SAND2012-2592P

Preparing Multi-physics, Multi-scale Codes for Exascale Computing

Richard Barrett
Center for Computing Research (1400)
SAND Number

Intel EPOCH MIC workshop
April 2, 2012

Collaborators: **Paul Crozier, Doug Doerfler, Simon Hammond, Michael Heroux, Paul Lin, Mahesh Rajan, Courtenay Vaughan, Alan Williams.**

Sandia National Laboratories is a multi-program laboratory operated and managed by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin company, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

National Nuclear Security Administration

Sandia National Laboratories

Our goals for this visit

- A strengthened understanding of the MIC architecture: current, emerging, and future, esp. within the context of our applications, leading to

- a strengthened understanding of the methodology employed for successfully exploiting the architecture's capabilities and performance potential.

The problem

- Codes are million lines,
- under constant development,
- controlled distribution,
- limited access to computing environments,
- etc.

The Mantevo Project

- focused on developing tools to accelerate and improve the design of high performance computers and applications by providing application and library proxies to the high performance computing community.

A miniapp is

- a proxy for a larger application,
- written by application code developers,
- representing a key performance issue impacting applications.
- O(k) SLOC, Open Source (LGPL)
- “Output” is information.

miniMD

- Lennard-Jones (LJ) atomic interaction molecular dynamics
- Proxy for LAMMPS LJ. http://lammps.sandia.gov/
- C, MPI; OpenMP
- Performance issues
  - managing threads operating on memory indirection accessed data.
  - frequent inter-process communication of relatively small messages.

miniMD
Force calculation

```
for (i = 0; i < nlocal; i++) {
   neighs = neighbor.firstneigh[i];
   numneigh = neighbor.numneigh[i];
   xtmp = x[i][0];
   ytmp = x[i][1];
   ztmp = x[i][2];
   for (k = 0; k < numneigh; k++) {
     j = neighs[k];
     delx = xtmp - x[j][0];
     dely = ytmp - x[j][1];
     delz = ztmp - x[j][2];
     rsq = delx*delx + dely*dely + delz*delz;
     if (rsq < cutforcesq) {
       sr2 = 1.0/rsq;
       sr6 = sr2*sr2*sr2;
       force = sr6*(sr6-0.5)*sr2;
       f[i][0] += delx*force;
       f[i][1] += dely*force;
       f[i][2] += delz*force;
       f[j][0] -= delx*force;
       f[j][1] -= dely*force;
       f[j][2] -= delz*force;
     }
   }
 }
```

LAMMPS(LJ) and miniMD performance

(Dell workstation, Intel dual-socket Xeon)

