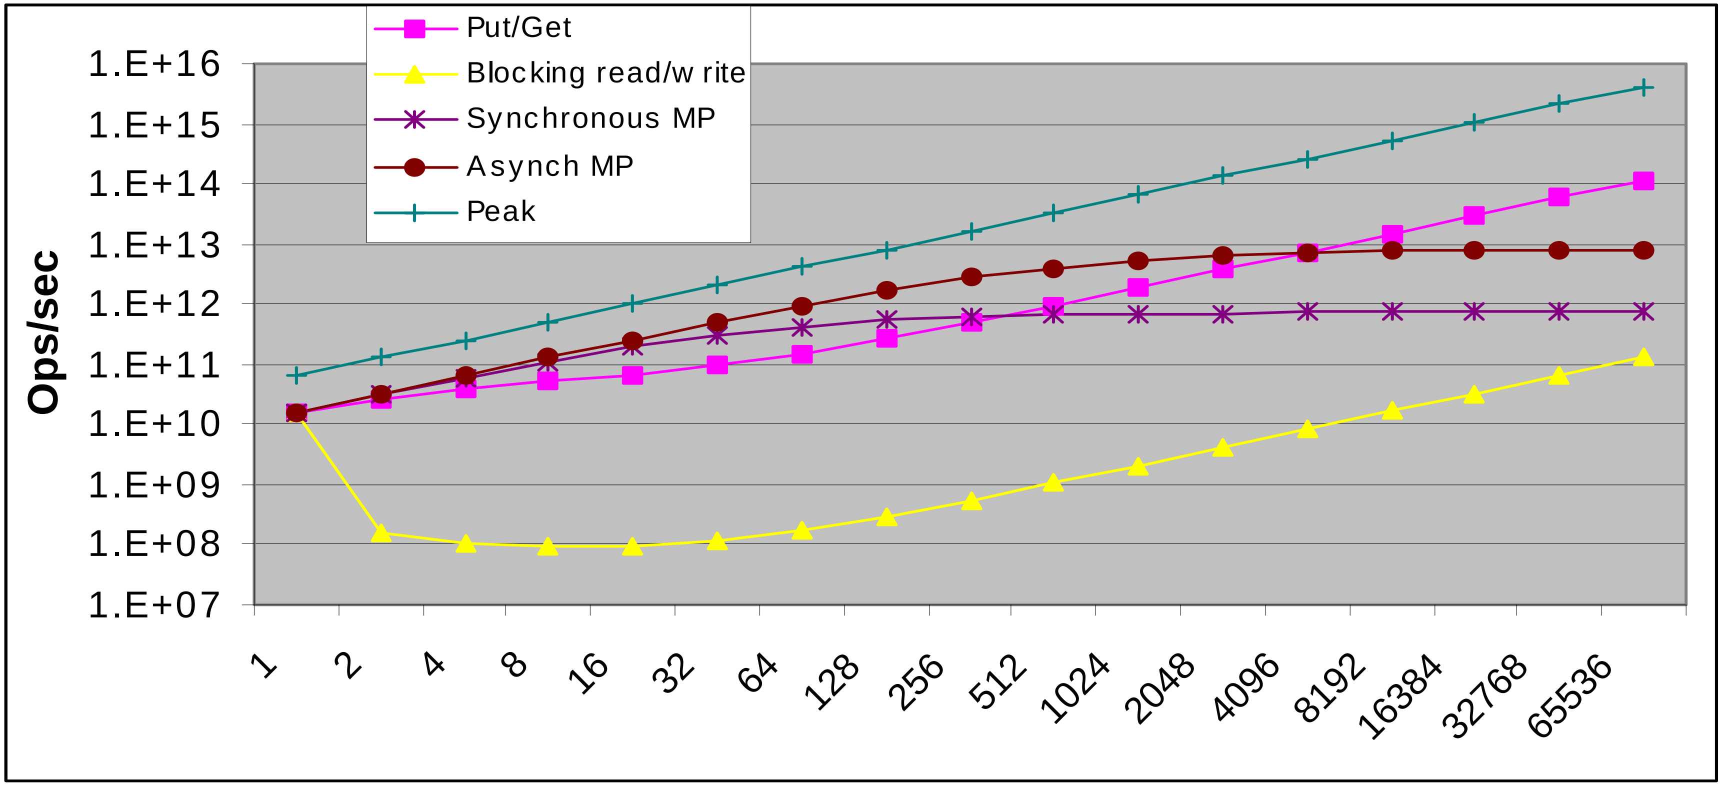

- 1 GHz * 8 pipes * 8 ALUs/Pipe = 64 GFLOPS/node peak
- 8 Address generators limit performance to 16 Gflops
- 500ns latency, 1 cycle put/get overhead, 100 cycle MP overhead
- Programmability differences too: packing vs. global address space

