

Particle-in-Cell Algorithms for Emerging Computer Architectures: GPUs

Viktor K. Decyk and Tajendra V. Singh

UCLA



GPUs are graphical processing units originally developed for graphics and games

- Programmable in 2007

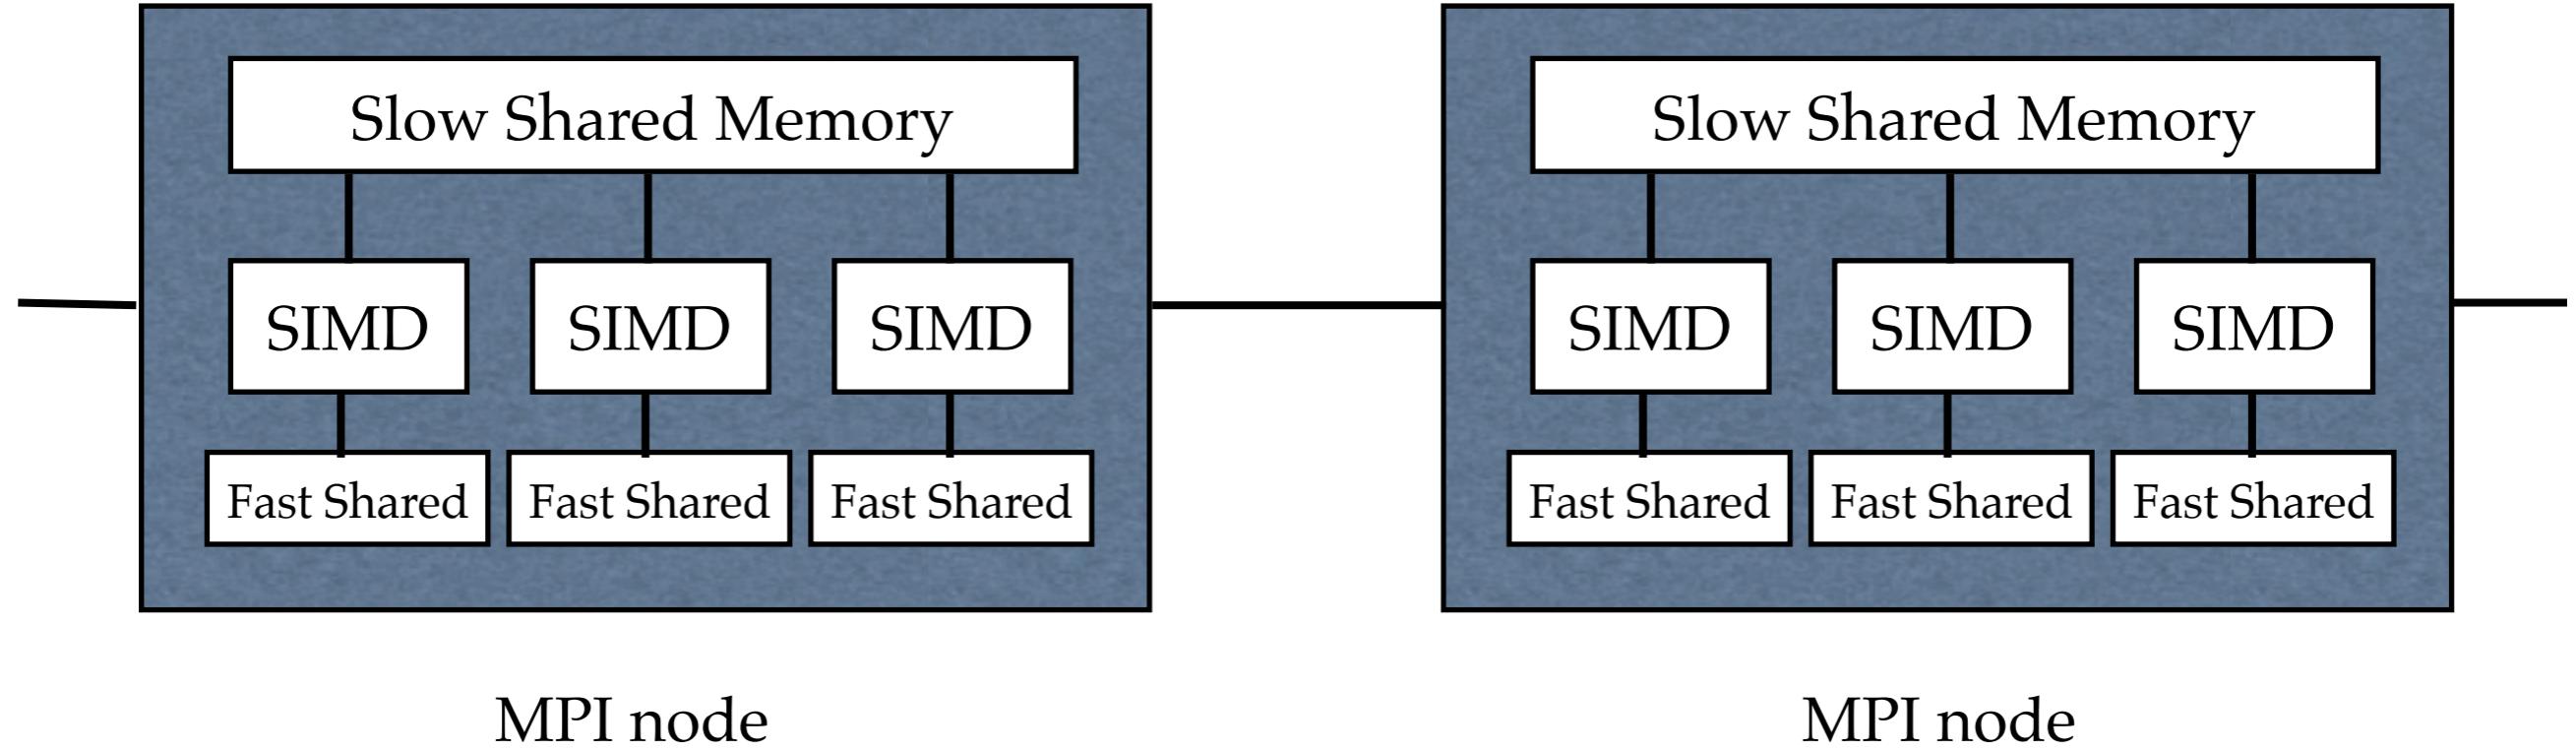
GPUs consist of:

- 12-30 SIMD multiprocessors, each with small (16-48KB), fast (4 clocks) shared memory
- Each multi-processor contains 8-32 processor cores
- Large (0.5-6.0 GB), slow (400-600 clocks) global shared memory, readable by all units
- No cache on some units
- **Very fast (1 clock) hardware thread switching**

GPU Technology has two special features:

- High bandwidth access to global memory (>100 GBytes/sec), but for ordered access
- Ability to handle thousands of threads simultaneously, greatly reducing memory “stalls”

Simple Hardware Abstraction for Next Generation Supercomputer



A distributed memory node consists of

- SIMD (vector) unit works in lockstep with fast shared memory and synchronization
- Multiple SIMD units coupled via global shared memory and synchronization

Distributed Memory nodes coupled via MPI

Each MPI node is a powerful computer by itself

A supercomputer is a hierarchy of such powerful computers

OpenCL programming model uses such an abstraction for a single node

This hardware model matches a variety of processors

On NVIDIA GPU:

- Vector length = block size (typically 32-128)
- Number of vector processors = 12-30
- Fast shared memory = 16-64 KB, plus L2 cache

On Intel MIC (PHI):

- Vector length for single precision = 16
- Number of vector processors = 50-60
- Fast shared memory = L1 Cache (32 KB), L2 Cache (512 KB)