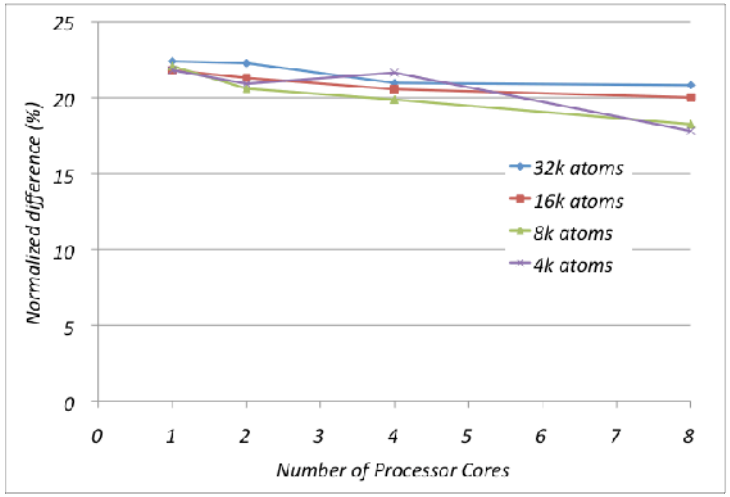

Total time

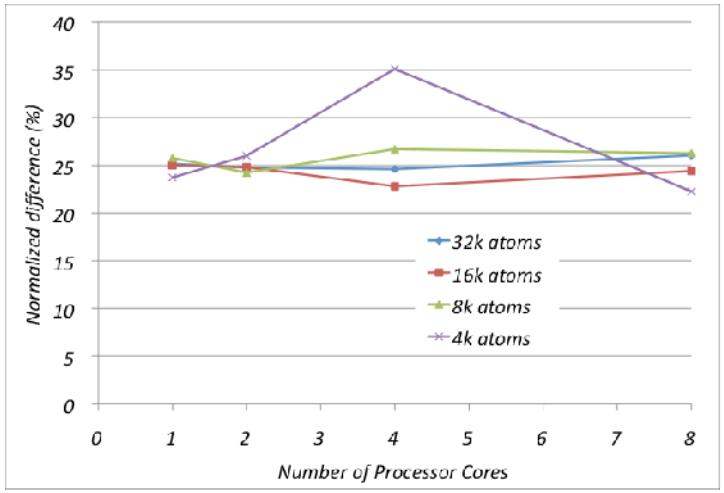

Force

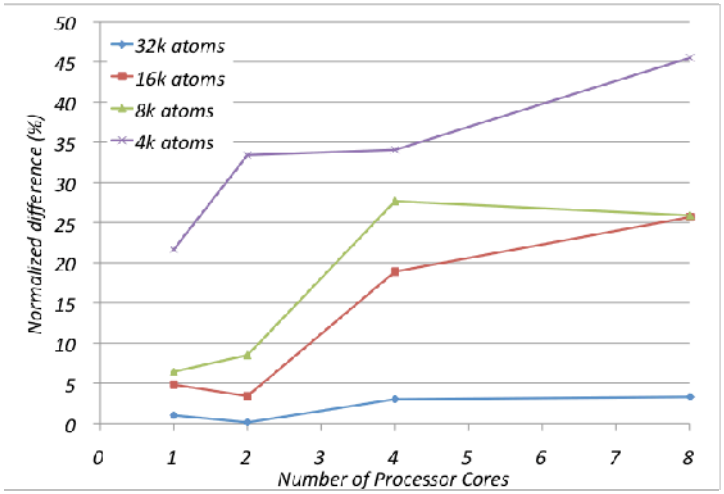

Neighbors

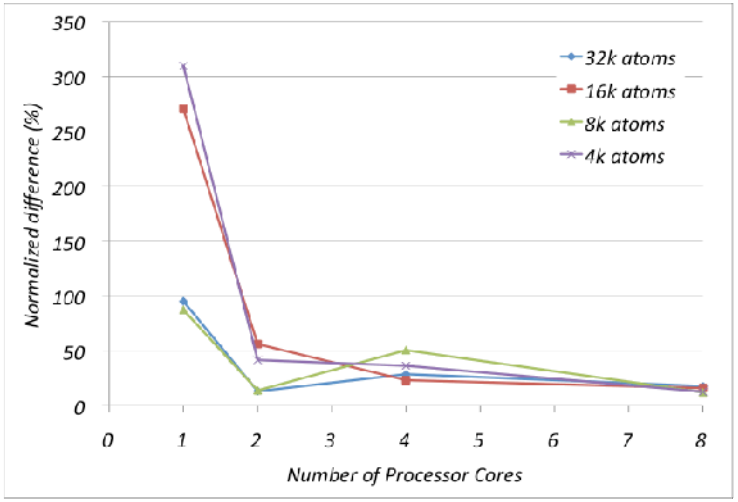

Communication

Sandia National Laboratories

miniFE

- Implicit Finite Element solver and matrix assembly
- Proxy for Charon, an electronic device simulation code. http://charleston.sandia.gov/Charon/
- C++, MPI; OpenMP, TBB, Cilk
- Greg Skinner has shown some great results using MKL sparse matrix-vector product. We are even more interested in the general case, that is, where the matrix structure and other issues are such that we cannot use a library function.