Effect of Memory Size

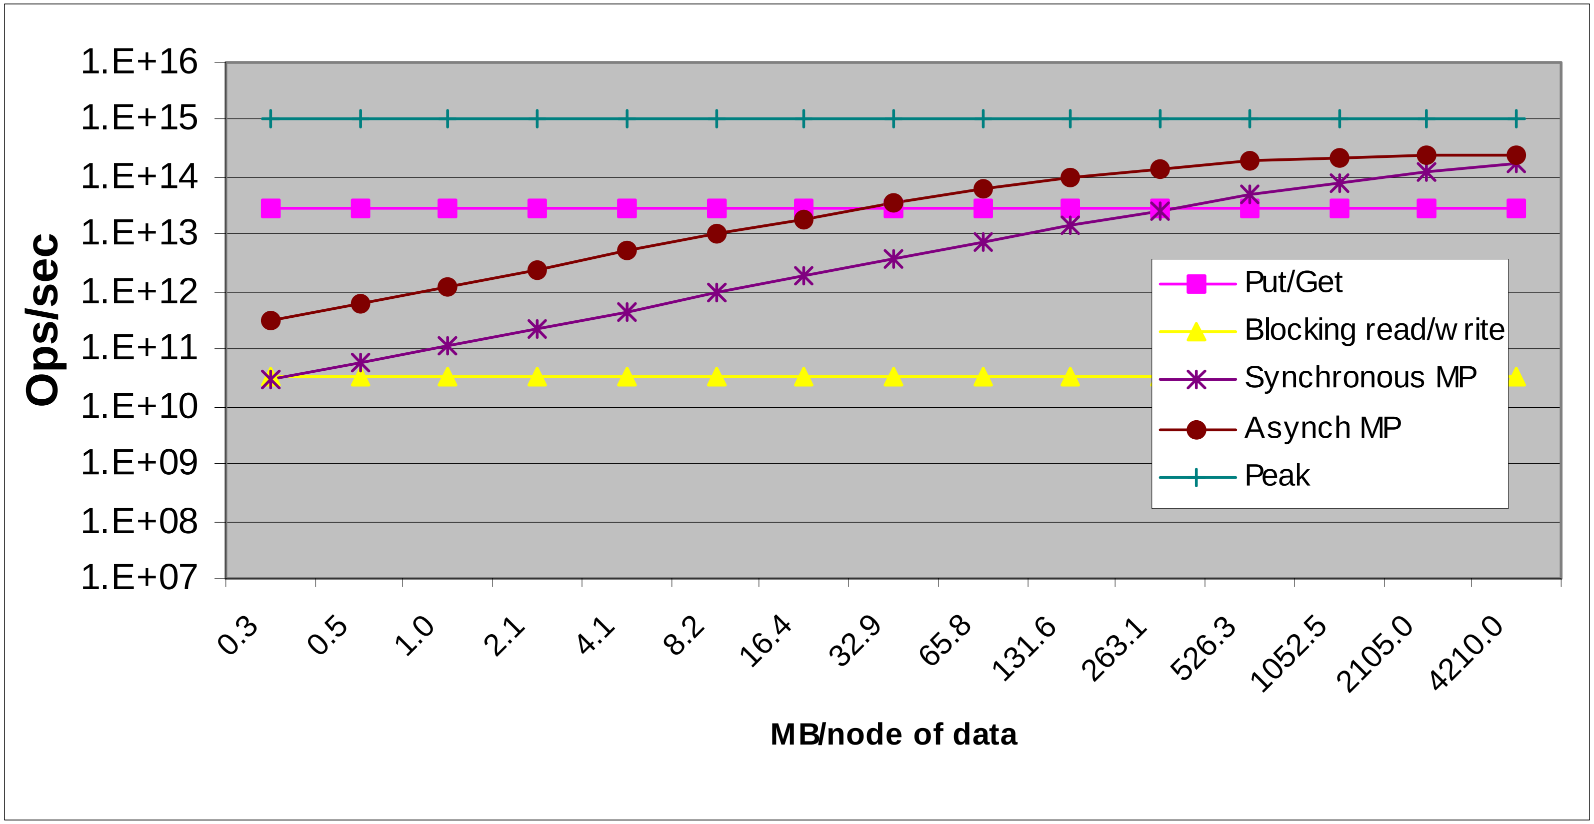

- **Low overhead is important for**
  - **Small memory nodes or smaller problem sizes**
  - **Programmability**

Parallel Applications in Titanium

- Genome Application
- Heart simulation
- AMR elliptic and hyperbolic solvers
- Scalable Poisson for infinite domains
- Genome application
- Several smaller benchmarks: EM3D, MatMul, LU, FFT, Join

MOOSE Application

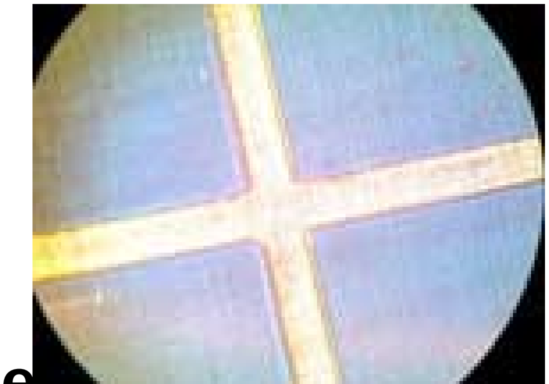

- **Problem: Microarray construction**
  - **Used for genome experiments**
  - **Possible medical applications long-term**
- **Microarray Optimal Oligo Selection Engine (MOOSE)**
  - **A parallel engine for selecting the best oligonucleotide sequences for genetic microarray testing**
  - **Uses dynamic load balancing within Titanium**

Heart Simulation

- Problem: compute blood flow in the heart
  - Model as elastic structure in incompressible fluid.
    - "Immersed Boundary Method" [Peskin and McQueen]
    - Particle/Mesh method stress communication performance
    - 20 years of development in model
  - Many other applications: blood clotting, inner ear, insect flight, embryo growth,…
- Can be used for design of prosthetics
  - Artificial heart valves
  - Cochlear implants

