

Performance of Large Scale Eulerian-Lagrangian Numerical Simulation for Particulate Flow in Rotating Reactors

Salar Tavakkol^{*}, Niklas Weber^{*}, Thorsten Zirwes^{**}, Jordan Denev^{**} and Henning Bockhorn^{*}

*Karlsruhe Institute of Technology, Engler-Bunte-Institute - Division of combustion technology (EBI-VBT)

** Karlsruhe Institute of Technology, Steinbuch Centre for Computing

Abstract

The aim of the current project is the detailed and efficient modeling of thermal processes in biomass particles and the surrounding gas phase in large-scale rotary drum reactors. The reactors operate in the temperature range between 300 and 400 degrees Celsius in a steam-filled environment under atmospheric pressure. The work presents the development of an Eulerian-Lagrangian model with the code OpenFOAM for dense particulate flows. In this work, numerical and parallelization aspects of the simulations are presented and discussed.

1. Introduction

Even with the current supercomputer resources, it does not seem to be feasible to use the resolved CFD-DEM (Computational Fluid Dynamics for the Eulerian gas phase and the Discrete Element Method for the movement of discrete particles) for the simulation of large-scale reactors or 3D systems [1]. The unresolved Lagrangian model known as Multi-Phase Particle-In-Cell (MP-PIC) method presents a good and accurate alternative to the DEM. The MP-PIC method is described in [2].

For a large number of particles, computing the movement of the discrete phase with the MP-PIC model is at least two orders of magnitude faster than with the DEM (called in OpenFOAM the DPM for Discrete Phase Model): while in the DEM the computational burden raises exponentially by increasing the number of particles in the domain, the increase is linear for the MP-PIC method. New solvers for OpenFOAM have been developed which can track the particles either with MP-PIC or with DPM and also can handle the thermally-driven processes in the particles like evaporation, boiling or devolatilization as well as chemical reactions. The main functions of each solver and the data exchange between them are depicted in Figure 1.

The present study focuses on the computational and parallelization aspects for the Eulerian and the MP-PIC solvers in view of their coupled use for modeling of the thermal conversion of biomass particles in rotary drum reactors.



2. Physical processes and requirements to the simulations

The conversion process in the reactor takes place for a duration of 30 to 45 minutes. Consequently, to reach a steady-state condition, the physical time that is simulated has to be at least 60 minutes. With a fine numerical grid, due to the Courant-Friedrichs-Lewy (CFL) restriction to the time step, this requires a large number of time-steps. The CFL number is defined as $CFL = \frac{u.\Delta t}{\Delta x}$ and in the present study the time step Δt has been adapted so that CFL < 0.25. In the above equation, Δx is the cell size in any spatial direction and *u* is the corresponding velocity component in this direction.



Fig. 1 Final layout of the developed Eulerian-Lagrangian solver

The MP-PIC method for the particles assumes implicitly that they are smaller than the half grid cell for the gaseous phase. Therefore, if the particles were large in size, they would restrict to some extent the refinement of the grid for the Eulerian solver. However, with the large dimensions of the rotating drum reactor - a length of 12 m and a diameter of 1.2 m, the particles, which are around 5-6 mm in diameter, are always much smaller than the grid for the gas phase.

In the present study, up to half a million particles are tracked and up to 2.7 million grid cells have been used for testing. Together with the large number of time-steps, this requires the parallel use of a large number of cores on the bwUniCluster.

The degree of filling of the reactor with particles presents a real challenge for the parallelization. The reason is that the filling evolves in time: at the beginning, the reactor is empty, then it is continuously filled up (e.g. from left to right) and at the end, it is filled (around 10 - 15% by volume) with particles. Because the domain decomposition is based on the static Eulerian grid, but the particle's load is non-even throughout the reactor and is changing in time, the scaling is expected to deviate from the ideal case. This is presented and discussed in the next chapter.



3. Results and discussion

Due to the large computation domain and the need to cover it with a reasonably fine numerical grid, the Eulerian computational time can become relatively high. Therefore, an optimization of the number of grid cells was made. Fig. 2 shows the impact of the mesh refinement on the performance of simulations with a fixed number of particles (approximately 6000). The reference values are taken from the finest grid simulated (2x2x2 cm). The results from the different grids were compared in terms of temperature, gas velocity, and composition of the gas phase. The comparison showed that the cubic cells give more accurate results comparing to long hexahedral cells. Further, there was a good agreement between the finest mesh resolution and the coarser geometry with cell sizes of 4x4x4 cm. The computation time for this coarser mesh was regarded also as suitable for the present study. Therefore, the grid size of 4x4x4 cm has been taken as the base case for further simulations.



Fig. 2 Execution time for different size of grid cells for 230 seconds of real-time

As a next step, the performance of the MP-PIC Lagrangian solver for an increasing number of particles have been evaluated. Fig. 3 shows the CPU-time for a single-core simulation. There is an almost linear increase in the computation time with the MP-PIC solver when the number of particles is increased. This result is logical, as this method does not require seeking the neighbor particles, which explains its linear behavior.

Finally, for a simulation with 375 000 cells that contains 450.000 particles, the scaling performance has been evaluated on multiple nodes (one node has 20 cores) of the BwUniCluster. Fig. 4 exhibits the results of tests with different number of cores for a fixed wall-clock CPU-time of 170 seconds. It highlights the fact that the efficiency of parallelization is lower than the ideal case and that it decreases considerably when using more than 560 processors. The reason is in the increasing communication



between the processors. Due to the relatively long waiting times in the multinode queue of the bwUniCluster, 280 processors have been found to be feasible to perform the long-run simulations of the large-scale drum reactor.



Fig. 3 CPU time used versus the number of particles in the domain



Fig. 4 Performance of parallel run for a case study with a high load of particles

References

[1] W. Zhong, A. Yu, G. Zhou, J. Xie, and H. Zhang, "CFD simulation of dense particulate reaction system: Approaches, recent advances and applications," *Chem. Eng. Sci., vol. 140, pp. 16–43*, 2016.

[2] P. J. O'Rourke and D. M. Snider, "An improved collision damping time for MP-PIC calculations of dense particle flows with applications to polydisperse sedimenting beds and colliding particle jets," *Chem. Eng. Sci., vol. 65, no. 22, pp. 6014–6028,* 2010.



Aerodynamic development of a formula student racecar using STAR-CCM+ on the bwHPC

Jonas Pangerl^{*}, Niklas Kaiser,^{*} and Rainer Stauch^{*}

*Hochschule Esslingen, University of Applied Sciences

Abstract

Automated CFD simulations are used to optimize the aerodynamic performance of a formula racecar. Different driving conditions are studied by the presented simulation models which allow a high throughput of simulations. The focus is on the modeling of cornering vehicle simulations in a curved wind tunnel using the moving reference frame (MRF) approach.

1. Introduction

As part of the student project "Formula Student", a group of students design, manufacture and test a formula vehicle following internationally valid regulations within one year. An improvement in vehicle performance can be achieved by the optimization of the aerodynamics. The design of the aerodynamic package is carried out with the help of CFD simulations, with the aim of generating the lowest possible drag force of the vehicle and a high, down-facing force, called downforce. The downforce allows greater lateral forces to be transmitted through the tire which results in faster cornering speeds.

The development of the aerodynamics package requires many iteration steps and hence, a lot of simulation runs. The goal is to be able to compare many different geometry variants within the shortest possible time. One challenge is to automate the simulation process at its best way, but without making any compromises in terms of robustness and flexibility. To increase the throughput of simulations, different models are used for the different development steps. As the level of development increases, the complexity of the simulation models increases, without reducing the level of automation. In the last racing season, during the 3.5-month development phase, more than 1800 CFD simulations were carried out on the bwUniCluster. Parallel to the development of the aerodynamics package, a constant improvement of the CFD simulation model takes place. In addition to improving the stability of the individual simulations, the aim is to achieve the correct mapping of increasingly complex driving conditions to the simulation models with high accuracy.



