Parallel Programming Experience with PPM based Gas Dynamics Code

Pei-Hung Lin

University of Minnesota

July 7, 2010

UNIVERSITY OF MINNESOTA

Outline

- Parallel programming model
- I/O model
- Machines scaled
- Performance

Parallel Programming Model

- Hybrid programming model
  - MPI + OpenMP
- MPI
  - Process (rank) management
  - Communication across nodes
- OpenMP
  - Thread management
  - Communication within node
  - Hyper-threading improves performance

Parallel Programming Model – IBM Cell

- Three levels of parallelism:
  - within-Cell
  - within-node
  - node-to-node
- Compute-communication overlap
  - DMA
  - DaCS
  - MPI

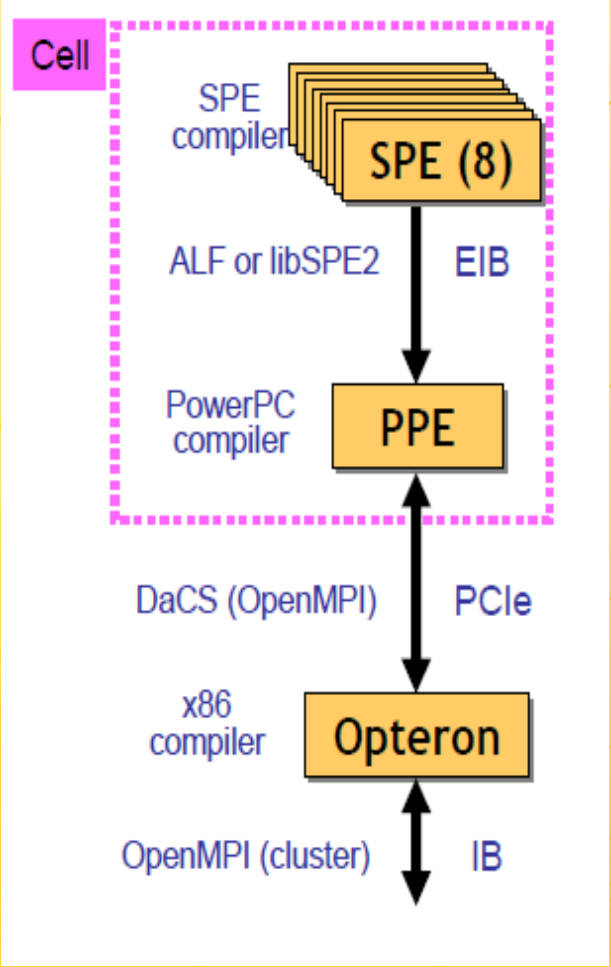

Parallel Programming Model – IBM Cell

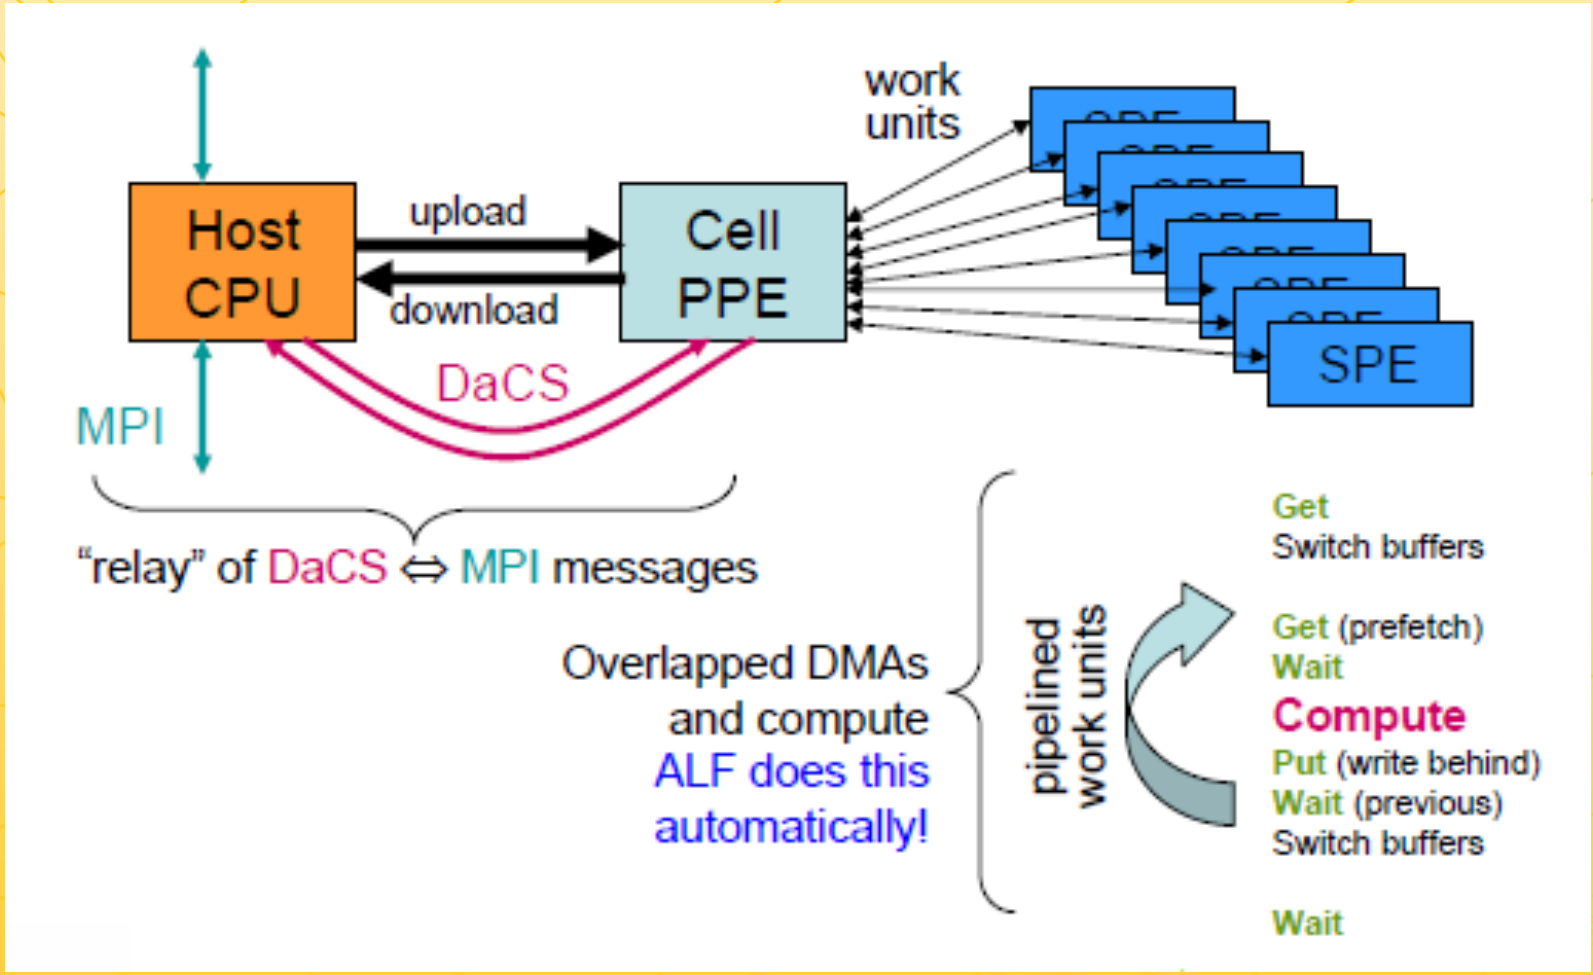

I/O model

- Aggregate file operations improve performance
- **Dedicated** MPI I/O ranks perform file ops
  - Doesn't hold back computation
- Load distribution:
  - Problem domain divided into rectangular blocks
  - MPI ranks working on each block are grouped into teams
  - Each team has one I/O server
  - The server reads/writes restart dumps and visualization data

Machines scaled

- Itasca (HP Linux Cluster) – 64*128+64 = 8256 cores
- Cerrillos (IBM cluster) – 1024+32 = 1056 cell processors = 8448 cores
- RoadRunner – 100, 000 cores (classified)