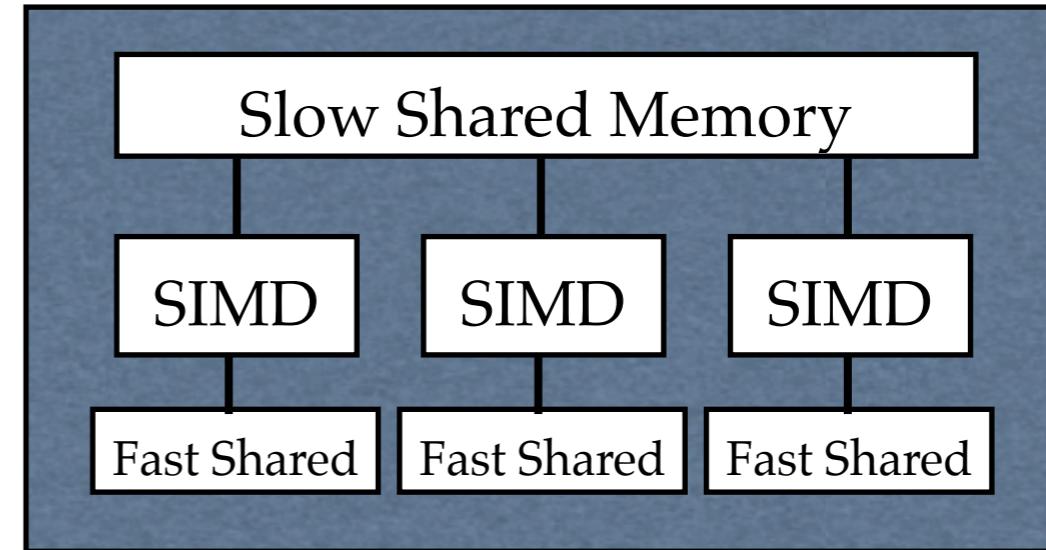
On Dual Intel Xeon multicore:

- Vector length for single precision with SSE2 = 4
- Number of vector processors = 4-16
- Fast shared memory = L1, L2 Cache

Designing new algorithms for next generation computers

This abstract machine contains a number of familiar hardware elements

- SIMD (vector) processors
- Small working memory (caches)
- Distributed memories



Scientific programmers have experience with each of these elements

Vector algorithms

- Calculation on a data set whose elements are independent (can be done in any order)
- Long history on Crays, Fijitsu, NEC supercomputers

Blocking (tiling) algorithms

- When data will be read repeatedly, load into faster memory and calculate in chunks
- Long history on RISC processors

Domain decomposition algorithms

- Partition memory so different threads work on different data
- Long history on distributed memory computers

Designing new algorithms for next generation computers

Programming these new machines uses many familiar elements,
but put together in a new way.

But some features are unfamiliar:

Languages (CUDA, OpenCL, OpenACC, OpenMP) have new features

GPUs require many more threads than physical processors

- Hides memory latency
- Hardware can use master-slave model for automatic load balancing among blocks

Optimizing data movement is very critical

Particle-in-Cell Codes

PIC codes integrate the trajectories of many particles interacting self-consistently via electromagnetic fields. They model plasmas at the most fundamental, microscopic level of classical physics.

PIC codes are used in almost all areas of plasma physics, such as fusion energy research, plasma accelerators, space physics, ion propulsion, plasma processing, and many other areas.

Most complete, but most expensive models. Used when more simple models fail, or to verify the realm of validity of more simple models.

Largest calculations:

- ~3 trillion interacting particles
- ~1.6 million processors (on Sequoia)

Particle-in-Cell Codes

Simplest plasma model is electrostatic:

1. Calculate charge density on a mesh from particles:

$$\rho(\mathbf{x}) = \sum_i q_i S(\mathbf{x} - \mathbf{x}_i)$$

2. Solve Poisson's equation:

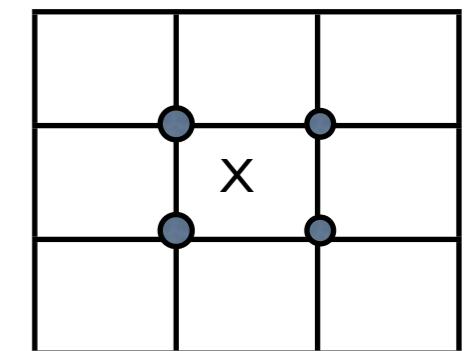
$$\nabla \cdot \mathbf{E} = 4\pi\rho$$

3. Advance particle's co-ordinates using Newton's Law:

$$m_i \frac{d\mathbf{v}_i}{dt} = q_i \int E(\mathbf{x}) S(\mathbf{x}_i - \mathbf{x}) d\mathbf{x} \quad \frac{d\mathbf{x}_i}{dt} = \mathbf{v}_i$$

