SAND2013-3294P

Photos placed in horizontal position with even amount of white space between photos and header

Photos placed in horizontal position with even amount of white space between photos and header

Discussion of some new and developing miniapps

Intel Xeon Phi dungeon
April 22-25, 2013
Hillsboro, OR

Richard F. Barrett
Scalable Computer Architectures
Sandia National Laboratories, NM
rfbarre@sandia.gov

SAND TBD.

Sandia National Laboratories

Exceptional service in the national interest

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

miniGhost tp

Sandia National Laboratories

- Representative of finite difference (or volume) stencil computations in a multi-material application.
- Domain decomposed across parallel (MPI) processes, then
- over-decomposed as tasks to be processed by threads.
- Each task applies the difference stencil.
- Task domains on an inter-process boundary exchange face data using blocking MPI calls (send/recv).
  - Scheduler must be able to swap tasks in and out based on MPI activity. This capability is being developed in SPR@Sandia.
- Tasks on a physical boundary apply boundary conditions.

miniGhost tp

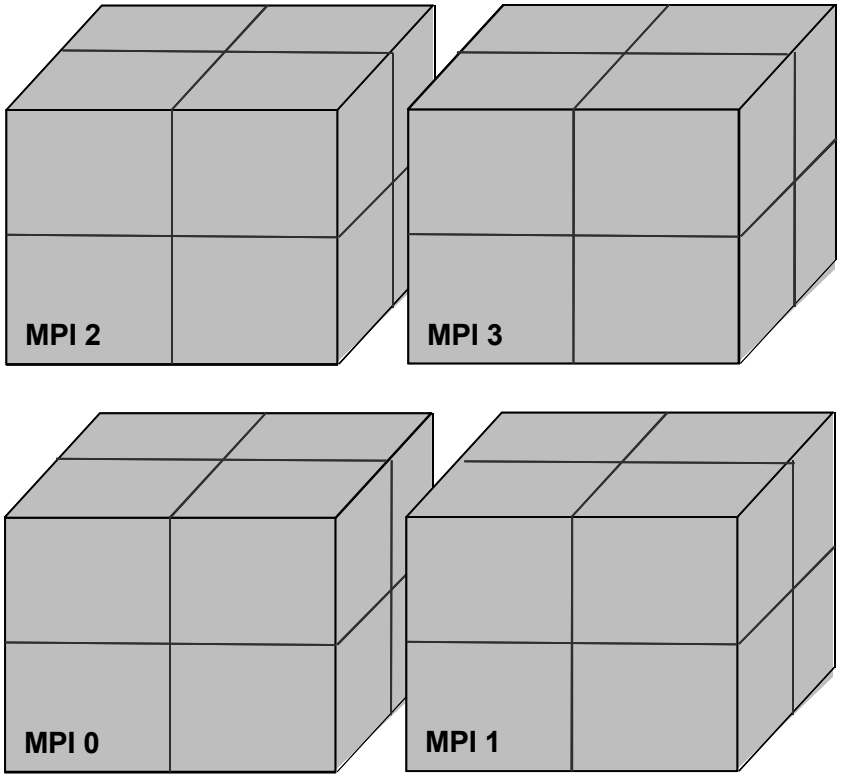

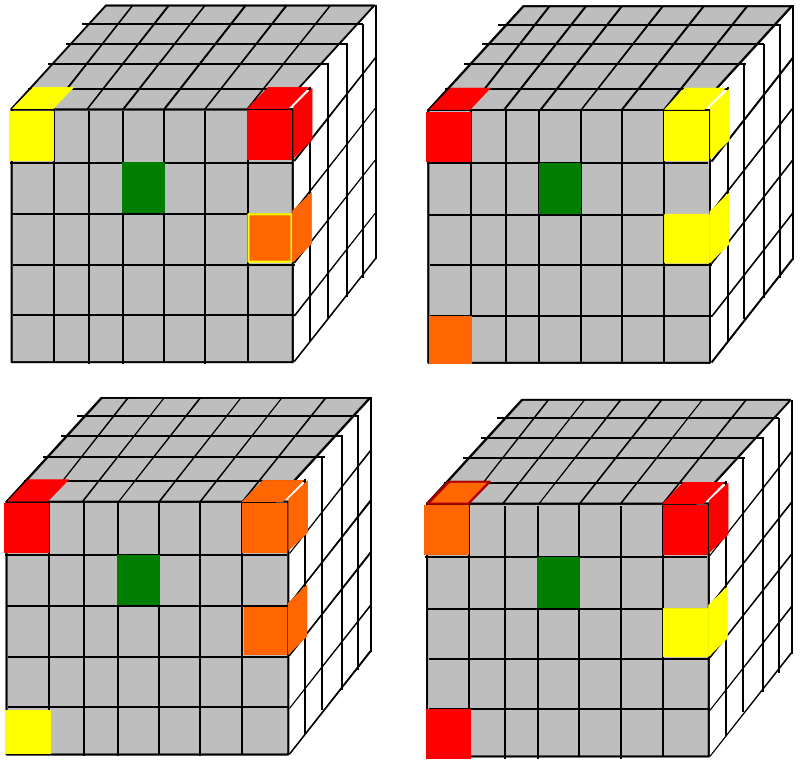

Data parallel

thread

Task parallel

Computation

Computation + MPI

Computation + BC

Computation + BC + MPI

miniGhost tp : status

Sandia National Laboratories

- C language, MPI internode.
- omp parallel single thread region spawns tasks
- Blocks defined within triply nested loop (i,j,k)
  - If block = local dimension, equivalent to reference implementation.
- qthreads implementation soon.
- Beginning to explore alternative mapping via DAGs.

miniGhost tp : Adding AMR

Sandia National Laboratories

- Diffusion over the 3d domain with random initial conditions and reflective boundary conditions.

  Two options:

  - uniform refinement, or
  - refinement based on the boundary or volume of an object being moved through the mesh and changing size. So, for example, a shock front can be simulated by refining based on a sphere which starts small and grows in size as the problem advances.