2. Simulation Models

In the development process of the aerodynamics of vehicles, the driving conditions on the road should be represented by the conditions in the wind tunnel. To achieve that, a change of the reference system is necessary. Whereas on the road the vehicle is moving and the air is quiescent, in the wind tunnel the vehicle, which is standing still, is exposed to an airflow. The same holds for a CFD simulation model, which represents a virtual wind tunnel. If all boundary conditions, like the surface of the road, are also transferred to the reference system, the results of the wind tunnel, and of the virtual wind tunnel, are not affected by the change.

To accelerate the development process, a simplified half car model is used for the driving straight ahead condition. This model can be divided into three subclasses with different complexity, as shown in Fig. 1. Increasing the complexity of the simulation, a full car model is simulated in the second step. This second model also has subclasses. The third model offers the highest complexity. Here, the vehicle is placed in a curved wind tunnel to simulate the cornering airflow. Different parameters such as pitch angle, roll angle, and ground clearance can be varied and investigated. The dependencies of the drag and the lift force and of the aerodynamic balance on the influencing parameters are required by the aerodynamics department and suspension team for development and are essential for a good vehicle performance.

		Model number	Flow direction	Characteristics		Number of Cores (PPN)	Number of Cells	Duration	Average Simulations per Month*
Halfcar	Pressure distribution halfcar with symmetry plane. View from below.	1	straight forward	no cooling systems	high Y+	1 & 2 (28)	20 mio	2-4 h	50
		2	straight forward	no cooling systems	low Y+	2 (28)	36 mio	4 h	100
		ŝ	straight forward	CHT + cooling systems	low Y+	2(28)	42 mio	4 h	100
ullcar	Pressure distribution fullcar at 15° yaw angle airflow. View from below.	4	staight forward	CHT + cooling systems	low Y+	4 (28)	84 mio	5 h	25
		ъ	Yaw angle	no cooling systems	low Y+	4 (28)	75 mio	4 h	75
ш		9	Yaw angle	CHT + cooling systems	low Y+	4 (28)	84 mio	5 h	100
Fullcar cornering	Pressure distribution in the curved windtunnel	7	cornering	CHT + cooling systems	low Y+	4 & 6 (28)	90 mio	6-8 h	50

* during the development process

Fig. 1: Different simulation models and their simulation details



3. Cornering simulation

For the modeling of the cornering airflow, a different setup of the simulation model is needed. For the simulation of a straightforward or of a yaw angle airflow, the reference system in the simulation model is chosen in opposition to reality. The car is standing still and the air flows around the car. For the modeling of a cornering vehicle the reference system cannot be switched in the same simple way as for the driving straight ahead condition as the curved airflow around the vehicle has to be taken into account in the simulation model, too. This makes this driving condition a special challenge. As shown in Fig. 2, the airflow must be kept on a curved path. For this purpose, the entire simulation region is rotated by means of an MRF (moving reference frame) approach around the center of the corner according to the driving speed. Strictly speaking, there is no rotation with the MRF, i.e. no movement of the vehicle in the simulation model (= virtual wind tunnel). In fact, the MRF only models the effect of the rotation on the fluid in the wind tunnel by the implemented forces. An analysis of the simulation then takes place in a rotating coordinate system of the MRF region and maps the flow around the car. For the inlet boundary condition of the airflow in the curved wind tunnel, a velocity condition having the value of 0 km/h and for the outlet boundary condition, a pressure boundary condition having an absolute pressure of 1 bar is selected. An inflow velocity greater than zero is not necessary, because the car moves virtually, according to the reality in the quiescent air, and the virtual airflow is generated by the rotation of the reference frame. The walls of the wind tunnel have no friction, the street is rotated with a tangential velocity condition against the direction of driving (opposite to the MRF rotation direction) at the same angular velocity as the airflow. The superposition of tangential velocity and MRF lead to a standing road on which the car seems to move. This simulation technique is used in the same way in professional motorsport, for example in Formula 1, and allows the study of the complex behavior and effects of aerodynamics of cornering vehicles [2]. Due to occurring lateral and longitudinal accelerations, the position of the car changes relatively to the road, which has great effects on the local airflow at different parts of the vehicle and thus leads to a change in drag and lift force. When designing the chassis and the setup of a racecar, the consideration of these effects is an important aspect and must be known as accurately as possible.



Fig. 2: Illustration of the cornering airflow in the virtual wind tunnel setup



4. Conclusion

For the first development steps of the cornering airflow simulations, a simplified geometry, the Ahmed Body, shown in Fig. 3, was used set up the straight wind tunnel to get results comparable to literature values.



Fig. 3: Ahmed Body [3]

The lift and drag coefficient as well as the moments on the different axles can be seen in Table 1. The coordinate system for the evaluation of the acting forces was placed at the front left edge of the Ahmed body. The influence of cornering airflow on the lift force (lift coefficient C_L) of the Ahmed Body is very small, in contrast to the influence on the drag force (drag coefficient C_D) and on the moment around the Z-axis (Moment Z) is high. The increasing drag can be explained with an increasing projected frontal surface to the incoming air and also with a larger joint vortex at the rear end of the car to be caused by the yaw angle airflow in this area. The increasing Z Moment, which means higher side forces affecting the car, can also be explained with an increasing with an increasing frontal surface projected to the flow direction.

	Velocity	Moment X	Moment Y	Moment Z	CD	CL
Straight forward	15 m/s	18,76 Nm	-0,81 Nm	-17,77 Nm	0,309	0,326
Cornering airflow	15 m/s	19,38 Nm	-0,56 Nm	-23,96 Nm	0,392	0,327

References

- [1] Keogh, J., Doig, G., S. Diasionos, S., and Barber, T.: The influence of cornering on the vortical wake structures of an inverted wing. Journal of Automobile Engineering, 229(13), pp. 1817-1829,
- [2] Keogh, J., Barber, T., Diasinos, S., Doig, G.: Techniques for Aerodynamic Analysis of Cornering Vehicles. SAE Technical Paper 2015-01-2022
- [3] Tunay, T, Sahin, B.,Ozbolat, V.: Effects of rear slant angles on the flow characteristics of Ahmed body. Journal: Experimental Thermal and Fluid Science 57 pp. 165-176

TP2x-HPC, AN OLD FEM GENERAL PURPOSE PROGRAM NOW AVAILABLE ON bwUNICLUSTER KIT Prof. Dipl.-Ing. Peter Groth

HS Esslingen ,University of Applied Sciences, Basic and Automotive Faculty

Abstract:

TP2x-HPC is based on TPS10 (T-Programm GmbH), 1971*- end of 90th market leader in Germany especially in automotive industry with about 500 company licenses. Development in continuation since 1995** - today. Parallelization (OMP/MPI) since 2016*** on Esslingen Grid, since 2018 on bwUniCluster KIT.

1. Introduction

TP2x.HPC supported all main FEM Analysis Types, as linear Elastostatics, non-linear Elastostatics, Dynamics, Potential Problems and Multi-Physics. Free license for State Education and Research world-wide, Fortran source code for projects.