Inverse interpolation (scatter operation) is used in step 1 to distribute a particle's charge onto nearby locations on a grid.

Interpolation (gather operation) is used in step 3 to approximate the electric field from grids near a particle's location.



When running in parallel, data collisions might occur

GPU Programming for PIC

Most important bottleneck is memory access

- PIC codes have low computational intensity (few flops/memory access)
- Memory access is irregular (gather/scatter)

Memory access can be optimized with a streaming algorithm (global data read only once)

PIC codes can implement a streaming algorithm by keeping particles ordered by tiles.

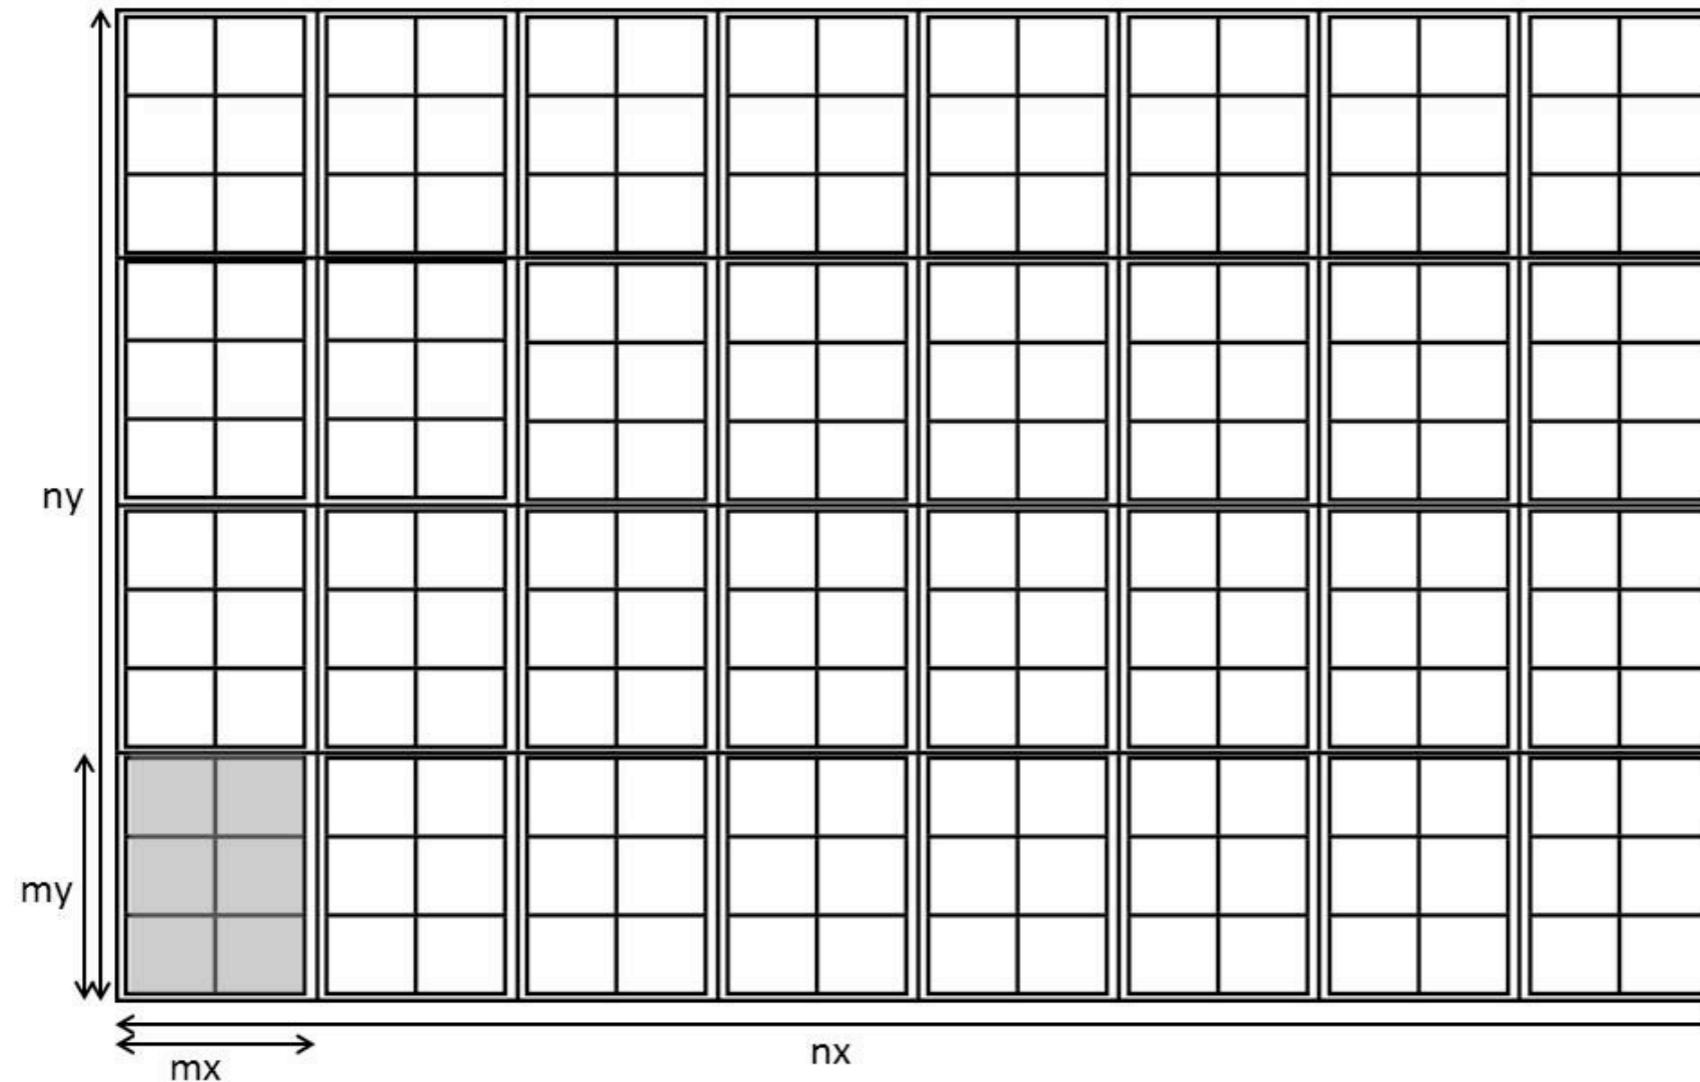
- Minimizes global memory access since field elements need to be read only once.
- global gather/scatter can be avoided.
- Deposit and particles update can have optimal memory access.
- Single precision can be used for particles

Challenge: optimizing particle reordering

Designing New Particle-in-Cell (PIC) Algorithms:
Particles ordered by tiles, varying from 2×2 to 16×16 grid points

We created a new data structure for particles, partitioned among threads blocks:

```
dimension ppart(idimp,npmax,num_tiles)
```



Designing New Particle-in-Cell (PIC) Algorithms: Push/Deposit Procedures:

Within a tile, all particles read or write the same block of fields.

- Before pushing particles, copy fields to fast memory
- After depositing charge to fast memory, write to global memory
- Different tiles can be done in parallel.

Each tile contains data for the grids in the tile, plus guard cells: an extra column of grids on the right, and an extra row of grids on the bottom for linear interpolation.

For push, parallelization is easy, each particle is independent of others, no data hazards

- Similar to MPI code, but with tiny partitions

For deposit, parallelization is also easy if each tile is controlled by one thread

- This avoids data collisions where two threads try to update the same memory

However, if each tile is controlled by a vector of threads, data collisions are possible.

