

Innovative HPC Methods Highly Scalable Molecular

and Application to Simulation (IMEMO)

Molecular simulations are an important tool for many research areas like chemical engineering, biochemistry, material sciences and nanotechnology. Despite its huge potential, the user community of this technology is relatively small. The main cause for this is the need to bring powerful computational resources to bear, in order to simulate sufficiently long times and apply complex models, necessary to make physically meaningful predictions. This poses great challenges to software. The national research project Innovative HPC Methods and Application to Highly Scalable Molecular Simulation (IMEMO) aimed exactly at these challenges and the opportunities offered by next generation

HPC systems for molecular simulations. The project was funded by the BMBF in the program "IKT 2020 - Forschung für Innovationen" category "HPC-Software für skalierbare Parallelrechner". In total nine partners from research and industry were involved in this three year project.

Software

Within the IMEMO project, two state-of-the-art molecular simulation software packages were investigated, both developed by the project partners: ms2 and ls1-MarDyn. The aim of ms2 is to determine thermodynamic properties of substances with high accuracy. This is especially of interest for toxic or ex-

plosive substances, as they are difficult to handle in experiments. Such simulations require only a few thousand molecules. This is a challenge for parallelization, as the main interest therefore lies in strong scaling. During the project new approaches for the underlying algorithms and computational kernels were examined. The scaling behaviour was significantly improved. Furthermore, ms2 was extended by many new physical models, allowing the study of more sophisticated substances and systems. ms2 was made available to the user community as open source software [6]. Another important research area is the simulation of condensation processes and flow

phenomena on the nanoscale – this is the domain of ls1-MarDyn. The simulation of condensation processes requires large systems and is characterized by heterogeneous density distributions of particles, resulting in load imbalances inside the program. This is a major limiting factor for the scalability of domain decomposition based MD codes. A variety of load balancing strategies were developed during the project [3]. These strategies now allow ls1-MarDyn to scale even in the difficult case of such heterogeneous particle distributions. Furthermore, the MPI parallelization was optimized helping to improve ls1-MarDyn's general scalability [5].




Figure 1: Isobaric vapour-liquid phase diagram of water and hydrazine at 0.1013 MPa: X experimental data by Lobry de Bruyn and Dito; (+) experimental data by Uchida et al.; (*) present simulation data with $\zeta=1.3$; (-) Peng-Robinson EOS with $k=-0.1325$.




Figure 2: Density profile of a drop on a surface. The results from the simulation were post processed in Matlab.

Applications

The most relevant results achieved with ms2 are related to the calculation of phase equilibria. The high efficiency and scalability of ms2 (also compared to other available software) makes simulation results highly competitive with experimental results. Figure 1 shows the isobaric vapour-liquid phase diagram of water and hydrazine at 0.1013 MPa from experiment, simulation and Peng-Robinson EOS. The mixture is azeotropic, having a temperature maximum. Hydrazine is highly toxic and dangerously unstable. However, it is a very important substance in chemical engineering, with applications ranging from polymers to rocket fuels. Relevant simulations were performed in the area of nanofluidics. Evaporation

and condensation are among the key topics addressed and to this day pose great questions. With ls1-MarDyn it is possible to simulate such systems on a molecular level. Figure 2 shows the density profile of a drop on a surface, simulated on the nanoscale.

Results

During the project the codes were evaluated on different systems. As can be seen from Figure 3, ls1-MarDyn can now scale up to the full 113,660 core Hermit system at HLRS. In our studies with the optimized ls1-MarDyn code we found a lower limit of roughly 500 molecules per processing element before the communication overhead becomes too high. The optimized version of ms2 is much faster than at

the beginning of the project. In special cases a speedup of one order of magnitude was achieved, compared to the original version, due to algorithmic optimizations. Also the scaling of ms2 on HPC systems was improved to a great extent. A hybrid MPI/OpenMP parallelization was the key to gain additional parallelism which was not available before. Furthermore, the MC simulations can now run with arbitrary numbers of processing elements as long as sufficient memory is available.

Conclusion

We presented the most important results of the IMEMO project. The initial versions of the two codes ms2 and ls1-MarDyn were improved in terms of scalability allowing to scale up to several hundred PEs (ms2) and even more than 100,000 PEs (ls1-MarDyn). At the same time both codes were extended to handle more physical properties of the simulated substances. Various interesting applications were investigated successfully within the project.