Charon and miniFE performance

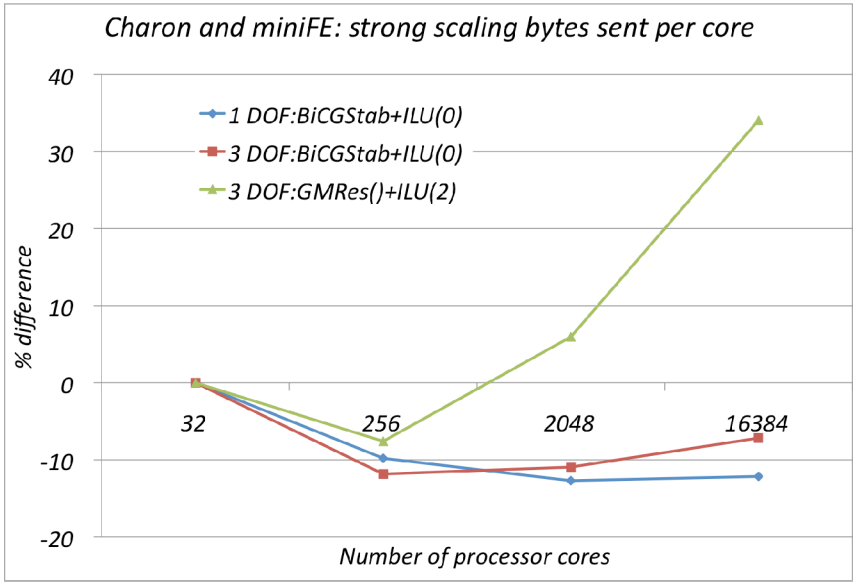

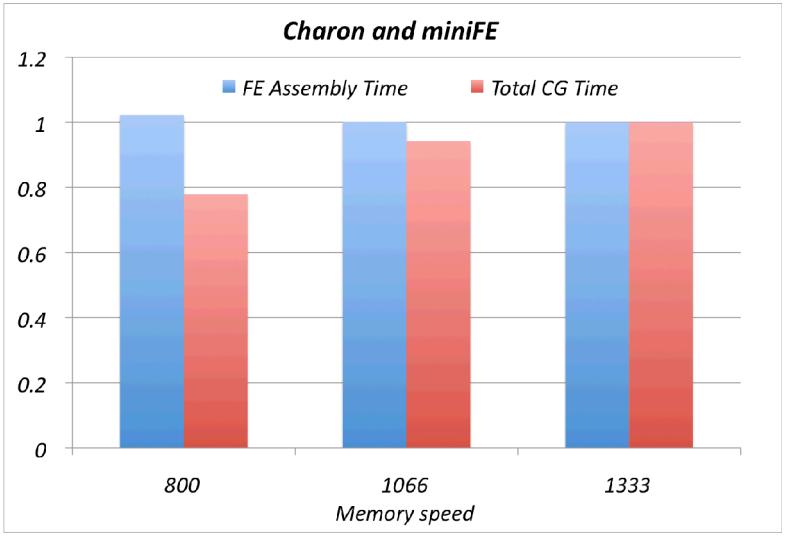

Sandia National Laboratories

minGhost

- Finite difference stencil computation
- Proxy for CTH, an Eulerian multi-material finite volume code
- Fortran, MPI; OpenMP, OpenACC
- Exploring alternatives to massive message aggregation BSP: How do alternative programming models impact?