2. Examples

2.1 Linear Elastostatics



- Toyota Camry model by George Mason University (FHWA) USA, 2.25 Mio. Elements with 12.8 Mio. Equations and 5.000 Weld Spots, linear static's: load-case torsion between axles, time: bwUniCluster KIT, 4 nodes 16 MPI Tasks: ca. 9 h, 16 nodes, 64 MPI tasks, ca. 6 h, single desktop 98 h





- historically bridge in Pfullingen, IGF, solid model, load-case some load positions, right to left (avi)

2.2 Non-linear Elastostatics, yielding



- car tank under rear seat, shell mesh 0,32 Mio.
DOFs, 10 load steps, increasing
10x10 mm² point load, Alfa Romeo (avi)



 disk spring with tongue, sector symmetric parabolic solid mesh, Steinbeis Transf. Zentr. ES breaking-down load, unload and reload
28 steps, yield stress with kin. hardening (avi)

2.3 Dynamics (extension to explicit crash calculation, partner ??)



2.5 Laminar Flow (extension to turbulent flow possible, partner ??)



- flow in a tube around a ball in mm/sec axissymmetric model, 3D possible

2.6 Acoustics



 2D sound pressure distribution in a passenger car, 3D and damping possible,
1.- 4. eigenfrequencies, BMW

2.7 Magnetostatics 2D + axissymmetric, (extension for solids possible, partner??)



 axissymmetric lifting magnet, induction lines yellow: air, magenta: magnet iron, green: coil, dark blue: lifted iron

2.8 Contact with friction



 friction clutch quarter-symmetric model shaft and rotation box steel, friction part between special material- developed by Freie Uni Berlin KfZ department, arrows = friction forces

2.9 Multi-Physics (extension to BEM solid elements possible, partner??)



BEM/FEM coupling at Microelectronics
TU München, Techn. Elektrophysik
BEM: electrical forces - FEM: mechanical
answer, airbag sensor deformation < gap (avi)

3.0 Pre- and Post processing

full integrated in TP2x:

Siemens NX Femap T-Systems Medina (bif/bof), Ansa (bif/bof), eFEA (Eigenentwicklung)

4.0 Partnership based on TPS10/TP2x since 1971

1971-1978 Basisentwicklung TPS10 2 Mio. € Zuschuss BMFT/BMBF 1979-1989 TU München Festkörperphysik, Erweiterung Mikroelektronik, BMBF TU München Bauwesen, Freiformat / FEDIS FEM Ein- Ausgabe Standard, BMBF TU Erlangen Boundary Elemente Betsy/Betti, VDMA TU Stuttgart Baumechanik, Schalenelement Laminat UNI Tübingen, Orthopädie, Knochenberechnung UNI Wien, Bauwesen, Gummielement, Tunnelelement UNI Brünn, Magnetostatik Lomonossow UNI Moskau, frei-frei Schwingungen 2000-2004 TU Berlin, Kfz-Technik, Fliehkraft-Kupplung, Kontakt/Reibung TU München Techn. Elektrophysik, Mikroelektronik BEM/FEM Kopplung, BMBF TU Chemnitz Mikroelektronik, Squeez-Film Dämpfung, BMBF

5.0 Documentation

German and English manual, 400 pages, 150 examples for each type of application, text book Springerverlag 2001 "FEM Anwendungen" (mit Femap) with description of all analysis types.



see flyer for more information





Modeling and simulation of liquid aerosol deposition in textile filters

Dennis Hoch, André Baumann, and Jennifer Niessner*

*Heilbronn University of Applied Sciences

Abstract (optional)

We study the improvement of filter cartridges for oil mist filtration in production halls as well as air dryer cartridges for air brakes of trucks and busses in cooperation with Junker-Filter GmbH, a local SME. By means of modeling and CFD simulation we optimize both filters with respect to energy efficiency (pressure loss) as well as filtration efficiency. While filters for production halls work at ambient temperature and pressure, air cartridges are operated at temperatures as high as 90°C and pressures up to 20 bar.

1. Introduction

Like fine dust, liquid aerosols may cause lung cancer, allergies, asthma or other diseases when the droplets are small. Therefore, we develop efficient mist filters with less energy consumption by means of CFD simulation that are especially capable of depositing the – hard to catch- fine droplets.

2. Description of the workflow

From nano CT scans of current filter materials CAD geometries of the real fibrous media are generated by means of image processing, see Fig. 1. CFD simulations using StarCCM+ (partly on bwUniCluster) on the so-called micro-scale characterize both the separation efficiency of the material (so-called fractional efficiency i.e. size-dependent separation efficiency) and allow for a determination of parameters. These parameters include porosity, permeability and constitutive relationships (relative permeability – saturation relationship and capillary pressure – saturation relationship) as well as a "deposition term" for oil (oil source term) for use in the macro-scale simulations. As single droplets and single fibers are resolved on the micro-scale the simulation is restricted to small sections of the filter (cubes with edge lengths of 1 mm) – even when using high-performance computing. The additional challenge is that these simulations are also used in order to generate the optimum filter material with respect to separation efficiency by means of stochastic generation of virtual fibrous media (fiber length, diameter, orientation, coating etc.).





Fig. 1: Workflow from nano CT scans to series production of oil mist filters

The macro-scale parameters and constitutive relationships enter the CFD simulations using ANSYS Fluent on the macro-scale where no droplets and fibers are resolved any more. Instead, averaged parameters are used (porosity, oil and air saturation in a cell etc.), see Fig. 2. Therefore, the complete filter can now be considered and the pressure loss across the filter can be determined. On the macro-scale the filter setup is to be optimized (several layers, alternating coarse, fine, oleophilic, and oleophobic layers) in order to reduce pressure loss and thus, energy needs. As many parameters are involved in the optimization process, many simulations are needed. Junker-Filter GmbH then produces the optimum filter as physical prototype.

3. Usage of bwUniCluster

For the micro-scale simulations, the StarCCM+ 12.06.011 installation by Esslingen University of applied sciences was used. As shared memory was critical, the fat nodes were used with processes per node ranging from 10-32, memory allocation of 15,000-64,000 MB and wall times between 5 and 20 hours. Per default, the platform MPI was used. It was recommended to further increase shared memory by setting MPI_GLOBMEMSIZE to 1400000000.

For the macro-scale simulations, ANSYS Fluent R19.1 was installed on bwUniCluster by Esslingen university of applied sciences. The installation was successfully tested. However, since the setup of the first simulation models was more time consuming than expected, the compilation of user defined functions on the cluster was challenging, and the simulation times were shorter than expected, most simulations were carried out on a local machine so far. However, for the optimization process planned in the near future, Fluent simulations on bwUniCluster will be carried out. Wall times, number of nodes etc. can so far not yet be quantified.

4. Results and Discussion

The geometry of current fibrous filter materials was determined using nano CT scans. Micro-scale simulations were carried through in order to obtain fractional efficiency as well as macro-scale parameters and constitutive relationships, see Fig. 2 (left hand side). For these micro-scale simulations both a Lagrangian and a discrete element



approach (DEM) were used in order to determine the droplet motion. Macro-scale simulations were then performed in order to investigate oil saturation and air pressure within the filter, see Fig. 2 (right-hand side).



Fig. 2: Micro-scale simulation (resolution of droplets and fibres, left hand side) and macro-scale simulation (volume-averaged quantities like oil saturation and air pressure - here in a 90° model, right-hand side)

By comparison to experimental data from a test rig at Heilbronn University, the simulations were validated. Both the pressure drop of the clean filter (see Fig. 3) and the pressure drop of the oil-laden filter (see Fig. 4) are in good agreement with experimental data.



Fig. 3: Validation of the CFD model by comparing the pressure loss of a clean filter to experimental data





Fig. 4: Validation of the CFD model by comparing the pressure loss during oil loading to experimental data

5. Conclusions

Thanks to high-performance computing on bwUniCluster, it is possible to speed up the product development process by virtual optimization of filter prototypes. Due to the virtual optimization the number of physical prototypes can be reduced and less (extremely time-consuming) experiments are necessary.

Acknowledgements

This work was funded by the federal state of Baden-Württemberg and the "European Regional Development Fund" (EFRE), Vorhaben: High-Tech-Aerosolnebelabscheider im Zero-Design (HAAZ) FKZ: 1096283.

Atomistic simulations of processes in energy conversion and storage

Axel Groß Institute of Theoretical Chemistry Ulm University 89069 Ulm/Germany

September 24, 2019

Abstract

The theoretical and numerical methods to describe atomistic processes in energy conversion and storage will be briefly presented. Furthermore, their main fields of application wll be sketched and illustrated with recent examples.

