GROMACS

Roland Schulz, PI: Jeremy C. Smith

2011

Biomass to Cellulosic Ethanol

Plant Cell Wall

- cellulose
- lignin
- matrix polysaccharides
  - hemicellulose
  - pectin
- proteins
- complex structure

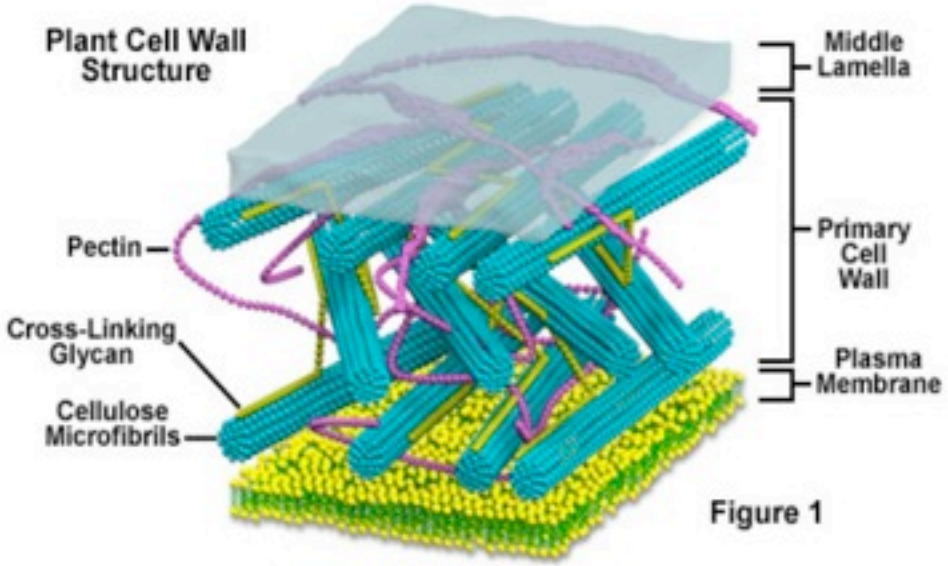

Figure 1

Feedstock → ***Pretreatment*** → ***Enzymatic Hydrolysis*** → ***Fermentation***

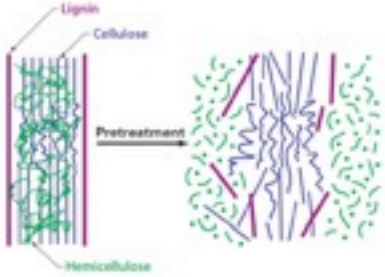

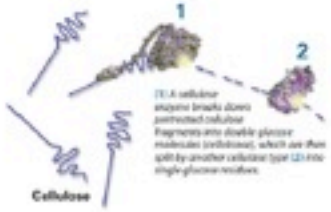

Molecular Dynamics

- Parallelization
  - Domain Decomposition
  - Load Balancing
  - Task Parallelism
- Non-Bonded Forces

  >90% of Time
- Short-Range: Compute Intensive
- Neighbor-List: Random Memory Access
- Long-Range: Communication Intensive

GROMACS

- GPL, GIT, Eclipse PTP
- ~500 citations/yr, 5k-10k users (also industry)
- ~500k SLOC, X86(SSE/SSE2/SSE3/… 32/64bit), PPC (Bluegene, Power6, Altivec), NVidia

