CMSC 714
Lecture 23
Molecular Dynamics

Alan Sussman

Notes

- Graded exams returned next Tuesday
- Group Project presentations scheduled for Dec. 6, and Dec. 8 (if needed)
  - final report due Monday, December 12

Molecular Dynamics

- Calculate trajectories of atoms and molecules by solving Newton's equations of motions
- Force calculations
  - Bonded interactions: bonds, angles, dihedrals
  - Non-bonded interactions: van der Waal's and electrostatic forces
- Number of atoms: thousands to millions
- Simulation step: ~1 femtosecond (10^{-15} sec)

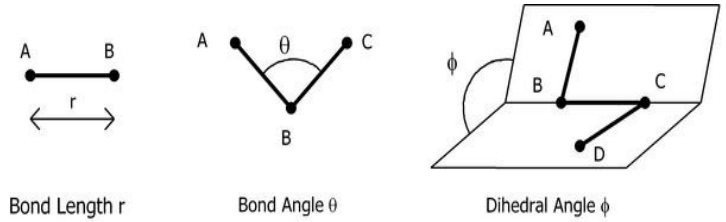

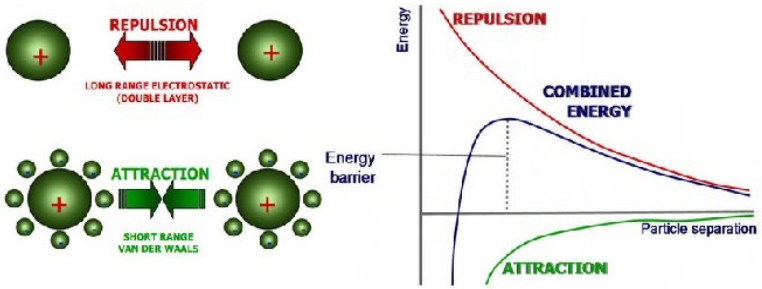

Sequential Algorithm

- At every step, calculate forces on each atom
  - Calculate bonded and short-range forces every step
  - Calculate long-range non-bonded forces every few time steps (using PME or P3M etc.)
- Particle mesh Ewald (PME) summation:
  - Calculate long-range interactions in Fourier space
- Calculate velocities and new positions
- Repeat …

Traditional approaches to parallelization

- Atom decomposition:
  - Partition the atoms across processes
- Force decomposition:
  - Distribute the force matrix to processes
  - Matrix is sparse and non-uniform
- Spatial decomposition:
  - Assign a region of the 3D simulation space to each process

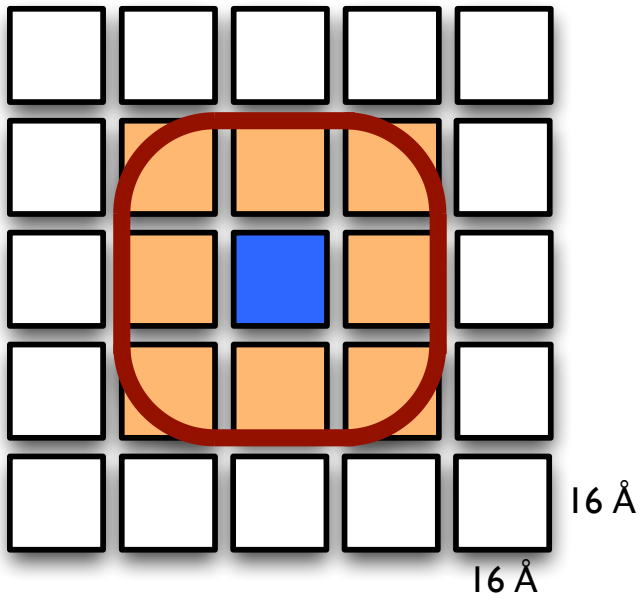

Hybrid parallelization

- Hybrid of spatial and force decomposition
- Decouple assignment of data and work to processes
- Distribute both atoms and the force calculations to different processes