CTH

DO I = 1, NUM_VARS

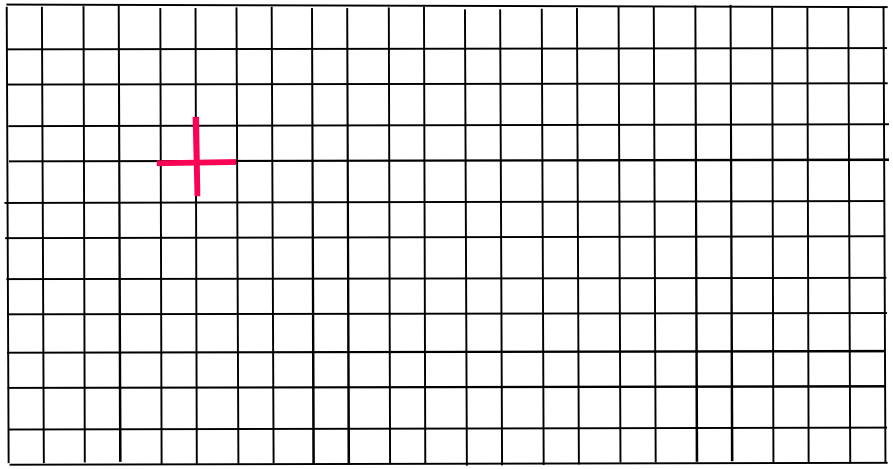

END DO

DO I = 1, NUM_VARS

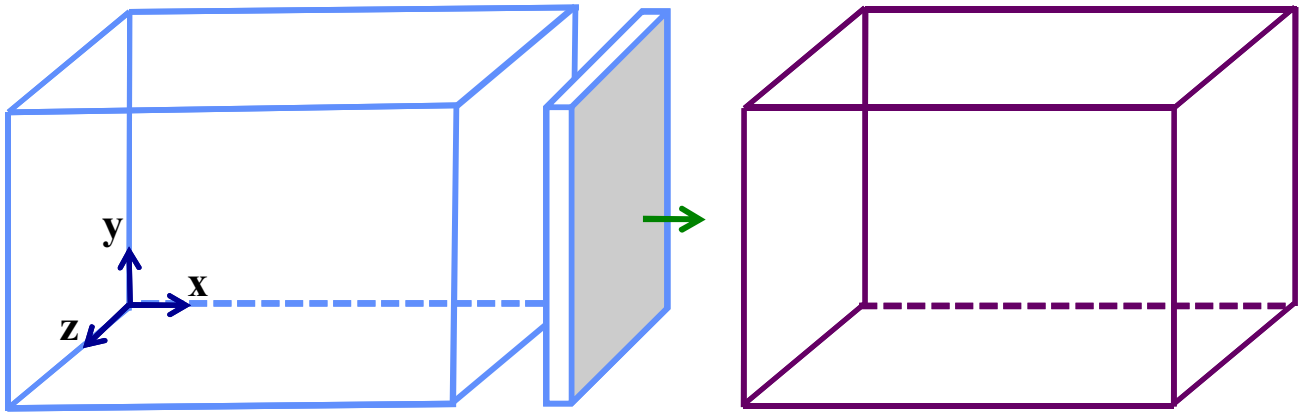

END DO

miniGhost performance

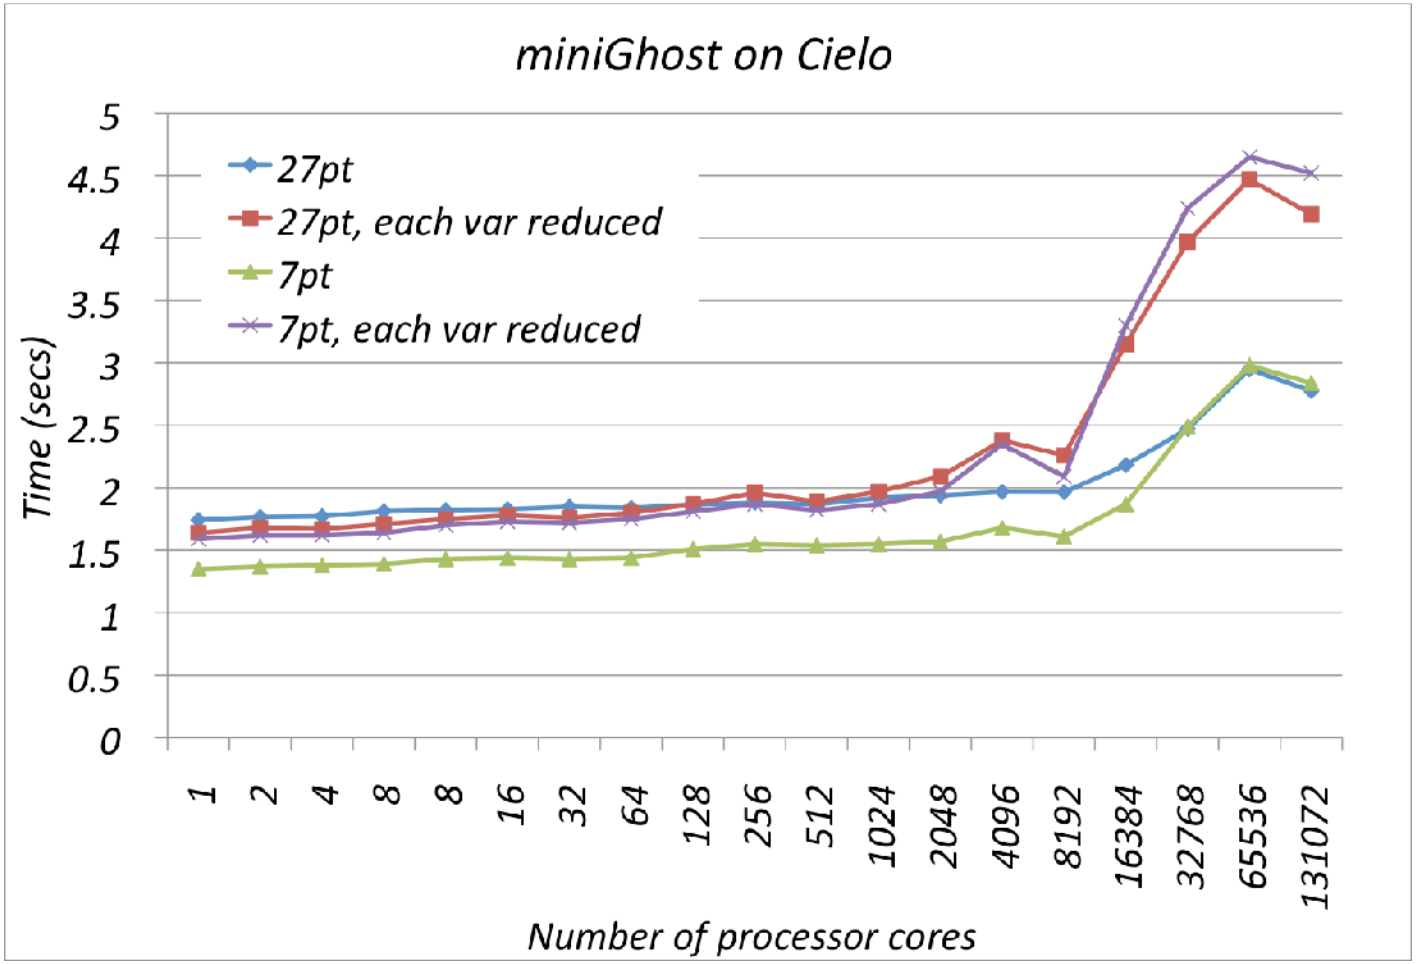

Runtime profiles