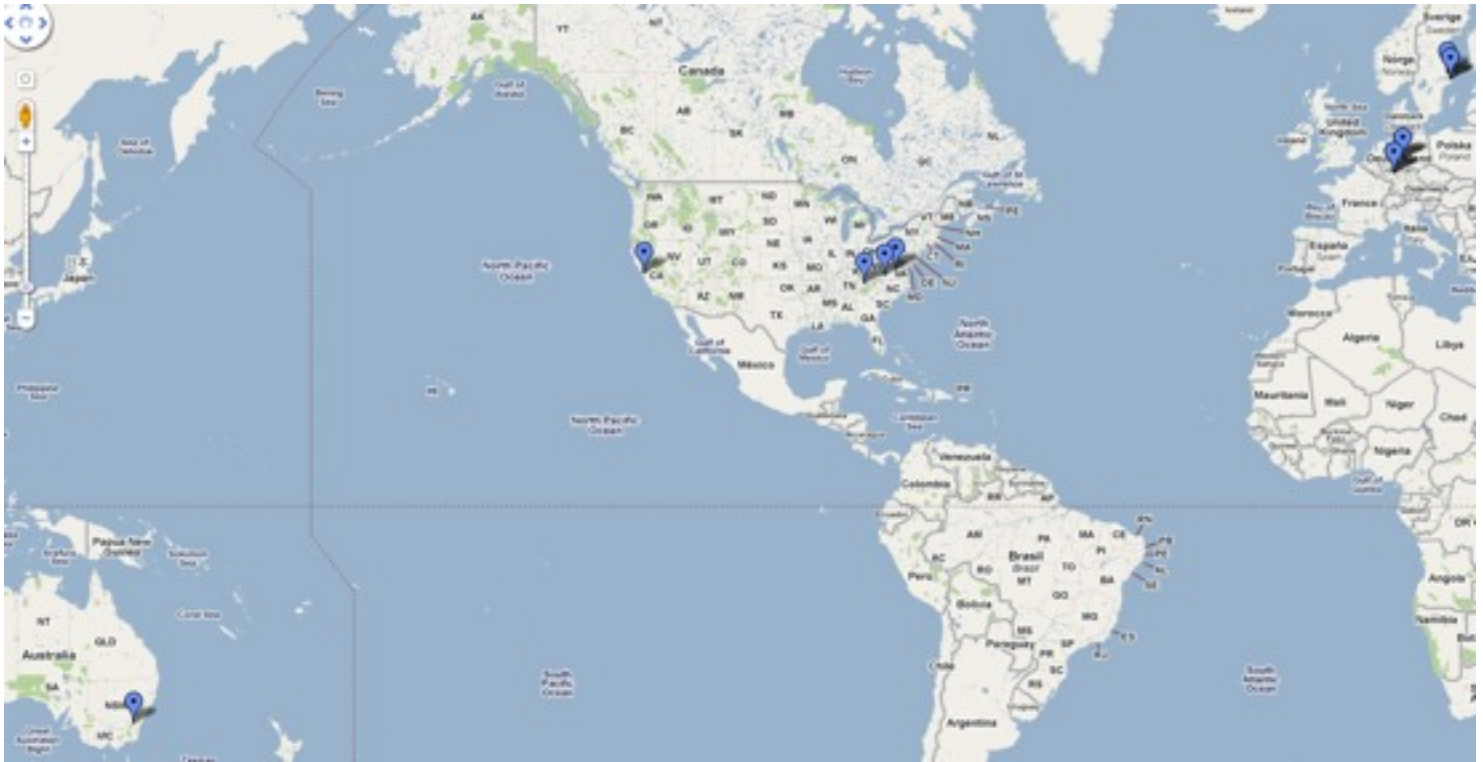

Programming Model

- MPI + OpenMP
- Optional:
  - MPI implementation using Pthreads
  - GPU (CUDA)
- C – moving to C++
- Cray, BlueGene, Cluster, PCs, other (e.g. PlayStation), Windows (VC)

Single Node Performance

- Assembler/Intrinsic SSE (significant faster than SPEC)
- single precision (optional)
  - wider SSE, number of steps for 1/sqrt(r^2)), exp, …
- Virial and pbc outside inner loop
- Own SSE transcendental functions
  - Similar performance on all compilers
  - Improves Intel SVML by up to 20%
- For water
  - unrolled loop
  - special neighbor-list
  - LJ for 1 atom/molecule

I/O

- usually <1MB/s (average, compressed)
- File size up to ~1TB
- Compression and communication important
- Using MPI-IO with custom sorting/compression

Visualization + Analysis

- VMD (Molecular Viewer)
- GROMACS Analysis Tools

Performance

- Main: CrayPat (easy to use)
- Also: Tau, HPCToolkit, MPIP, MPE

Debugging

- DDT (easier to use)
- Sometimes Totalview, Eclipse PTP (free)

Other Tools

- Eclipse PTP
  - Remote Development
  - Debugger
  - Performance Tools (e.g. Tau)

Scaling PME

- PME benefits from task parallelism / MPMD
  - Smaller AlltoAll communicator, tuning problem