- Atomic updates (which treat an update as an uninterruptible operation) are one approach

Designing New Particle-in-Cell (PIC) Algorithms: Maintaining Particle Order

Three steps:

1. Create a list of particles which are leaving a tile, and where they are going
2. Using list, each thread places outgoing particles into an ordered buffer it controls
3. Using lists, each tile copies incoming particles from buffers into particle array

- Less than a full sort, low overhead if particles already in correct tile
- Can be done in parallel
- Essentially message-passing, except buffer contains multiple destinations

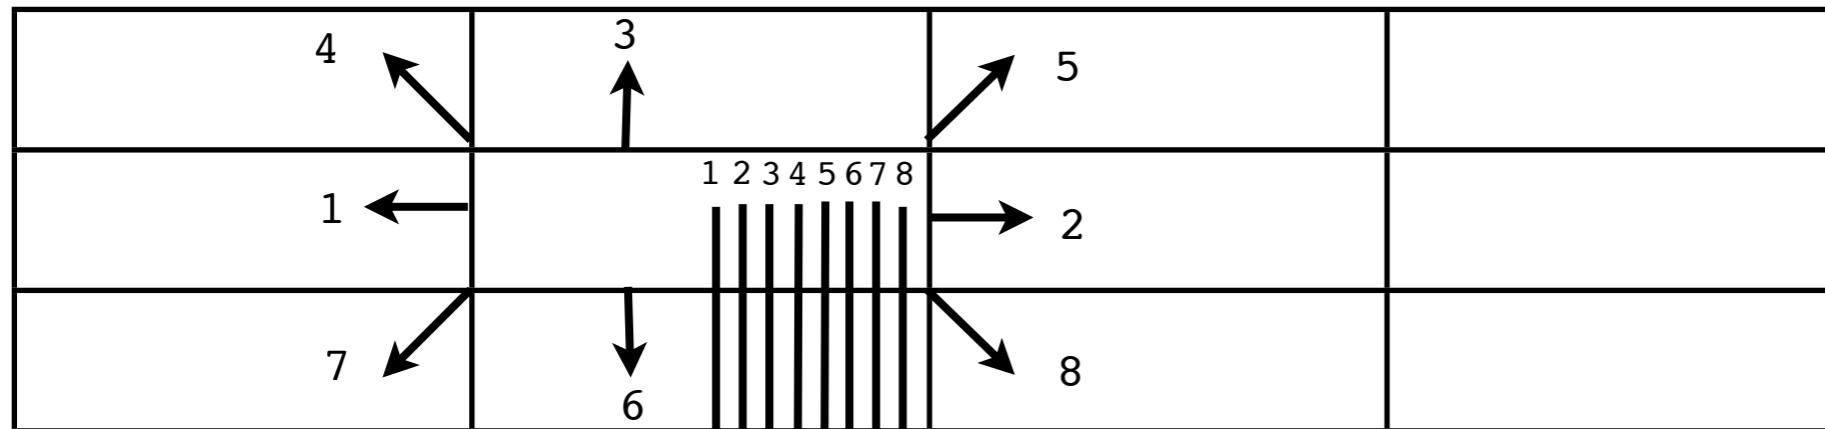
In the end, the particle array belonging to a tile has no gaps

- Incoming particles are moved to any existing holes created by departing particles
- If holes still remain, they are filled with particles from the end of the array