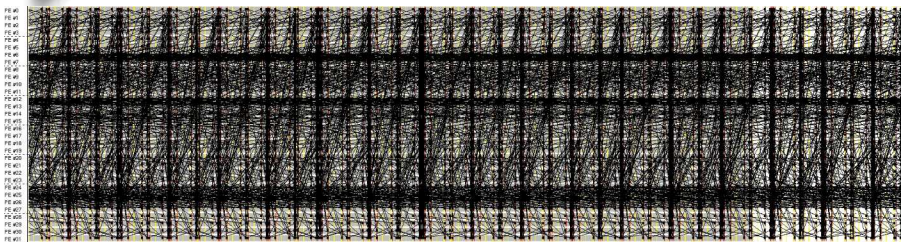

Charon

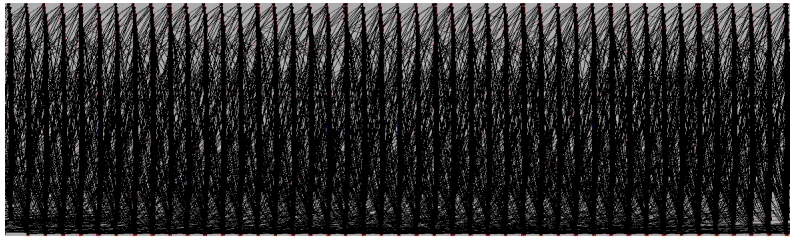

miniFE

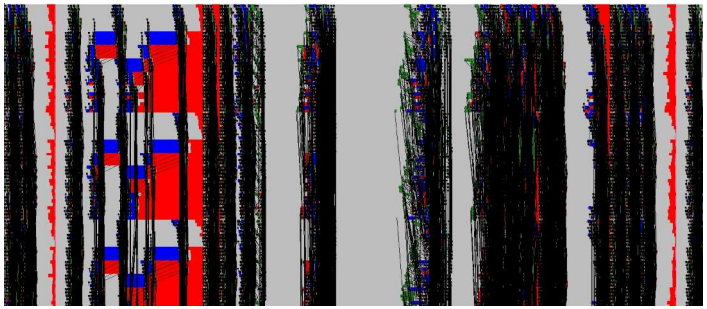

CTH

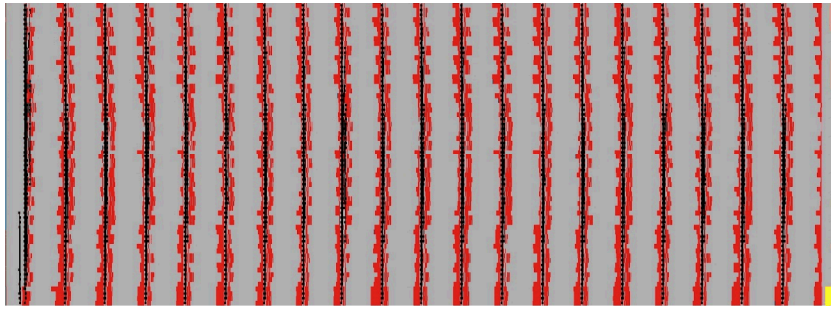

miniGhost

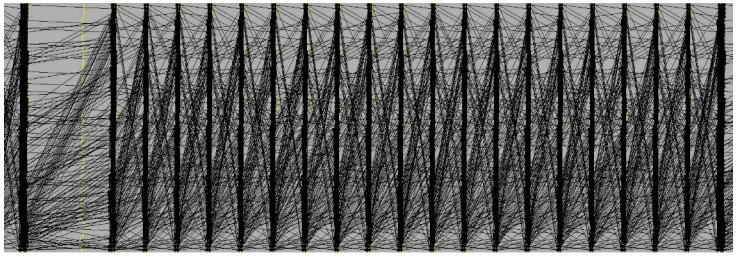