Reaction Field

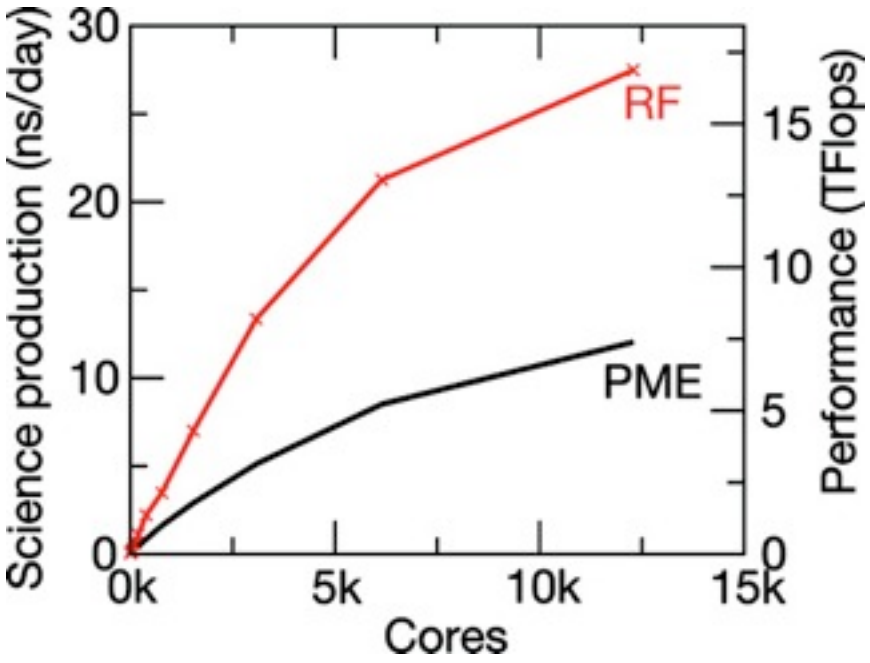

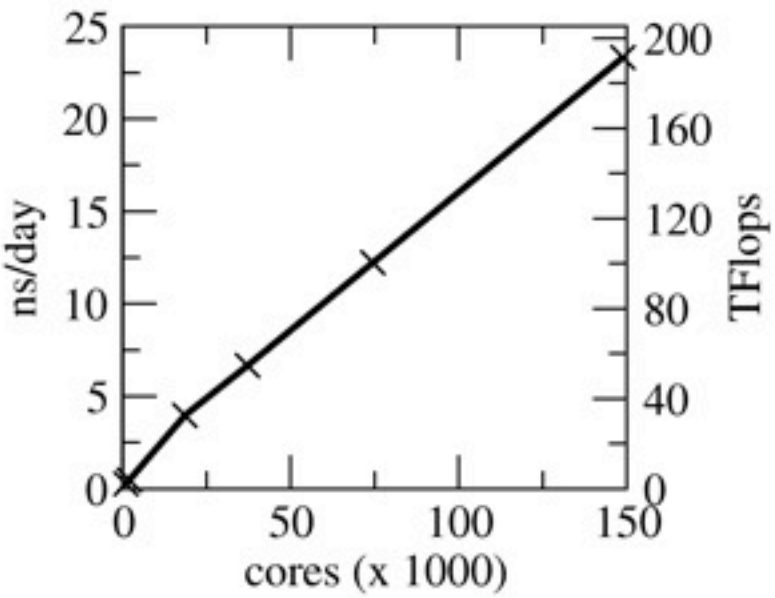

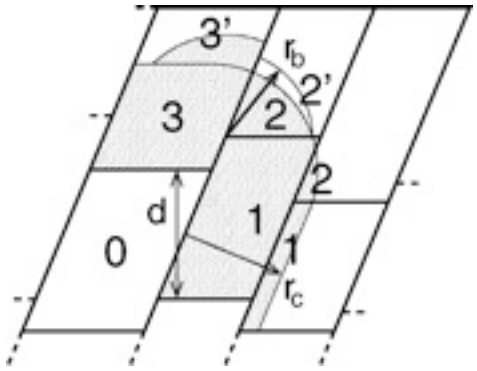

- Load balancing critical
- Water – solute difference
- Imbalance from 200% to 75%
- 44% speed improvement

Plans

- GPU+MPI
- Better Loadbalance
- Task Parallelism

Pains

- Strong Scaling is Difficult
- MPI (not always optimal, latency & message rate)
- User-friendliness of Tools
- Large range of Architectures/Topologies
- Network unpredictable on Cray

Roadmap

- Cell wall model containing Cellulose,Lignin,Hemicellulose
- Algorithm improvements: Multigrid

Thanks

- Center for Molecular Biophysics
- E. Lindahl, B. Hess, D. van-der-Spoel, et al.
- Genome Science & Technology
- NCCS, Incite