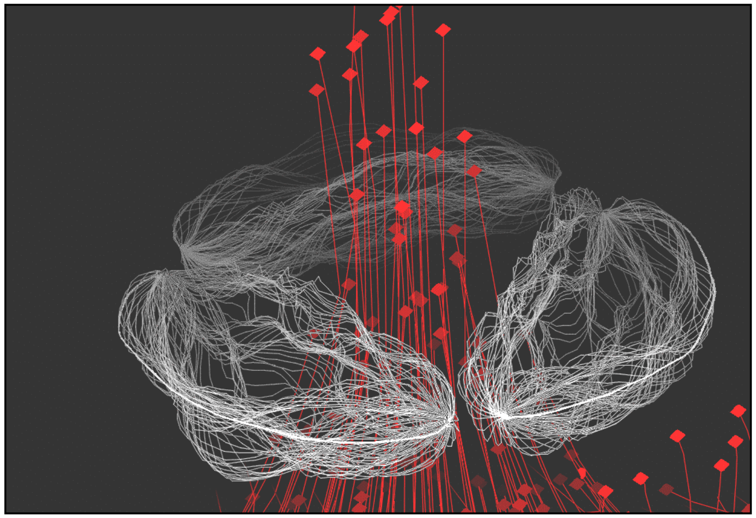

Scalable Poisson Solver

- MLC for Finite-Differences by Balls and Colella
- Poisson equation with infinite boundaries
  - arise in astrophysics, some biological systems, etc.
- Method is scalable
  - Low communication
- Performance on
  - SP2 (shown) and t3e
  - scaled speedups
  - nearly ideal (flat)
- Currently 2D and non-adaptive
- Point charge example shown
  - Rings & star charges
  - Relative error shown

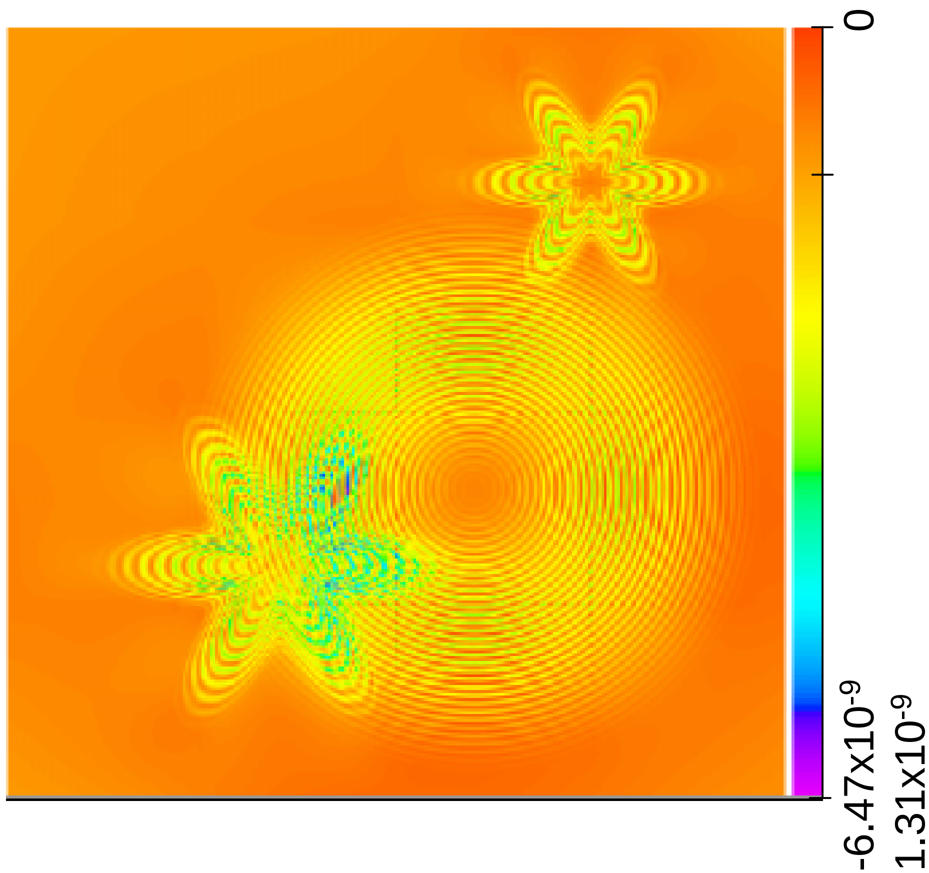

AMR Gas Dynamics

- Developed by McCorquodale and Colella
- 2D Example (3D supported)
  - Mach-10 shock on solid surface at oblique angle