Acknowledgement

The IMEMO project was supported by the German Federal Ministry of Education and Research (BMBF) within the program "IKT 2020 - Forschung für Innovationen" in the call "HPC-Software für skalierbare Parallelrechner". The project ran for 39 months starting on the 1st of October 2009 and ending 31st December 2011. The consortium consisted of five project partners.

Partners

- High Performance Computing Centre Stuttgart (HLRS)
- Lehrstuhl für Thermodynamik und Energietechnik (ThEt), Universität Paderborn
- Lehrstuhl für Thermodynamik (LTD), Technische Universität Kaiserslautern
- Lehrstuhl für Informatik mit Schwerpunkt Wissenschaftliches Rechnen (SCCS), Technische Universität München
- Institut für Techno- und Wirtschaftsmathematik (ITWM), Fraunhofer Gesellschaft Kaiserslautern

In Collaboration with

- BASF SE
- Bayer Technology Services GmbH (BTS)
- Evonik Industries AG
- IBM Deutschland Entwicklung GmbH

References

- [1] <http://www.imemo-projekt.de>
- [2] <http://ms-2.de>
- [3] Buchholz, M. Framework zur Parallelisierung von Molekulardynamiksimulationen in verfahrenstechnischen Anwendungen, Dissertation, Institut für Informatik, Universität München 2010
- [4] Technische Universität München. Verlag Dr. Hut, München, August 2010
- [5] Niethammer, C. Performance Evaluation and Optimization of the ls1-MarDyn Molecular Dynamics Code on the Cray XE6, CUG 2012
- [6] Deublein, S., et al. ms2: A Molecular Simulation Tool for Thermodynamic Properties, 2011

- Christoph Niethammer¹
- Colin W. Glass¹
- Martin Bernreuther¹
- Stefan Becker²
- Thorsten Windmann³
- Martin T. Horsch²
- Jadran Vrabec³
- Wolfgang Eckardt⁴

¹ University of Stuttgart, HLRS

² Technische Universität Kaiserslautern

³ Universität Paderborn

⁴ Technische Universität München

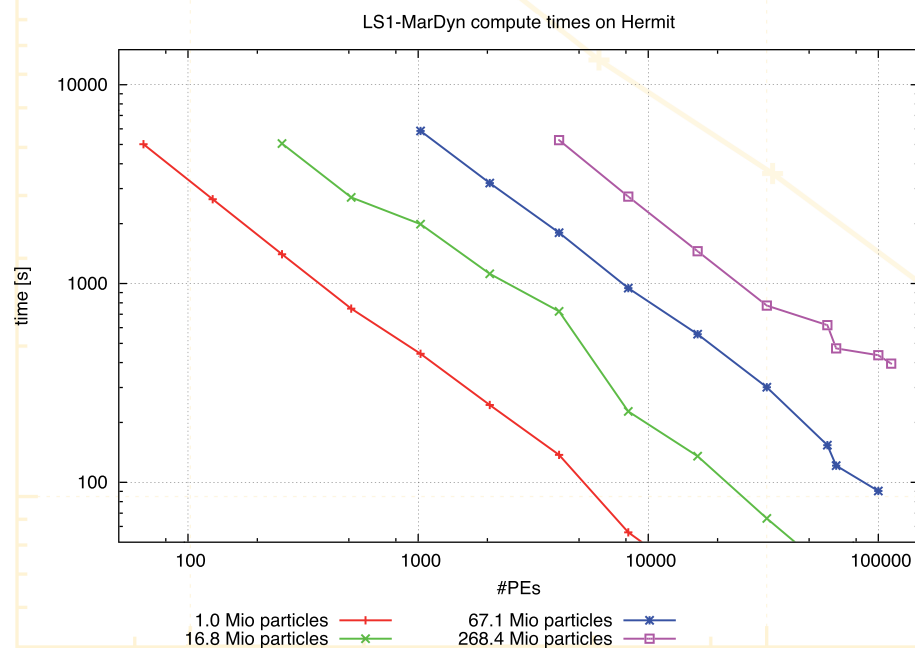


Figure 3: Scaling of ls1-MarDyn on hermit for different numbers of particles.