Scaling of Multimillion-Atom MD Simulation on a Petascale Supercomputer

Roland Schulz[1], Benjamin Lindner[1], Loukas Petridis[1], Berk Hess[2], Erik Lindahl[2] and Jeremy C. Smith[1]
Center for Molecular Biophysics, Oak Ridge National Laboratory; Royal Institute of Technology, Stockholm

Abstract: This poster shows performance and scaling for all-atom molecular dynamics (MD) simulation of multimillion-atom biological systems. In the era of petaflop supercomputers, such simulation is limited by the parallel efficiency of the MD algorithms. Highly tuned SSE compute kernels enable node performance similar to GPU performance. The bottleneck for highly-parallel all-atom simulations is the computation of the electrostatic interactions. Shown is the scaling of MD using the reaction field method to the full Jaguar XT5 size. Also presented is a new highly scalable Particle Mesh Ewald decomposition.

Atomistic Cellulose Lignin Aggregation

Efficient decomposition of lignocellulosic biomass is key to production of 2nd generation bio-ethanol

Atomic detailed model of cellulose and lignin

Reaction field (RF) for electrostatic, GROMACS 4.5

Good agreement between RF and PME

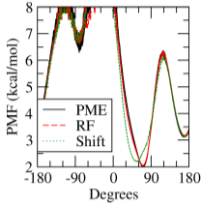

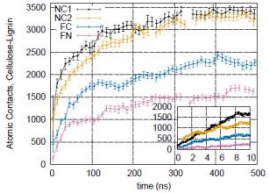

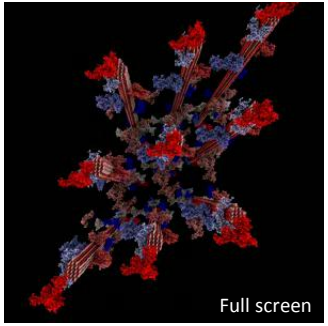

Aggregation after 250ns.
Each lignin clustere is shown in a different color.

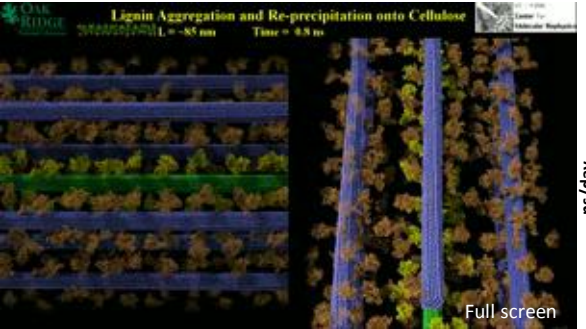

Performance: Reaction Field

- 3.3M atoms
- Ligno-cellulose
- Load-imbalance improved from 200%->75%, 44% speed improvement

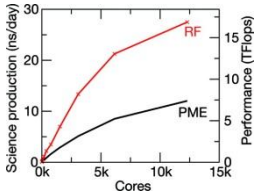

- 100M atoms
- Pure Water: >300TF

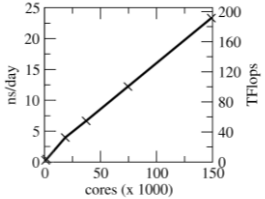

Parallel IO: 3M atoms, 8192 cores
Buffering frames, Parallel sort +compress
21ns/day improves to 34ns/day

Performance: PME

- 2.2 M atoms, MPMD
- Comparison 1.1M atoms NAMD 9.7ns/day on 4048 cores XT5 (data: teragrid)

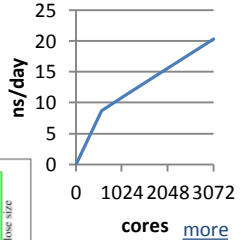

more

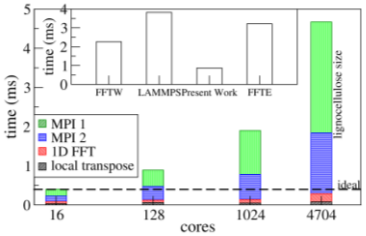

AllToAllV **much** slower
Requires padding
Available as library
Single + real important
Task placement

Strong Scaling 128³ (ms)

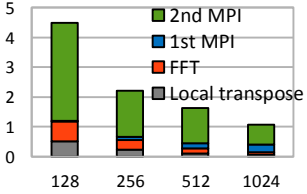

FFT-Improvements obtained using MPI+OpenMP versus only OpenMP

data	Nodes x cores	Using Threads (ms)	No Threads (ms)
64^3	16x16 Opteron	1.0	14.0
64^3	16x8 XEON,QDR	0.32	0.77
648x 128^2	180 x 12 Jaguar	4.1	11.8

GPU/CPU performance

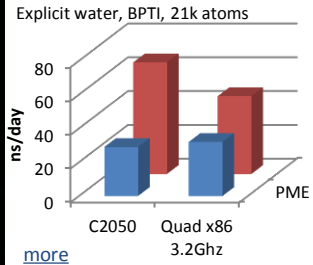

more

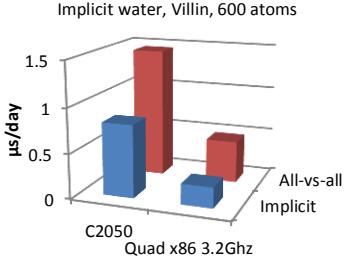

Conclusion

With RF >20ns for 100M atoms, much improved scaling with PME, fast 3D FFT library, PME using threads, GPU performance much dependent on system size and electrostatic method, fair CPU comparison important

Acknowledgements: David van der Spoel, Rossen Apostolov, Szilard Pall, Peter Eastman, Vijay Pande, GST, NCCS, Incite, FWP ERKP704.