LAMMPS

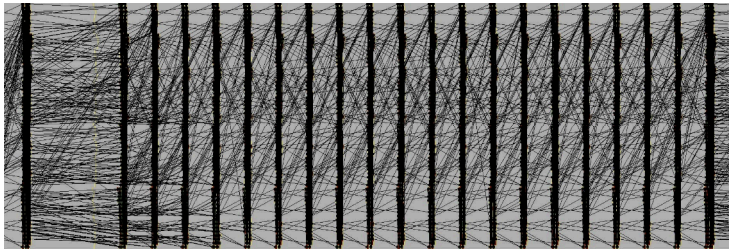

miniMD

Communication Patterns

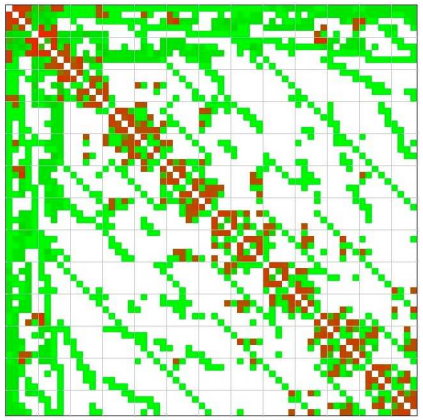

Charon

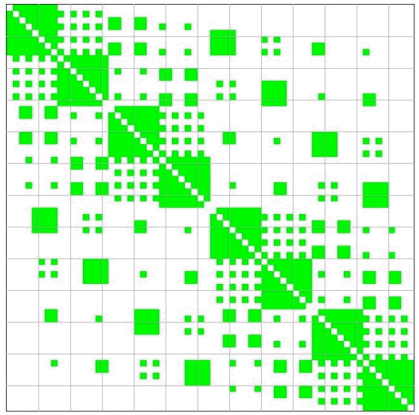

miniFE

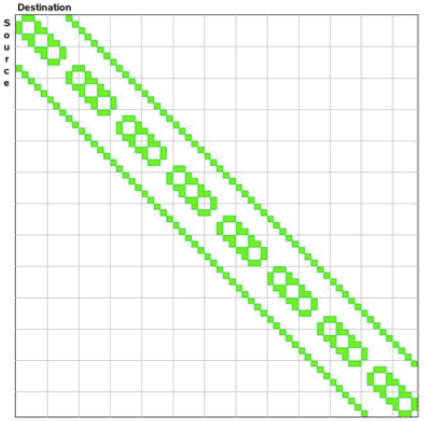

CTH (meso)