Straightforward to implement with one thread per tile

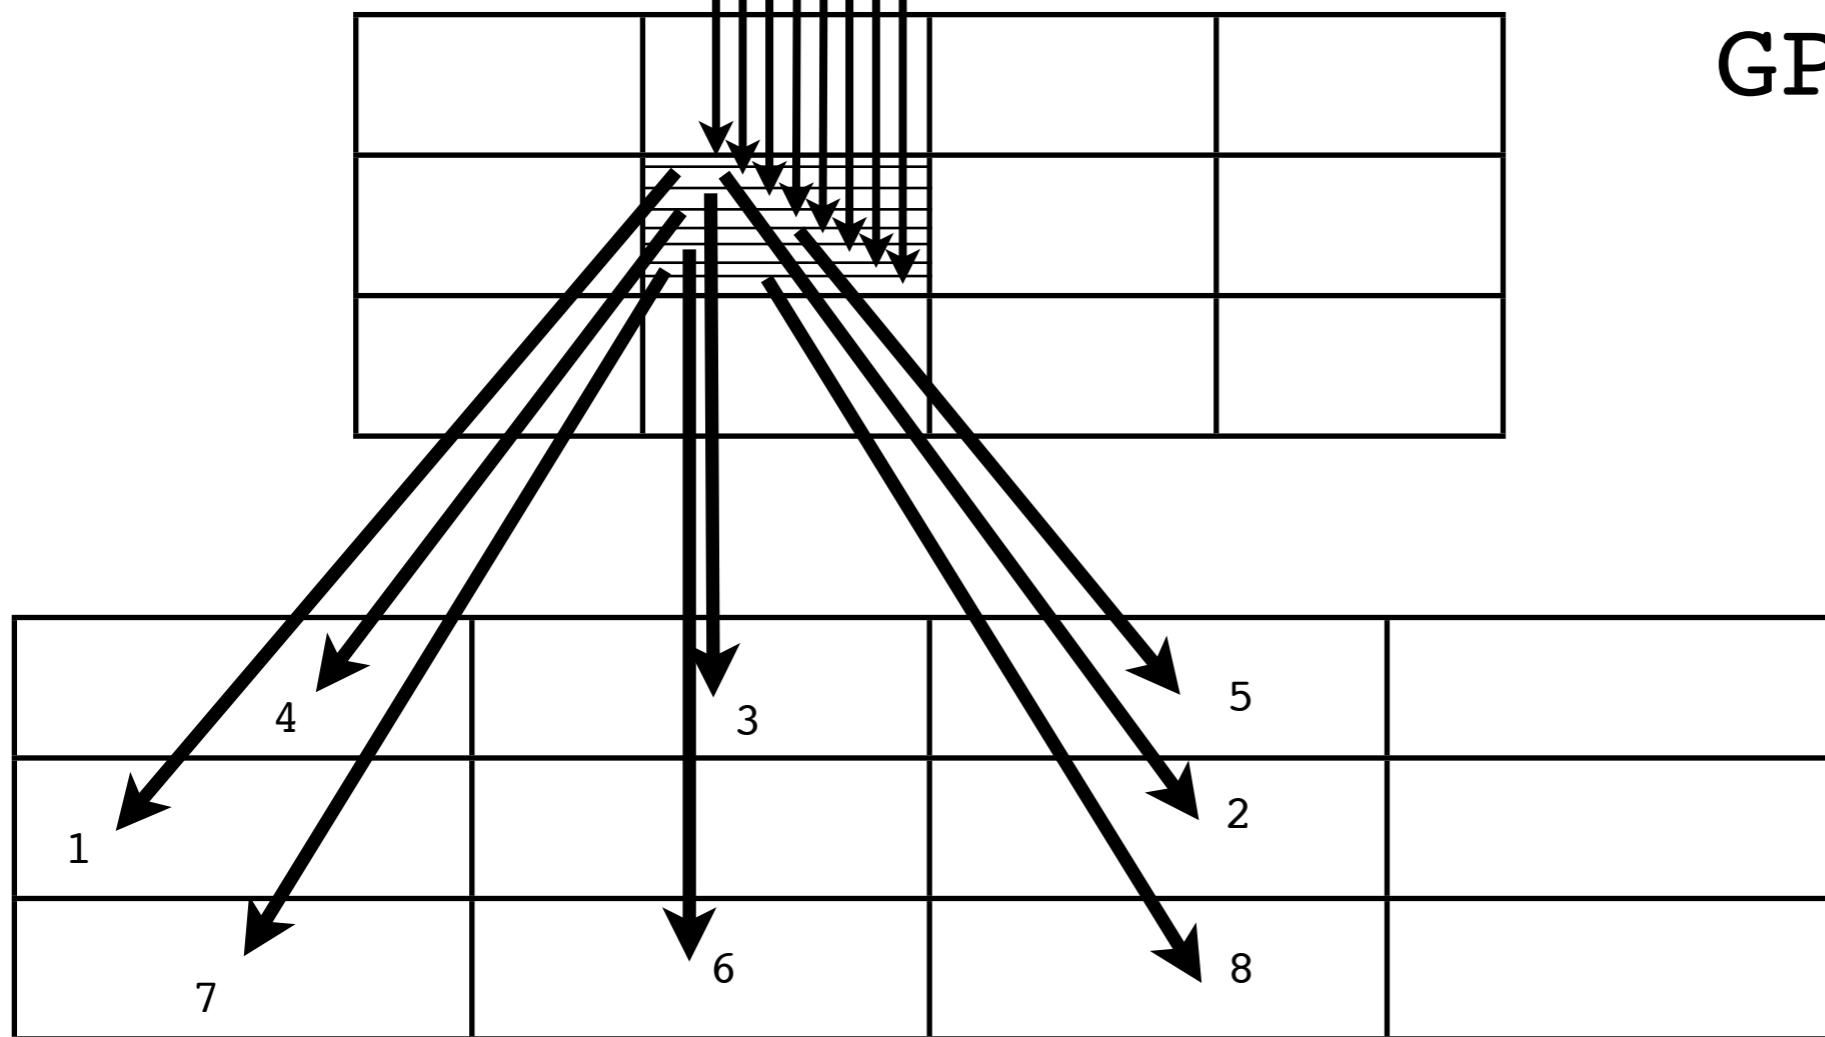
- Much more complex with multiple threads per tile