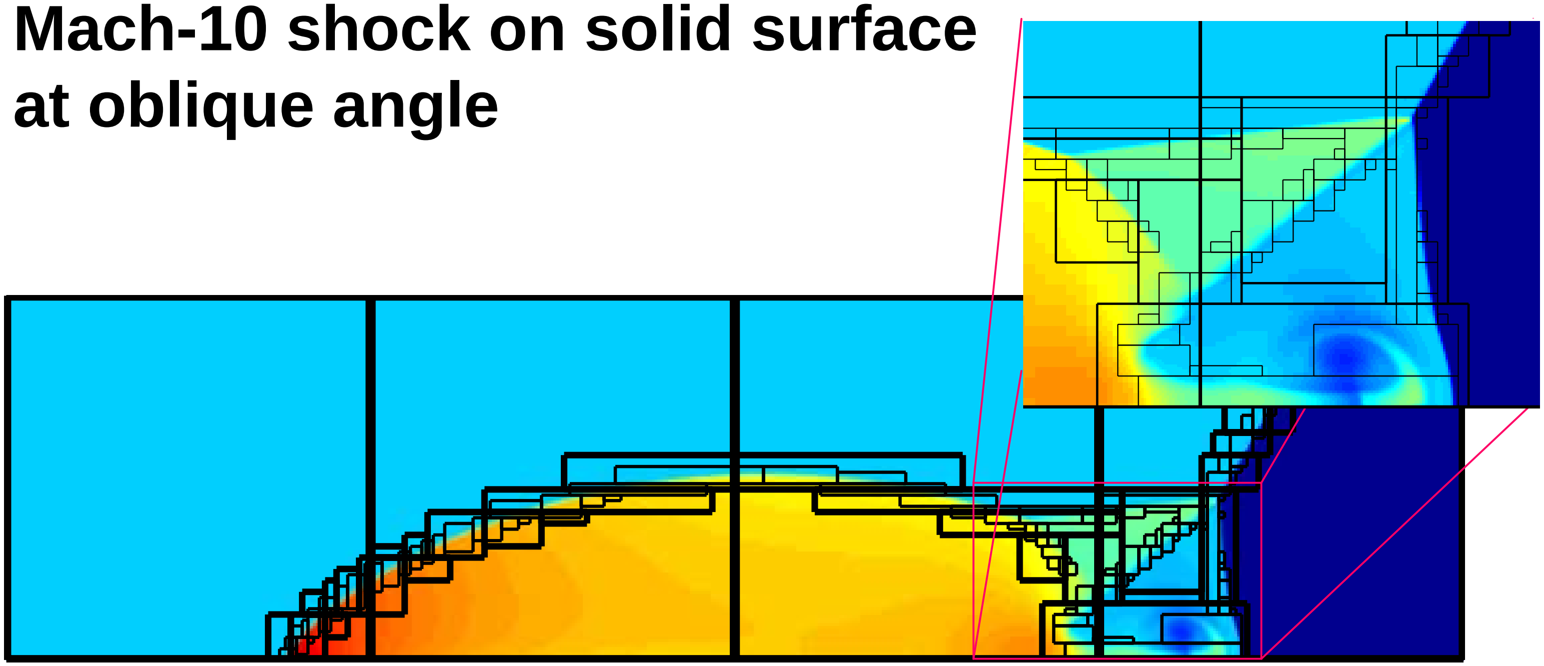

- Future: Self-gravitating gas dynamics package

UPC Application Investigations

- **Pyramid**
  - **3D Mesh generation [Shewchuk]**
  - **2D version (triangle) critical in Quake project**
  - **Written in C, challenge to parallelize**
- **SuperLU**
  - **Sparse direct solver [Li,Demmel]**
  - **Written in C+MPI or threads**
  - **UPC may enable new algorithmic techniques**
- **N-Body simulation**
  - **“Simulating the Universe”**

Summary

- UPC Killer App should
  - Leverage programmability: hard in MPI
  - Use fine-grained, irregular, asynchronous communication
- Libraries
  - Must allow for interface to libraries
  - MPI libraries, multithreaded libraries, serial libraries
- Compilation needs
  - High performance on at least one machine
  - Portability across many machines