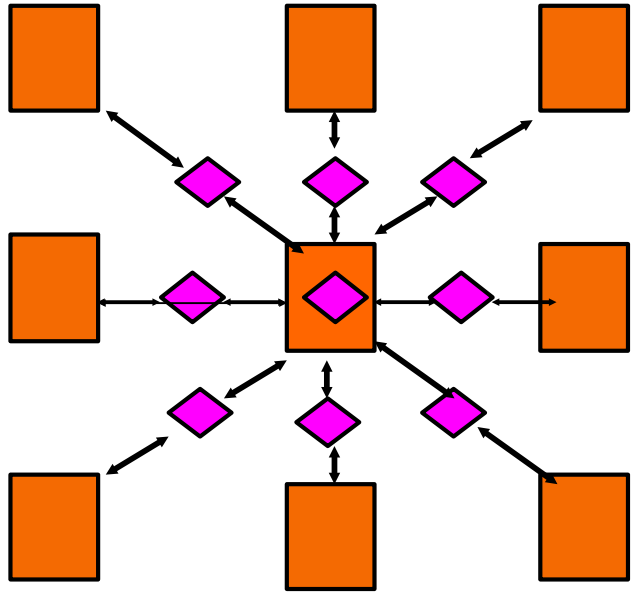

Neutral territory (NT) methods

- Desmond's mid-point method

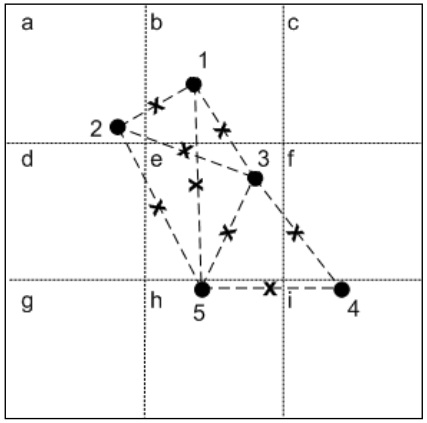

Midpoint method

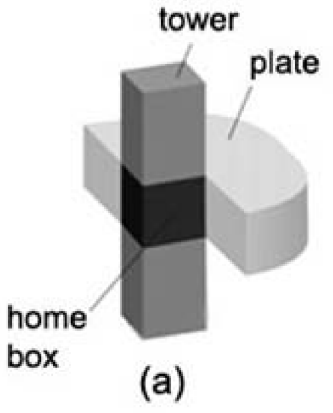

NT method

Particle mesh Ewald method

- Replace direct force calculations by:
  - Calculate short-range forces in real space
  - Calculate long-range forces in Fourier space
- Create a 3D mesh/grid representing charge densities of atoms
  - Compute a 3D Fast Fourier Transform (FFT)
- FFT computes the discrete Fourier transform (DFT) or inverse DFT
  - Reduces the complexity from $O(N^2)$ to $O(N \log N)$

Parallelization of PME (3D FFT)