1 Introduction

Processes in energy conversion and storage are of strong current interest due to the importance of these processes for our future energy technology [?]. Due to the improvement in the computer power and the development of numerically more efficient methods, these processes can now be reliably addressed by atomistic simulations, allowing close collaborations between theory and experiment, as will be demonstrated in this contribution.

2 Theoretical and Numerical Issues

The atomistic simulation of processes in energy conversion and storage requires a quantum chemical treatment as bond-breaking and bond-making processes are involved. Since the systems of interest in this field are rather complex, density functional theory (DFT) calculations within the Kohn-Sham formalism correspond to the method of choice as DFT combines numerical efficiency with a satisfactory reliability for a large number of systems. Especially for extended system, periodic DFT calculation are most efficient which are often based on a plane-wave representation of the electronic oneparticle states. The potential energy terms in the Kohn-Sham Hamiltonian are evaluated in real space whereas the kinetic energy terms are processed in momentum or Fourier space. Correspondingly, most of the numerical effort is associated with Fast Fourier Transform routines. Still, these methods do typically not scale too well with the number of cores.

3 Results and Discussion



Figure 1: Illustration of the door-opener mechanism responsible for the fast diffusion of oxygen atoms on a CO-covered surface [?]. Source: https://www.en.uni-muenchen.de/news/newsarchiv/2019/wintterlin_katalysatoren.html

Atomistic simulations can serve several purposes, as will be illustrated in

this contribution. First of all, they can be used in close collaborations with experimental colleagues to identify processes that our beyond the resolution of the experimental tools. A recent example is the diffusion on a crowded surface which is relevant in heterogeneous catalysis. Here DFT calculations have clarified that the surprisingly fast motion of oxygen atoms on a ruthenium surface almost fully covered by CO molecules can be clarified by a door-opener mechanism [?] which is illustrated in Fig. ??.



Figure 2: Correlation between atomistic diffusion processes and fire hazards in batteries [?].

Furthermore, atomistic simulations can be used in a modeling approach. These calculations do not necessarily need to be fully realistic, but can relate specific atomistic processes to macroscopic phenomena. In such an approach it has been show that self-diffusion barriers can obviously be used as a so-called descriptor for dendrite growth in batteries [?]. which represent severe hazards in battery operation, as illustrated in Fig. ??.

In addition, atomistic simulations can be used as a tool for high-troughput screening in order to identify promising new materials for a certain application. As an example, the determination of the open-circuit potential for fluoride batteries will be presented [?] which allows to suggest promising anode and cathode materials in these batteries.

4 Acknowledgements

This work contributes to the research performed at CELEST (Center for Electrochemical Energy Storage Ulm-Karlsruhe).

References

- F. Gschwind, G. Rodriguez-Garcia, D. Sandbeck, A. Gross, M. Weil, M. Fichtner, and N. Hörmann. Fluoride ion batteries: Theoretical performance, safety, toxicity, and a combinatorial screening of new electrodes. J. Fluorine Chem., 182:76 – 90, 2016.
- [2] A.-K. Henß, S. Sakong, P. K. Messer, J. Wiechers, R. Schuster, D. C. Lamb, A. Groß, and J. Wintterlin. Density fluctuations as door-opener for diffusion on crowded surfaces. *Science*, 363:715–718, 2019.
- [3] M. Jäckle, K. Helmbrecht, M. Smits, D. Stottmeister, and A. Groß. Selfdiffusion barriers: possible descriptors for dendrite growth in batteries? *Energy Environ. Sci.*, 11:3400–3407, 2018.
- [4] O. M. Magnussen and A. Groß. Toward an atomic-scale understanding of electrochemical interface structure and dynamics. J. Am. Chem. Soc., 141:4777–4790, 2019.

QUANTUM TRANSPORT AND NONEQUILIBRIUM PHENOMENA: FROM NANODEVICES TO MACROSCOPIC SYSTEMS

Martin Žonda, Jakob Bätge, Christoph Kaspar, Yaling Ke, Junichi Okamoto, Rajagopala Seelam, Rudolf Smorka, Sebastian Stumper, Sebastian Wenderoth and Michael Thoss Institute of Physics Albert Ludwig University of Freiburg

September 16, 2019

1 Introduction



Illustration of a molecular junction with graphene electrodes

The theoretical research of our group covers a wide range of aspects and phenomena in the fields of theoretical condensed matter physics, chemical physics and nanophysics. Our main focus is on nonequilibrium processes in many-body quantum systems. This includes, but it is not limited to, the study of quantum transport through nanostructures (e.g. quantum dots or molecules) crucial for future electronics; transient states in quantum chains far from equilibrium, related

to engineering of materials with new exotic properties; or the fundamental problem of thermalization of isolated systems important, e.g., for the research of quantum computers. These problems call for a variety of methods with substantial computational requirements, each suited for different time and size scales.

2 Description of Methods

A major part of our work is dedicated to the development and improvement of efficient and flexible numerical methods. We briefly introduce the main ones and discusses their typical requirements and utilization in the bwHPC clusters.

The hierarchical quantum master equation approach (HQME) provides a numerically exact description of nonequilibrium charge transport in nanosystems which can include the influence of electronic-vibrational coupling and strong mixture of quantum dynamics of electrons and nuclei. The solution of HQME leads to a system of 10^3 - 10^7 (depending on precision) linearly coupled differential equations. Computationally, the method can be decomposed to numerous linear algebra tasks such as sparse matrixmatrix and matrix-vector multiplication, QR factorization or vector-vector addition. The standard numerical libraries such as Intel MKL and CUDA cuBLAS/cuSPARSE can be utilized to tackle many of these tasks with high efficiency. The requirements then depend on the addressed problem but typically we need 16-120GB of RAM, 8 cores and 1-8 days of computational time on JUSTUS or 12GB of RAM, 1 CPU core, 1 GPU and 10 days of computational time on BINAC.

The multi-configuration time-dependent Hartree approach (MCTDH) is a variational basis-set method using a very efficient representation of the wave function and its time-propagation which overcomes the usual restriction of many exact methods where the required basis grows exponentially with the number of degrees of freedom. This makes the method suitable not only for addressing overall dynamics of molecules but also problems like many-body localization in quantum systems. Technically, the method is equivalent to solving a time-dependent first order differential equation where the most time consuming tasks are matrix multiplications. A typical calculation requires 10GB of memory and up to seven days of running time on a single core on JUSTUS. A combination of the nonequilibrium Green's function with Monte-Carlo methods (neGFMC) can be used to study transport through layered systems of correlated electrons and its influence on phase transitions, such as metal-insulator of order-disorder ones. Computationally, the most challenging tasks are exact diagonalizations of small clusters ($\approx 10^3 \times 10^3$) and multiplications of dense matrices of the same size which are, however, performed more then 10⁷ for every Monte Carlo (MC) run. The most efficient way to run these calculations is to use parallel MKL libraries for the matrix manipulations at 1 – 4 cores and a poor man's parallelization for collecting the proper MC statistics. Typically, we require 1GB RAM, 1-7 days of calculation time and tens to hundred independent runs for every parameter set.

Besides these methods we also utilize a various standard techniques, packages and libraries. Here we list the most important ones.

To study the time evolution in interacting quantum many-body systems we use time dependent density matrix renormalization group (DMRG), the time-dependent variational principle (TDVP) and exact diagonalization ED. The first two methods are based on locally updated matrix product states and require 1GB and 2-7 days for a typical calculation. ED is more extensive as it typically requires 10GB of memory but this can be easily scaled up to 100GB, by increasing the system size. A typical single-core calculation takes from a day to two weeks. To study supercurrent and quantum phase transitions in quantum dots coupled to superconducting leads we employ numerical renormalization group open source implementation NRG-Ljubljana which requires a commercial software Mathematica. A typical single-core calculation can be run with 6GB of memory and takes 2 days. Further, to obtain potential energy surfaces of molecules we employ the ab initio quantum chemistry package GAMESS where a parallel calculation using 4 nodes (64 cores) takes up to 4 hours. The JUSTUS cluster is our preferred choice for all these methods.

