

The Fast Multipole Method in molecular dynamics

KTH Royal Institute of Technology, Stockholm, Sweden

ADAC6 workshop Zurich, 20-06-2018



Berk Hess

Funding



Horizon 2020







Key Software

Key Workflows and Platforms

Partners





BioExcel





Molecular Dynamics of biomolecules

Powerful because:

structure and dynamics in atomistic detail

Challenging because:

- fixed system size: strong scaling needed
- always more simulation needed than possible
- force-field (model) accuracy can be limiting
- results can be difficult to analyse







Molecular Dynamics



Simple (symplectic) integrator

Fundamental issue: sequential problem need **10⁹ - 10¹⁵** steps! $\Delta t \sim 1$ fs, timescales of interest μs - s

 $v_i(t +$

$$\frac{\mathbf{x}_i}{2} = \mathbf{F}_i = -\nabla_i V(\mathbf{x}) , \quad i = 1, \dots, N$$

$$\frac{\Delta t}{2}) = v_i(t - \frac{\Delta t}{2}) + a_i(t)\Delta t$$

$$x_i(t + \Delta t) = x(t) + v_i(t + \frac{\Delta t}{2})$$



Computational cost mostly in computing the forces





The MD strong scaling challenge

single node):

- at ~ 200 atoms per CPU core ⇒ we run in L1/L2 cache
- at ~ 3000 atoms per GPU
- but systems of interest are usually > 100000 atoms

Extreme demands on hardware and parallelisation libraries (MPI, OpenMP): ideally overheads should be < 1 microsecond

- To reach 10⁹ 10¹⁵ steps we need to minimise the walltime per step
- The best MD codes in the world are at **100 microseconds** per step (on a



GROMACS

- Good parallel scaling
- Compiles and runs efficiently on nearly any hardware
- Open source, LGPL
- C++, MPI, OpenMP, CUDA, OpenCL
- C++ and Python API in the works

 Started as hardware+software project in Groningen (NL) Now core team in Stockholm and many contributors world-wide Focus on high performance: efficient algorithms, SIMD & GPUs

Incorporate important algorithms for flexibly manipulating systems



Highly optimised vectorization:

Small C++ SIMD layer with highly optimised math functions for: SSE2/4, AVX2-128, AVX(2)-256, AVX-512, ARM, VSX, HPC-ACE

CPU kernels reach 50% of peak

CUDA and OpenCL need separate kernels







Atom clustering (regularization) and pair list buffering



We can optimize the size of this buffer

Larger buffer means more calculations, but we can update the neighbor list less frequently



Atom clustering and pair list buffering



We can optimize the size of this buffer

Larger buffer means more calculations, but we can update the neighbor list less frequently





Intra-rank parallelisation: OpenMP

GROMACS has efficient parallelization of all algorithms using MPI + OpenMP

OpenMP is (performance) portable, but limited:

- No way to run parallel tasks next to each other
- No binding of threads to cores (cache locality)

Need for a better threading model, requirements: • Extremely low overhead barriers (all-all, all-1, 1-all)

- Binding of threads to cores
- Portable



Spatial decomposition

8th-shell domain decomposition:

- minimizes communicated volume
- minimizes communication calls
- implemented using aggregation of data along dimensions
- 1 MPI rank per domain
- dynamic load balancing

Question: still the best approach for networks that support many messages?

Center of Excellence for Computational Biomolecular Research



Hess, Kutzner, Van der Spoel, Lindahl; J. Chem. Theory Comput. 4, 435 (2008)







New in GROMACS-2018: PME on GPU!





GROMACS strong scaling

Lignocellulose benchmark

- 3.3 million atoms
- No long-range electrostatics!
- machine: Piz Daint at CSCS
 - Cray XC50
 - NVIDIA P100 GPUs



Center of Excellence for Computational Biomolecular Research





Strong scaling issues

At 100s of microseconds per step:

- The 3D-FFT in PME electrostatics
- MPI overhead
 - we need MPI_Put_notify()
- OpenMP barriers take significant time
- Load imbalance
- CUDA API overhead can be 50% of the CPU time
- Too many GROMACS options to tweak manually

140000 atom ion channel on Piz Daint xc





Long-range electrostatics

All examples up till now without long-range electrostatics, but needed in most cases We need all-vs-all Coulomb interactions

All MD codes use Particle-mesh Ewald (PME) for electrostatics PME uses a 3D-FFT, 64³ - 128³ grid points Issues:

- The 3D-FFT requires two grid transposes that require All-to-All communication in slabs
- Pencil decomposition clashes with 3D decomposition
- 3D-FFT scaling is the bottleneck for MD

GROMACS Multiple-Program Multiple-Data parallelisation improves scaling by a factor 4

8 PP/PME ranks



3D-FFT communication







Long-range electrostatics methods for molecular dynamics

	in use in MD	computational cost	pre-factor	communication cost
Ewald summation	since 70s	O(N ^{3/2})	small (FFT)	high
PPPM / Particle Mesh Ewald	since 90s	O(N log N)	small (FFT)	high
Fast Multipole Method	not yet	O(N)	larger	low
Multigrid Electrostatics	2010	O(N)	larger	low

we recently solved the energy conservation issue



information moves from red to blue

FMM scheme

M2M multipole to multipole treecode & FMM

P2M particle to multipole treecode & FMM

Rio Yokota Tokyo Tech

source particles





The accuracy of FMM is controlled by

- the order of the expansion(s)
- the theta parameter

Main issue for FMM in MD:

- We want low accuracy (or better: high speed)
- We want energy conservation

Center of Excellence for Computational Biomolecular Research

θ=0.50





*θ***=0.35**



θ=0.20







work by **Rio Yokota**







Main physics issue with FMM in MD:

- All atoms have partial charges, but molecules are mostly locally neutral
- The sharp boundaries of FMM cells introduce charges at the boundaries



a neutral water molecule crossing 2 FMM cells

at long distance water is a dipole





FMM regularization: add smooth overlap between cells





orig. idea: Chartier, Darrigand Faou, BIT Num. Math. 50, 23 (2010)

$$V = s_1(x)V_1(x) + s_2(x)V_2(x)$$

$$F = -\nabla V$$

$$= s_1(x)F_1(x) + s_2(x)F_2(x)$$

$$-s'_1(x)V_1(x) - s'_2(x)V_2(x)$$



- Davoud Saffar Shamshir Gar has analyzed the accuracy of regularization
- Currently: 2D FFM with full regularization
- Results for charge pairs with $\theta = 0.35 =$ second nearest neighbors, 8 atoms per cell











Center of Excellence for Computational Biomolecular Research



Computational cost of regularization

- At 0.473 nm (10.6 atoms) cell free with LJ
- 1/2 regularization width = 1/4 cell Second nearest neighbor direct • Direct pair cost factor $(2.25/2)^3 = 1.42$
- Multipole lowest level: factor $1.5^3 = 3.4$
- But we get factor 5 more accuracy with P=4 And we get energy conservation!



Conclusions

- PME 3D-FFT is limiting scaling

- FMM regularization improves accuracy

Is there a better threading/tasking model than OpenMP on the horizon?

 FMM has better communication characteristics FMM is a good fit for SIMD/GPU acceleration FMM regularization gives energy conservation FMM fits wells with GROMACS pair computation

We are implementing optimized, regularised FMM





Acknowledgements

Erik Lindahl Mark Abraham Szilard Pall Aleksei lupinov



Rio Yokota, Tokyo Tech John Eblen, ORNL **Roland Schulz**, Intel Mark Berger, NVIDIA rest of the GROMACS team world-wide





Horizon 2020 European Union Funding