- 1D or slab decomposition

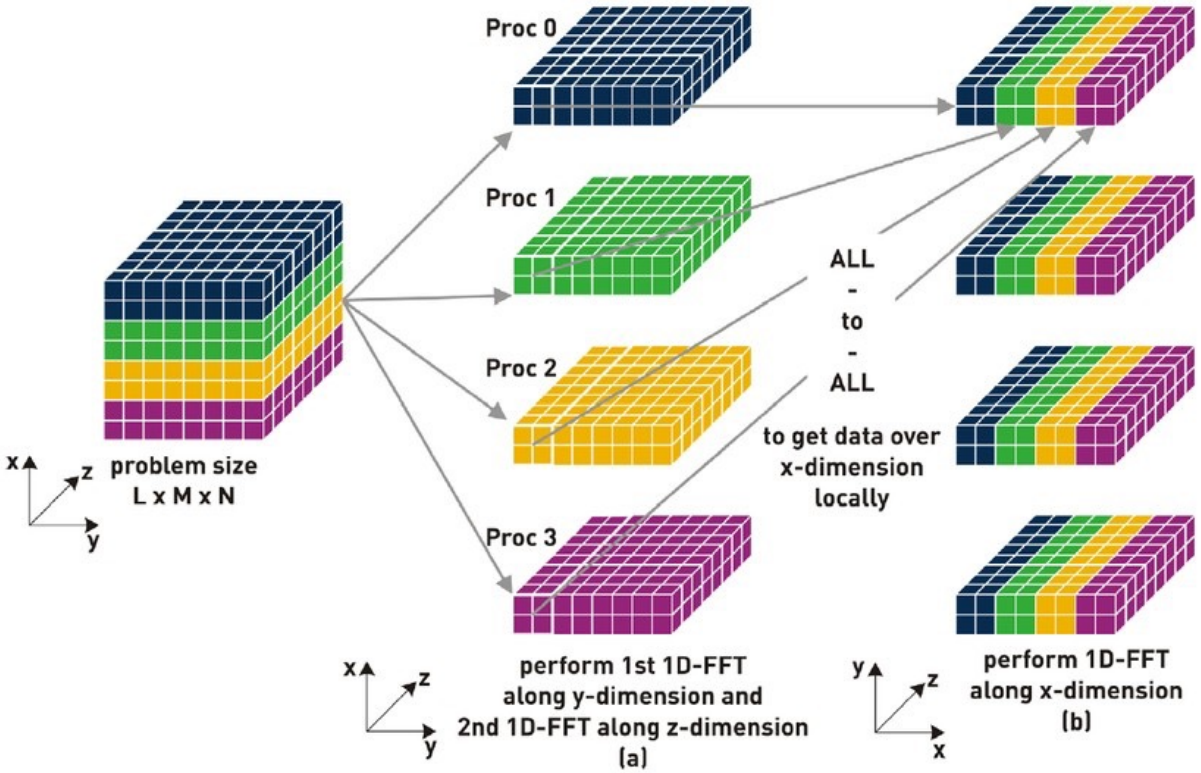

Parallelization of PME (3D FFT)