3 Results and Conclusions

Four of our most recent publications, for which we have utilized the bwHPC facilities, can serve as typical examples of our work.

In Ref. [1], we investigated intramolecular singlet fision in pentacene dimers linked by various spacers which holds the potential to improve the efficiency of solar cells. We suggest new ways of enhancing singlet fision employing chemical substitution or solvent effects in these dimers.

An optically excited one-dimensional ionic Hubbard model is studied in Ref. [2]. This is motivated by the recent development of experimental techniques working with intense laser fields, which can drive different systems far from equilibrium and allow that way an analysis of before unknown transient states. We demonstrate a distinct nature of optically excited states in Mott and band types of insulators.

Refs. [3, 4] are dedicated to charge transport through large system of correlated electrons. We investigate the influence of various typical phases found in such materials, e.g. charge density waves or Mott insulator, on the nonequilibrium current and show that the current can be used to identify various phase transitions. By studying the problem of localization we demonstrate that the presence of impurities in these systems plays a different role for thin and broad systems.

The authors acknowledge support by the state of Baden- Württemberg through bwHPC and the German Research Foundation (DFG) through Grant No. INST 40/467-1 FUGG.

References

- S. Rajagopala Reddy, Pedro B. Coto, and Michael Thoss. Quantum dynamical simulation of intramolecular singlet fission in covalently coupled pentacene dimers. J. Chem. Phys., 151(4):044307, 2019.
- [2] Junichi Okamoto. Time-dependent spectral properties of photoexcited one-dimensional ionic Hubbard model: an exact diagonalization study. *arXiv e-prints*, page arXiv:1909.02471, Sep 2019.
- [3] Martin Zonda and Michael Thoss. Nonequilibrium charge transport through falicov-kimball structures connected to metallic leads. *Phys. Rev.* B, 99:155157, Apr 2019.
- [4] Martin Zonda, Junichi Okamoto, and Michael Thoss. Gapless regime in the charge density wave phase of the finite dimensional falicov-kimball model. *Phys. Rev. B*, 100:075124, Aug 2019.



Network remodeling induced by transcranial brain stimulation: a computational model of tDCS-triggered cell assembly formation

Han Lu*,**, Julia V Gallinaro*, and Stefan Rotter*

*Bernstein Center Freiburg, Biology Faculty, University of Freiburg **Institute of Cellular and Integrative Neurosciences, University of Strasbourg

1. Introduction

Transcranial direct current stimulation (tDCS) is a variant of noninvasive neuromodulation, which promises treatment for brain diseases like major depressive disorder [1,2]. In experiments, long-lasting aftereffects were observed, suggesting that persistent plastic changes are induced [3]. The mechanism underlying the emergence of lasting aftereffects, however, remains elusive.

2. Description or Method or Numerical Issues

Here we propose a model, which assumes that tDCS triggers a homeostatic response of the network involving growth and decay of synapses [4,5,6]. The cortical tissue exposed to tDCS is conceived as a recurrent network of excitatory and inhibitory neurons, with synapses subject to homeostatically regulated structural plasticity. We systematically tested various aspects of stimulation, including electrode size and montage, as well as stimulation intensity and duration with NEST.

3. Results and Discussion

Our results suggest that transcranial stimulation perturbs the homeostatic equilibrium and leads to a pronounced growth response of the network. The stimulated population eventually eliminates excitatory synapses with the unstimulated population, and new synapses among stimulated neurons are grown to form a cell assembly. Strong focal stimulation tends to enhance the connectivity within new cell assemblies, and repetitive stimulation with well-chosen duty cycles can increase the impact of stimulation even further. One long-term goal of our work is to help in optimizing the use of tDCS in clinical applications.





Fig. 1 Modeling the effect of tDCS on cortical networks. (A) It is assumed that transcranial stimulation leads to a weak polarization of the neuron's membrane potential (left). For a point neuron, this is achieved by injecting a current of suitable strength into its soma (right). (B) Firing rate modulation with the angle θ for three different values of Δ Vm (dotted lines on C). (C) Firing rate of a neuron, the ongoing activity of which is modulated by tDCS, for different values of θ and membrane polarization Δ Vm. The contour lines correspond to 7 Hz, 8 Hz, and 9 Hz in white, orange, and maroon. (D) Electrode montages used in tDCS. (E) The region of interest subject to tDCS is modeled as a recurrent network of excitatory and inhibitory neurons. (F) Excitatory-to-excitatory synapses require the combination of a bouton (empty triangle) and a spine (red dot). The growth rate of both types of synaptic elements depends linearly on firing rate. (G) The network is grown from scratch before each tDCS stimulation experiment.





Fig. 2 tDCS triggers the formation of cell assemblies. (A) A subgroup comprising 10% of all excitatory neurons in a larger network is stimulated by tDCS. (B) Average firing rate of directly stimulated (blue) and unstimulated (gray) excitatory neurons before, during, and after applying a depolarizing stimulus. (C) Average connectivity among stimulated neurons (blue), among unstimulated neurons (dark gray), and between neurons belonging to different groups (light gray) upon depolarizing stimulation. (D–E) Similar to (B–C), but for a hyperpolarizing stimulus. Shaded areas on (B–E) indicate the stimulation period. (F) Illustration explaining the process of structural plasticity that happened after a depolarizing tDCS. The stimulation triggers the removal of interpopulation synapses, and accelerates the growth of synapses among stimulated neurons, leading to the formation of cell assemblies.





Fig. 3 Interactions between subpopulations and cell assembly formation in more complex stimulation paradigms. (A) Tri-group scenario: 30% of all neurons in a network (G1) are depolarized by 0.1 mV, another 30% (G2) are hyperpolarized by -0.1 mV, and the rest of 40% receives no stimulus. (B) Bi-group scenario: 30% (G1) are hyperpolarized by -0.1 mV, and the remaining 70% (G2) are depolarized by 0.1 mV. (C,E) Group averages of firing rates in G1 (blue) and in G2 (yellow) before, during, and after stimulation. (D,F) Group averages of the connectivity within G1 (blue), within G2 (yellow), and between G1 and G2 (gray).





Fig. 4 Comparison of tDCS effects with different electrode montage, as well as stimulus focality and intensity. (A) Bi-group stimulation scenario. (B) Uni-group stimulation scenario. (C) Tri-group stimulation scenario. (D–F) Integrated G1 cell assembly connectivity (IG1) at different focality and intensity levels for scenarios (A–C). (G,H) Difference between D and E, as well as F and E, respectively. (I) Integrated G1 cell assembly connectivity integrals (IG1) for different stimulation intensity levels for a specialized bi-group scenario, where G1 and G2 comprise half of the excitatory population, respectively. The white squares correspond to situations where the difference between stimulation intensities of both groups amounts to 0.8 mV.





Fig. 5 Repetitive stimulation boosts network remodeling. (A) A subnetwork of excitatory neurons (10%) is stimulated with a train of DC stimuli. Stimulation time is t1, followed by a pause of duration t2. (B, C) Average firing rate and connectivity during a train of stimuli. (D) For the same total stimulation time (6, 000 s), the boosting depends on the exact repetition protocol. (E) The peak connectivity reached depends on the stimulation intensity; an asymmetric repetitive protocol (t1 = 75 s, t2 = 150 s) was used for all simulations here.





Fig. 6 Comparison of three different scenarios for repetitive DC stimulation. (A) 10% of excitatory neurons were stimulated, using the same temporal protocol (t1 = 150 s, t2 = 150 s) in each case, but different amplitudes and polarities were employed, as indicated by the three different curves. (B) Evolution of average connectivity for the different stimulation scenarios; color code matches the stimulus curves in panel A. (C) Histograms of the connectivity reached after three cycles in the different scenarios extracted from 30 independent depolarizing simulations; mean values and standard deviations are shown in the inset.