GPU Particle Reordering



GPU Tiles

Particles buffered
in Direction Order



GPU Buffer

GPU Tiles

Evaluating New Particle-in-Cell (PIC) Algorithms on GPU: **Electrostatic Case**
2D ES Benchmark with 2048x2048 grid, 150,994,944 particles, 36 particles / cell
optimal block size = 128, optimal tile size = 16x16. Single precision

GPU algorithm also implemented in OpenMP

Hot Plasma results with dt = 0.1		
	CPU:Intel i7	GPU:Fermi M2090
Push	20.7 ns.	0.552 ns.
Deposit	9.4 ns.	0.247 ns.
Reorder	1.6 ns.	0.114 ns.
Total Particle	31.7 ns.	0.914 ns.

The time reported is per particle/time step.

The total particle speedup on the Fermi M2090 was 35x compared to 1 CPU.

Field solver takes an additional 12% on GPU, 11% on CPU.

Evaluating New Particle-in-Cell (PIC) Algorithms on GPU: **Electromagnetic Case**
2-1/2D EM Benchmark with 2048x2048 grid, 150,994,944 particles, 36 particles / cell
optimal block size = 128, optimal tile size = 16x16. Single precision

GPU algorithm also implemented in OpenMP

	CPU:Intel i7	GPU:Fermi M2090
Push	68.4 ns.	0.980 ns.
Deposit	43.2 ns.	1.327 ns.
Reorder	0.6 ns.	0.068 ns.
Total Particle	112.2 ns.	2.375 ns.

The time reported is per particle/time step.

The total particle speedup on the Fermi M2090 was 47x compared to 1 CPU.

Field solver takes an additional 16% on GPU, 11% on CPU.

Conclusions

PIC Algorithms on emerging architectures are largely a combination of previous techniques

- Vector techniques from Cray
- Blocking techniques from cache-based architectures
- Message-passing techniques from distributed memory architectures
- Programming to Hardware Abstraction leads to common algorithms
- Streaming algorithms optimal for memory-intensive applications

Scheme should be portable to architectures with similar hardware abstractions

Further information available at:

V. K. Decyk, and T. V. Singh, “Particle-in-Cell Algorithms for Emerging Computer Architectures,” Computer Physics Communications, 2013.

<http://www.idre.ucla.edu/hpc/research/>