- 2D or pencil decomposition

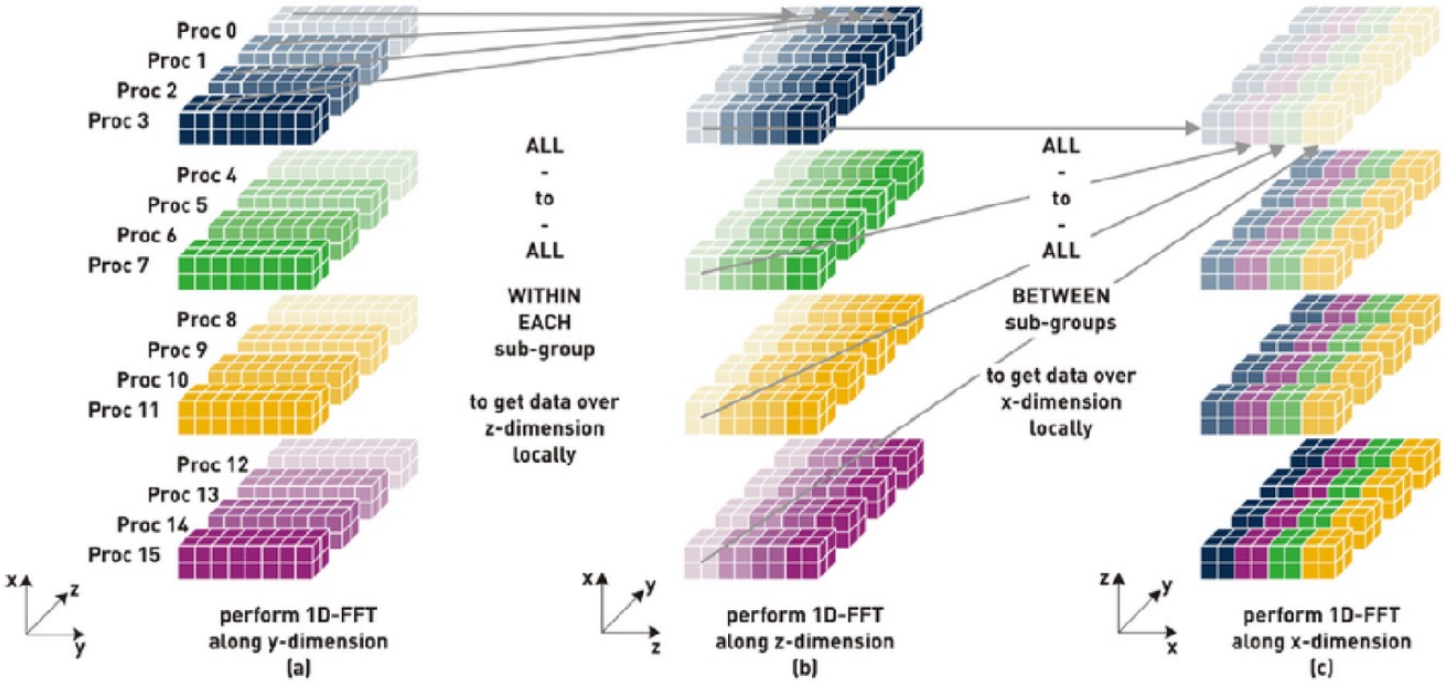

NAMD

- MD code for large biomolecular systems
  - Implemented in parallel using CHARM++
- Targets large physical systems, for a large number of time steps
  - Tens to hundreds of thousands of molecules/atoms
  - Simulating maybe ten nanoseconds (10^{-9}) of real time, at one femtosecond (10^{-15}) per time step
- CHARM++
  - C++ language/programming model that provides object-based message-driven execution model
  - Load balances objects across nodes/processors dynamically, and allows users to develop their own load balancers – objects can migrate across nodes at runtime
  - Objects interact via asynchronous method invocation
  - Users have to decide on granularity of objects, to balance complexity, locality, load balancing, etc.

NAMD

- To get sufficient parallelism, computation decomposed both spatially (into cubes with size determined by cutoff radius) and by forces (with separate non-bonded force computation objects)
- Methods needed to get good load balancing are quite complex, and bind some objects to nodes while others can migrate
  - Initial load balancing uses recursive coordinate bisection
  - Dynamic load balancing is measurement-based to refine the initial (and subsequent) assignments
  - Uses proxy (ghost) objects as needed to provide data for local computations that need data from objects owned by other nodes/processors
  - Uses PME for electrostatic computations, so needs efficient 3D FFTs using CHARM++, via modified FFTW routines
- Extensive performance results show good weak scaling up to 512 processors for large molecular systems
  - But not clear how the methods would scale on modern multicore systems since best results don't use all processors on each node (3 is better than 4 on PSC Lemieux)
  - They also show better performance directly using Elan message passing library for Quadrics network compared to MPI

Anton

- Special purpose machine built for molecular dynamics simulations, built at D.E. Shaw Research
  - to simulate biological processes that occur on very small time scales (10^{-15} sec), such as protein folding, interaction between proteins, etc.
  - and simulate those processes for a long time
- Molecular dynamics
  - force calculation followed by integration step to move particles
  - biomolecular forces have 3 parts
    - bonded forces – small atomic groups with covalent bonds
    - van der Waals forces – all pairs of atoms, but fall off quickly with distance (so only need close ones)
    - electrostatic forces – all pairs of atoms, fall off slowly with distance – divide into 2 parts to avoid all pairs computation

Anton

- Anton machine
  - up to 512 nodes in 8x8x8 torus
  - each node has 2 parts on 1 chip
    - high throughput interaction subsystem (HTIS) for range-limited interactions , using 32 hardwired pairwise point interaction pipelines (PPIPs)
    - flexible subsystem with 8 programmable geometry cores (GCs) for less structured part of MD computation, 4 Tensilica processors, 4 data transfer engines
    - plus DRAM controllers, 6 network interfaces, and host interface for I/O
- Most of computational time mapped to PPIPs, which run those computations maybe 100x faster than standard microprocessor core
- And computations spatially decomposed across nodes, with some twists to deal with communication as particles move between spatial domains
- Uses fixed-point arithmetic, with various bit widths, for several reasons:
  - performance – fixed-point hardware fast and small
  - fixed point arithmetic is truly associative
  - gain determinism – run same simulation again get exact same results bit-for-bit (doesn't really help, since MD is a chaotic system, so need ensemble)
  - computations are reversible

Anton

- Performance results show can run a large chemical system at much higher rates than any previous system
  - can run multiple microseconds of simulation time per day of wall clock time
    - maybe 500 times faster than 512 node Intel Xeon cluster
    - and have run simulated systems up to over 1000 microseconds, which showed interesting behavior of the molecules
  - and results are validated very well
    - both against “known” results and using statistical error tests