Simulation	CPU-time, h	Number of cores	bwHPC Cluster
Run 1	2000	100	NEMO
Run 2	96	1	NEMO

Table 1. We simulated for different length of tDCS stimulation so both single core stimulation and openMPI with multiple cores were used.

4. Conclusions

We explored the plastic changes in network structure that can be induced by transcranial direct current stimulation (tDCS), exploiting the homeostatic response of synaptic growth and decay. We demonstrated that weak subthreshold DC stimulation induces changes of neuronal firing rates and, thus, triggers network remodeling and cell assembly formation. Depolarized neurons first reduce the number of excitatory input synapses during stimulation, but then create new excitatory synapses predominantly with other stimulated neurons after stimulation is off. Interestingly, hyperpolarization also causes new synapses being formed preferentially among



stimulated neurons. Stimulation triggers a profound and sustainable reorganization of network connectivity and leads to the formation of cell assemblies. With the help of our model, we explored different parameters of tDCS stimulation and found that strong and focused stimulation generally enhances the newly formed cell assemblies. We also observed that repetitive stimulation with well-chosen duty cycles can boost the induction of structural changes, and repetitive stimulation with alternating polarization may induce even higher connectivity changes.

5. Acknowledgements

The authors thank Claus Normann, Lukas Frase, Andre Russowsky Brunoni, Sandra Diaz-Pier, and Benjamin Merkt for useful discussions. We also thank Uwe Grauer from the Bernstein Center Freiburg as well as Bernd Wiebelt and Michael Janczyk from the Freiburg University Computing Center for their assistance with HPC applications.

References

[1] Loo, C. K., Alonzo, A., Martin, D., Mitchell, P. B., Galvez, V., Sachdev, P.: Transcranial direct current stimulation for depression: 3-week, randomised, sham-controlled trial. British Journal of Psychiatry (2012) 200(1) 52-59

[2] Nitsche, M. A., Boggio, P. S., Fregni, F., & Pascual-Leone, A.: Treatment of depression with transcranial direct current stimulation (tDCS): A review. Experimental Neurology (2009) 219(1) 14-19

[3] Nitsche, M. A., & Paulus, W.: Excitability changes induced in the human motor cortex by weak transcranial direct current stimulation. Journal of Physiology (2000) 527(3) 633-639

[4] Butz, M., Steenbuck, I. D., & van Ooyen, A.: Homeostatic structural plasticity increases the efficiency of small-world networks. Frontiers in Synaptic Neuroscience (2014) 6

[5] Diaz-Pier, S., Naveau, M., Ostendorf, M., & Morrison, A.: Automatic generation of connectivity for large-scale neuronal network models through structural plasticity. Frontiers in Neuroanatomy (2016) 10

[6] Gallinaro, J. V., & Rotter, S.: Associative properties of structural plasticity based on firing rate homeostasis in recurrent neuronal networks. Scientific Reports (2018) 8(1) 3754



Simulation Foundry: scalable and F.A.I.R. molecular modelling

Gudrun Gygli and Jürgen Pleiss

University of Stuttgart, Institute of Biochemistry and Technical Biochemistry

Abstract

We provide a simulation foundry for the standardised production of molecular dynamics data. Our simulation foundry is a workflow that performs the computational equivalent of a 96 well plate experiment. We have demonstrated the feasibility with binary solvent mixtures of water and methanol at different temperatures, studying the hydrogen bonds, mean square displacement, and diffusion coefficients.

1. Introduction

Quality and reproducibility of data can be ensured by automated and open workflows (the simulation foundry), thus making data findable accessible, interoperable, and reusable (F.A.I.R.). The F.A.I.R. principles support provenance, reproducibility, and open science. Simulation foundries are applied to perform parameter studies for many different molecular systems, conditions, and replicates.

2. Methods

The simulation foundry is organised in two bash scripts, for simulations and for the analysis of trajectories. Upon completion of the analysis, a report is automatically created, containing a Methods section and a Results section. The Methods section includes citations as hyperlinks (Fig.1) and can be included directly in a publication as supplementary information.



2.2 Methods

Molecular dynamics simulations were performed with the GROningen MA-Chine for Chemical Simulation (GROMACS) package (GROMACS 5.1.4gnu-4.9/2016.5, GROMACS1, GROMACS2, GROMACS3, GROMACS4, GRO-MACS5, GROMACS6, GROMACS7), on the high performance infrastructure of the homer. A bash (version 4.4.19) workflow was used to simultaneously launch 2 replicates of each of Parameters-number solvent mixtures of molecules at different temperatures (Temperature, degrees celsius), thus 36 simulations were performed in total. The oplsaa and spce water model were used (oplsaa, spce). Non-standard molecules were parametrized using the LigParGen webserver with the 1.14*CM1A charge model (LigParGen1, LigParGen2, LigParGen3).

A workflow (DOI) was used to automatically perform the ensemble of simulations (simulation foundry). The simulation box was 8 by 8 by 8 nm large, and neutralised by adding 0 None ions.

Minimisation was performed for 50000 steps using the steepest descent algorithm and the steep integrator (DOI, DOI). Equilibration was performed for 10 ns using the md integrator (DOI), the V-rescale thermostat (DOI) with isotropic coupling (tauT=0.1), and electrostatic interactions were calculated with the PME. While equilibration was performed under NPT conditions, production was performed under conditions by adding a Berendsen pressure coupling (DOI, tauP=0.6) for 10 ns. Check the mdp files of the workflow for more parameters.

Fig. 1 Illustration of a preliminary methods section created by the workflow. Hyperlinks are in blue and link to the reference via the DOI. Missing references are given as "DOI" and can be added by the user in the analysis bash script.

3. Results and Discussion

A simulation foundry (an automated workflow) is a complete documentation of the methods, parameters, and data: it assigns provenance to the data and facilitates the reporting of the results of simulation ensembles or parameter studies. Additionally, open and FAIR simulation foundries make it possible to exchange methodologies between groups and therefore allows for an assessment of reliability and reproducibility of a simulation study. It is of the utmost importance that workflows are open, editable, and flexible to support quality control, method validation, and customisation.

The results of the parameter study are presented in a report created by the simulation foundry. A visual representation of the parameters and results provides an overview of the parameter space covered by the simulation study (Fig. 2).

Simulation foundries make molecular dynamics simulations repeatable, replicable, and reproducible.





Fig. 2 Overview of the parameter space covered by simulations with methanol/water mixtures for the property density. Note than simulations were created for multiple mole fractions of water and temperatures, covering a wide range of the parameter space. In total, 176 simulations were performed.

4. Conclusions

The simulation foundry is a technical tool to support scientific work and make molecular dynamics simulation data F.A.I.R. by providing rich metadata to describe the entire process of data creation.


Analysis of large-scale microbiome sequencing datasets on BinAC

Sascha Patz*, Caner Bağcı*, Monika Balvočiūte*, Benjamin Albrecht*, Basil Britto Xavier**, Byoung S. Jeon***, Mohamed El-Hadidi****, Surbhi Malhotra-Kumar**, Largus T. Angenent***, Daniel H. Huson*

* Algorithm in Bioinformatics, Center for Bioinformatics (ZBIT), University of Tübingen, Germany ** Faculty of Medicine and Health Sciences, Laboratory of Medical Microbiology, University of Antwerp, Belgium *** Environmental Biotechnology, University of Tübingen, Germany **** Bioinformatics Research Group, CIS and ITCS, Nile University, Egypt

Abstract

Metagenomic data analysis relies often on a high sample and data size to reduce group variance and improve statistical significance. Such large-scale analysis requires significant computational resources and run time. Thus, high-performance computing is needed. In this context, we use the BinAC cluster for analyzing metagenomic studies and for software development, such as performance optimization of MEGAN6/-LR or of our recently announced software MAIRA for mobile and real-time long-read analysis.