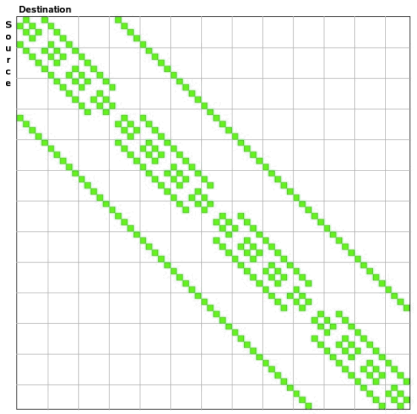

miniGhost

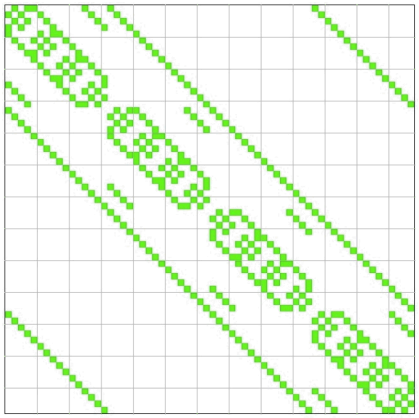

LAMMPS

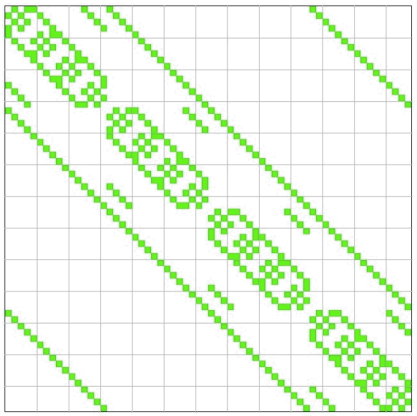

miniMD

ALEGRA data structure

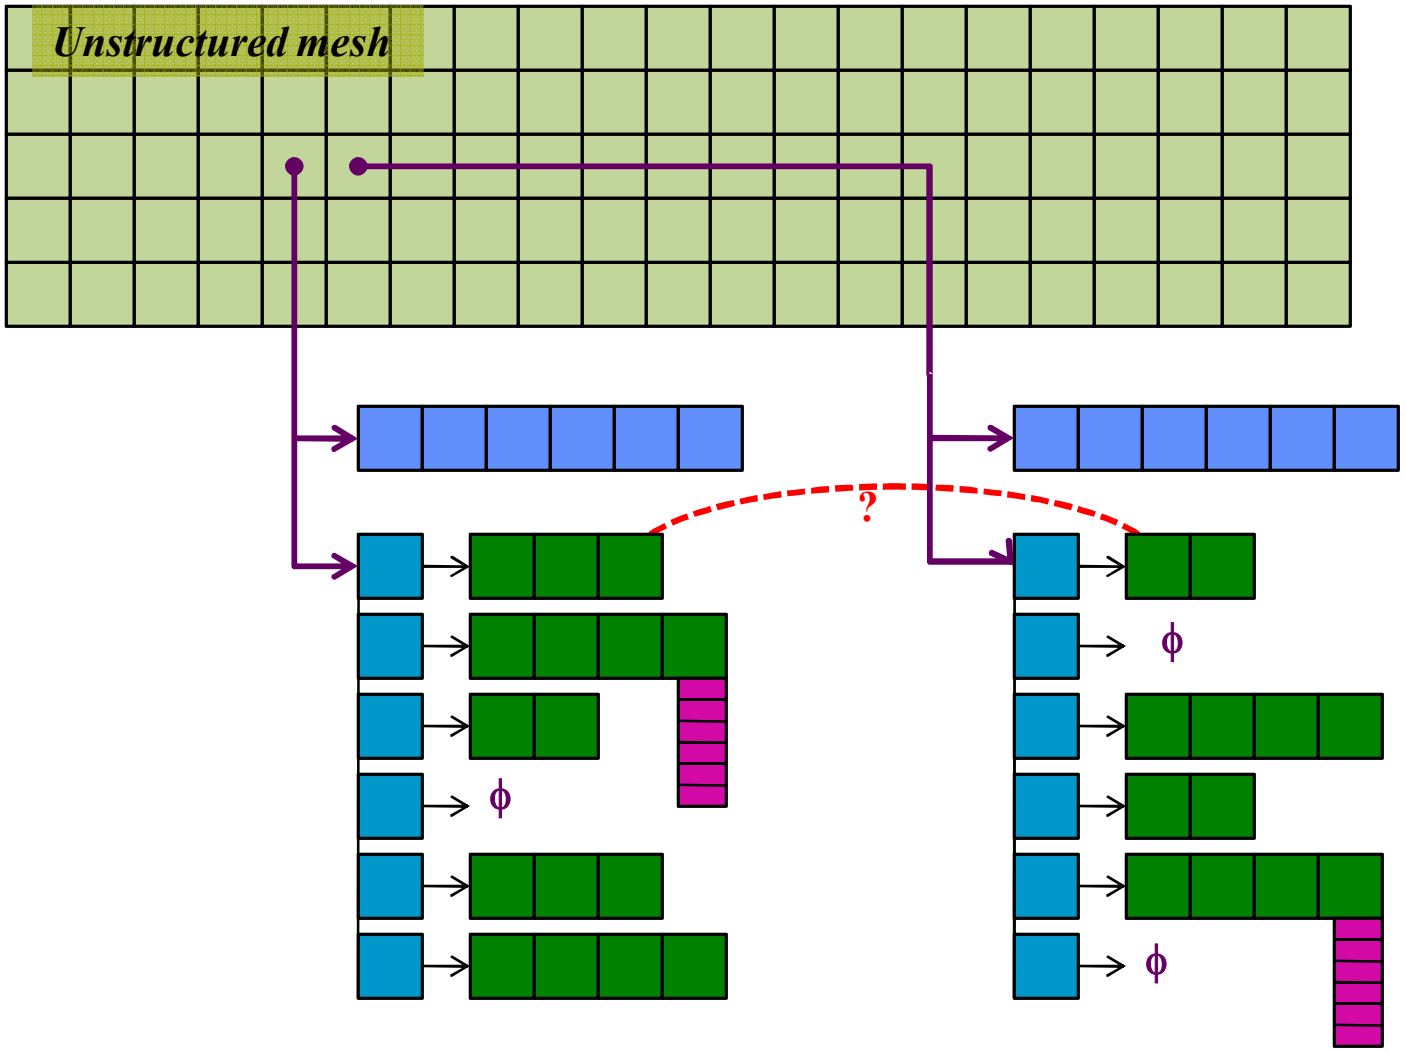

Some references

- "Improving Performance via Mini-applications", Heroux et al, Sandia technical report SAND2009-5574, 2009. http://software.sandia.gov/mantevo
- "Assessing the Validity of the Role of Mini-Applications in Predicting Key Performance Characteristics of Scientific and Engineering Applications", R. Barrett, P. Crozier, D. Doerfler, S. Hammond, M. Heroux, P. Lin, H. Thornquist, T. Trucano, C. Vaughan. In preparation (June 30).
- "Navigating An Evolutionary Fast Path to Exascale", Barrett, Barragy, Doerfler, Hammond, Luitjens, Roweth, Vaughan, submitting to SC12.
- "MiniGhost: A Miniapp for Exploring Boundary Exchange Strategies Using Stencil Computations in Scientific Parallel Computing", Barrett, Heroux, Vaughan, SAND report 2012-2437.
- "A case for application-specific processors", W. Alkohlani and J. Cook, New Mexico State Univ. In preparation.
- "Energy Impacts of Accelerating Data Assembly in Finite Element Methods on GPUs", Li Tang, Sharon Hu, S.Hammond, and R. Barrett, submitting to SC12.
- "On the Viability of Checkpoint Compression for Extreme Scale Fault Tolerance", D. Ibtesham, D. Arnold, K. Ferreira, and P. Bridges, Resilience Workshop@EuroPar'11.