- Refinement within blocks.
- A block is refined into 8 blocks.
- Neighbors must be within one level of refinement.
- Computation is self-contained within a block.
- Communication aggregated to BSP model.
  - Excellent candidate for task parallelism version.

Data analytics

Miniapps in development:

- PathFinder
- Enumerate all triangles in a graph
  - Preliminary steps:
    - Traverse adjacency lists
    - Do previous with “Manhattan Loop Collapse”
    - Build bucket structure for triangle enumeration
    - Traverse that structure
- Boolean Satisfiability (SAT) : e.g., (A∧B)∨C
- Analysis of combustion simulation data using sublinear sampling of statistics on critical interfaces to determine parameters for the next iteration
  - Scientific-computing communication patterns and random access consistent with other big data

PathFinder

Sandia National Laboratories

searches for specified signatures in graphs.

- A signature is a specific path or subgraph that
- represents some higher-order feature such as a specific piece of functionality.
- Typical problems require identifying multiple signatures.
- Nodes are labeled with characteristics. Signatures represent a sequence of characteristics that the program is expected to execute at runtime.
- Identify a path through the graph where each characteristic in a signature appears in order.
  - Thousands of signatures in thousands of files

PathFinder cont'd

Sandia National Laboratories

Adding complexity:

- Noisy graphs, requiring "fuzzy" matches: missing edges, nodes, or entire sequences of nodes and edges.
- Nested, directed, and potentially cyclic graphs.
- Interior CFGs have edges to the exterior graph.
- Interior graphs are fairly localized, but the exterior graphs are not.
- Branching factors and structures of the two layers are different.
- Irregular metadata and signatures.

PathFinder

Sandia National Laboratories

Status

- Map of the MapReduce implemented
  - C language (OO style)
  - OpenMP data parallel
    - Outer loop omp parallel do
  - OpenMP task parallel in progress.
  - Qthreads next.
- Reduce of MapReduce to be added
  - simple MPI_Allgatherv