1. Introduction

During the last decades the performance of sequencing technologies has increased dramatically, causing sequencing to be used in many different ways. Microbiologists use sequencing to understand microbial communities and their interactions in their environments and/or hosts. Best-known examples are the Human Microbiome and the Earth Microbiome Projects (HMP, EMP). While 16S rRNA gene based microbial profiling delivered only taxonomic resolutions, modern whole metagenome shotgun and long-read metagenomic sequencing approaches additionally provide functional genomic information¹. Here, the research group of Algorithms in Bioinformatics at the University of Tübingen aims to improve or establish new algorithms and tools for metagenomic analyses. The BinAC cluster is mainly used for alignment-based taxonomic and functional sequence and read annotation, binning and comparison, relying on huge amounts of public available data resources, like the NCBI nr database. Hence, it contributes to metagenomic analyses of medical, environmental and bioreactor microbiomes, and to the optimization and development of metagenomic analysis tools, like MEGAN6/-LR^{2,3} (http://ab.inf.uni-tuebingen.de/software/megan6/), and the recently announced stand-alone application MAIRA⁴ (paper submitted, http://ab.inf.uni-tuebingen.de/data/software/maira/download/welcome.html) for mobile and real-time analysis of long-read sequence data derived from Oxford Nanopore or Pacific Biosciences technology (ONT, PacBio).



2. Projects, Tools and Cluster-Parameters

The following subsections will describe briefly the main research projects that are supported by the use of the BinAC cluster (Fig. 1).



Fig. 1: Collaborations and outcome of the research group Algorithms in Bioinformatics of the University of Tübingen that have benefited most by the BinAC cluster

Software development

MAIRA: The development of MAIRA required the establishment of a reduced taxonomic marker gene database, to reduce computational complexity while maintaining consistent accuracy in assigning metagenomic reads. Starting with all proteins of known bacterial strains (grouped by genus), we aligned all-against-all using DIAMOND⁵ on BinAC. Further, proteins of around 2420 genera were then clustered by their similarity with an in-house Java script and subsequently reduced to obtain genus-and species-specific marker genes, that we stored in respective databases. Both are implemented in MAIRA for fast and accurate taxonomic long-read analysis.

MEGAN6/-LR: Binning of metagenomic reads to taxonomic and functional classes requires regularly updates and improvements of the taxonomic and functional annotation systems, as e.g. implemented in MEGAN6/LR. With the aim to facilitate the functional binning step by enabling metagenomic read classification across different systems, we are currently comparing all proteins stored in InterPro (15.7 million), KEGG (8.5M), EggNOG (6.3M) and SEED (3.6M) databases on the BinAC cluster.



Metagenomics Analyses

Our collaborators provide metagenomic sequence data, that has to be analyzed regarding microbial and functional diversity. Often huge data sets have to be processed in an automatic way. Therefore, we have established an analysis pipeline on BinAC (Fig. 2), that allows submission of multiple jobs in parallel to improve runtime. In the next subsections two currently ongoing projects are presented.



Fig. 2: Example of a metagenomics pipeline running on the BinAC cluster for standard microbial taxonomic and functional analysis

Medical Microbiology: In collaboration with the University of Antwerp⁶ we are investigation the enhancing restoration of the gut microbiome and antibiotic resistance⁷ of clinical patients after selective oropharyngeal decontamination (SOD), selective digestive decontamination (SDD), fecal microbiota transplantation (FMT) in combination with and without antimicrobial compounds. Currently we are analyzing 800 samples containing each in average 10 million reads after host filtering with a median read length of 80 bp according to the metagenomics pipeline on BinAC shown in Figure 2.

Environmental Biotechnology: Here, we are analyzing the waste water bioreactor's microbial and functional diversity of around 50 samples to uncover new microbes and pathways that might be involved in the production of medium chain carboxylic acids (MCCAs)⁸ that are important metabolites and precursors in ketogenic diet and biofuel production. Among short-read analysis, we are applying additionally metagenomics assembly techniques, like MetaVelvet or metaSPAdes (Fig. 2), to reconstruct syntenic metabolic gene clusters and correlate their occurrence to taxonomic annotations under multiple conditions by applying DIAMOND and MEGAN6/-LR.

Computational Resources used on the BinAC cluster

On BinAC we run jobs with parameter settings displayed in Table 1 in parallel within 1-2 days, whereas the same analysis on only one server would take around 50 days.



Sixth bwHPC Symposium, 30th September 2019, Karlsruhe

	5
RESOURCES	APPLIED OPTIONS
bwHPC Cluster	BinAC
CPU-time (average)	1.5h
Number of cores (average)	28
Queue (mostly used)	short
Modules (mostly used)	Java JDK, DIAMOND, BWA, CheckM, FastQC
Local Programs (mostly used)	Trimmomatic, FASTX, MEGAN6/-LR, metaSPADes, MetaVelvet

rad. I: Common parameter setting for lods running on the BinAC cluste	[ab.1:	: Common	parameter	setting	for job	s running	on the	BinAC cluste	r
---	--------	----------	-----------	---------	---------	-----------	--------	---------------------	---

3. Conclusions

The current use of the BinAC cluster in the group Algorithms in Bioinformatics comprises mainly metagenomics analyses and software development. The possibility to run computationally expensive jobs on different nodes in parallel provides significant improvements in runtime and total project working hours. Thus, we were able to establish and improve computational complexity and performance of our software (MEGAN6/-LR, MAIRA), based on e.g. the novel taxonomic marker gene databases and of the metagenomic analyses that we perform in collaboration with our international partners. However, due to increasingly high queuing durations of submitted jobs, we would highly appreciate the extension of nodes and cores within the cluster.

References

[1] Sarhan, MS., Hamza, MA., Youssef, HH., Patz, S., Becker, M., ElSawey, H., Nemr, R.;,Daanaa, HA., Mourad, EF., Morsi, AT., Abdelfadeel, MR., Abbas, MT., Fayez, M., Ruppel, S., Hegazi, NA.: Culturomics of the plant microbiome and the dawn of plant-based culture media – A review. *Journal of Advanced Research* (2019)

[2] Huson, DH., Auch, A., Qi, J., Schuster, SC.: MEGAN Analysis of Metagenomic Data. *Genome Research* 17(3):377-386 (2007)

[3] Huson, D., Albrecht, B., Bagci, C., Bessarab, I., Gorska, A., Jolic, D., Williams, RB.: MEGAN-LR: New algorithms allow accurate binning and easy interactive exploration of metagenomic long reads and contigs. Biology Direct, 13(6) (2018)

[4] Albrecht, B., Bağcı, C., Huson, DH.: MAIRA- real-time taxonomic and functional analysis of long reads on a laptop. (submitted, 2019)

[5] Buchfink, B., Xie, C., Huson, DH.: Fast and sensitive protein alignment using DIAMOND. *Nature Methods* (2014)

[6] Tacconelli, E., Górska, A., De Angelis, G., Lammens, C., Restuccia, G., Schrenzel, J., Huson, DH., Carević, B., Preoţescu, L., Carmeli, Y., Kazma, M., Spanu, T., Carrara, E., Malhotra-Kumar, S., Gladstone, BP.: Estimating the association between antibiotic exposure and colonization with extended-spectrum β -lactamase-producing Gram-negative bacteria using machine learning methods: a multicentre, prospective cohort study. *Clinical Microbiology and Infection* (2019)

[7] Willmann, M., El-Hadidi, M., Huson, DH., Schütz, M., Weidenmaier, C., Autenrieth, IB., Peter, S.: Antibiotic selection pressure determination through sequence-based metagenomics. *Antimicrobial Agents and Chemotherapy* (2015)

[8] Kucek, LA., Spirito, CM., Angenent, LT.: High n-caprylate productivities and specificities from dilute ethanol and acetate: chain elongation with microbiomes to upgrade products from syngas fermentation. *Energy & Environmental Science* 9, 3482–3494 (2016)



Prediction of Protein Stability Changes and Binding Free Energies from Molecular Simulations

Daniel Markthaler^{*} and Niels Hansen^{*}

* Institute of Thermodynamics and Thermal Process Engineering, University of Stuttgart, D-70569 Stuttgart, Germany

1. Introduction

To speed up the design process of new enzyme variants with customized properties such as increased tolerance towards elevated temperatures, theoretical approaches are required which make the reliable prediction of protein stability changes feasible. In rationalized drug design, the precise estimation of binding affinities between a protein target and potential drug candidates is crucial for efficient screening for new therapeutic agents. Molecular **Dynamics** (MD) simulations based on thermodynamically calibrated force fields has been established as a valuable tool in addition to experimental techniques for both kind problems. This work demonstrates the application of atomistic, explicit-solvent MD in combination with state-of-the-art free energy calculation methods for (i) the prediction of protein stability changes due to amino acid mutations [1] and (ii) the calculation of binding free energies [2,3].

2. Applied Methods

Both presented applications rely on different free energy calculation methods, derived from statistical-mechanical principles:

(i) The impact of a particular amino acid mutation onto the overall protein stability is assessed from the relative free energy difference ($\Delta\Delta G$), as obtained from two alchemical free energy calculations - one performed for the native (or folded) state and one for the denatured (or unfolded) state of the protein (c.f. Fig. 1). The representation of the protein's unfolded state requires special considerations. During an alchemical perturbation, the force field parameters of a selected amino acid in state A (wild type) are sequentially transformed into those of the corresponding mutated amino acid in state B (mutant), as controlled by a switching or coupling parameter $\lambda \in [0,1]$.

(ii) The binding free energy (ΔG_{bind}) for a ligand/receptor pair is determined by transferring the unbound ligand through a series of intermediate steps to the binding pose of the receptor [2,3].

To obtain reliable and model-consistent free energy estimates, we develop robust simulation protocols based on enhanced sampling techniques (Umbrella Sampling, Hamiltonian Replica Exchange) and (semi-) automatized workflows.





Fig. 1:

Thermodynamic cycle illustrating the different routes followed in experiment (vertical) and simulation (horizontal) to characterize changes in protein stability due to amino acid mutations [1]. Vertical processes represent the physical unfolding while horizontal processes represent so-called alchemical perturbations. The relative free energy change upon mutating $\Delta\Delta G$ can be equally obtained from both routes.

3. Results and Discussion

Prediction of protein stability:

Comparison of computed protein stability changes ($\Delta\Delta$ G) with experimental denaturation data for a small β -sheet protein shows good overall agreement with only a few outliers (c.f. Fig. 2). An important finding of this work was that subtle local variations in the protein starting structure might lead to substantial deviations in the calculated free-energy changes. For such cases, the usage of multiple starting structures in combination with Hamiltonian Replica Exchange (HRE) was found to be essential to obtain converged free-energy estimates. The quality of the estimated free energy differences depends on the number and the spacing of selected states (characterized by the values $\lambda_i \in [0,1]$) of the transformation amino acid A \rightarrow B. While each λ -state could be run independently, in principle, the usage of HRE requires switching of configurations (based on the Metropolis-Hastings criterion) and therefore communication between the different states. For the considered mutations it was found that 20 equally distributed λ -states (also called "replicas") are sufficient to obtain sufficient phase space overlap with simulation times from 60 to 100 ns per each λ -state (using a MD time step of 0.002 ps) for reaching converged results.

"Typical" job specifications:

- Software: GROMACS 2016
- Cluster: BinAC (short queue)
- 20 replicas distributed over 5 (10) nodes with 4 (2) replicas per node and
 7 cores per replica → 28×5 (10) = 140 (280) cores in total
- ~ 20 ns per day * \rightarrow 3-4 days CPU time for requested simulation time
- Combined MPI / OpenMP parallelization scheme
- GPU acceleration preferable, but typically not profitable due to long queue times of the GPU queue

* Performance can differ significantly, depending on the considered system size and simulation settings



Fig. 2:

Correlation between experimental [4,5] and computed $\Delta\Delta G$ estimates for almost 100 single amino acid mutations of a 34-residue β -sheet protein.



4. Conclusions

Prediction of protein stability:

- It can be expected that the observed conformational sensitivity of the calculated free-energy differences is not a particular issue of the considered system but applies equally to other force fields / mutations / proteins, also in case of rather small proteins such as in the current study.
- Combination of the simulation with HRE is an effective method to speed up convergence but expensive hardware resources required.

5. Acknowledgements

Simulations were performed on the bwHPC computational facilities (BinAC) funded by the Ministry of Science, Research and Arts and the Universities of the State of Baden-Württemberg, Germany, within the framework program bwHPC and the DFG through grant no INST 39/9631 FUGG. We gratefully acknowledge Dr. Jens Krüger for excellent technical support.

References

[1] D. Markthaler, H. Kraus, N. Hansen, J. Chem. Inf. Model. 58 (2018), 2305-2318.

[2] D. Markthaler, J. Gebhardt, S. Jakotorweihen, N. Hansen, *Chem. Ing. Tech.* 89 (**2017**), 1306-1314.

- [3] D. Markthaler, J. Gebhardt, N. Hansen, submitted to Living J. Comput. Mol. Sci. (2019)
- [4] M. Jäger, M. Dendle, J. W. Kelly, Protein Sci. 18 (2009), 1806.
- [5] K. Dave, M. Jäger, H. Nguyen, J. W. Kelly, M. Gruebele, J. Mol. Biol. 428 (2016), 1617.

Simulating ultra-high energy cosmic ray air shower cascades with CORSIKA

Ralf Ulrich, Maximilian Reininghaus Institute for Nuclear Physics Karlsruhe Institute of Technology

September 19, 2019

Abstract

Many fundamental questions in astroparticle-, astro- and particle physics are linked to the research of ultra-high energy cosmic ray air shower cascades. For a detailed analysis the precise and microscopic simulation of such cascades is of paramount importance and requires dedicated algorithms and significant computing resources.

1 Introduction

In astroparticle physics the impact of ultra-high energy cosmic ray particles on the atmosphere of the Earth are a key aspect to understand. The flux of such particles can be extremely low, reaching much below 1 particle per km² per century, however, the produced extensive air shower cascades are huge and extremely energetic [1]. From the observation of such ultra-high energy particles we can learn about the astrophysics of our local universe in terms of high-energy objects, large-scale structure and also magnetic fields [2]. Furthermore, the collisions of such particles with matter in the universe, or in the atmosphere, can be used to study particle physics at energies far beyond what is accessible at accelerators [3]. Precise simulations of air shower cascades are a key tool for this research.

2 Simulating extensive air shower cascades

The worldwide leading tool for air shower simulations is the CORSIKA [4] package developed in Karlsruhe since about 30 years. As input it takes a single ultra-high energy cosmic ray particle with arbitrary impact direction. The code then performs a detailed Monte Carlo simulation of the energy splitting and transport from this single particle to millions or billions of secondary particles produced in collisions with atoms of the atmosphere. The highest energies in this cascade are well beyond what can be achieved at any human-built particle accelerator. Extensive air showers are the only natural source of muons on Earth and are a very efficient generator of so-called *atmospheric neutrinos*. Even small uncertainties in the cascade, in particular at high energies, has the potential to yield large inaccuracies in the final particle distributions.

Cosmic ray particles hit the atmosphere from 10^9 eV (GeV) to $\approx 10^{21} \text{ eV}$ (ZeV). The effort to simulate such events is linear in energy, as well is the resulting output file size. There are essentially two computational boundaries involved: at the lowest energies the computing time per shower is tiny, but the number of typical events in data is huge often requiring billion of simulation to describe the data. At the highest energies the computing time per event becomes excessive making even single events a major challenge: a $10^{18.5} \text{ eV}$ air shower requires about 1.5 years on a single 2 GHz CPU and produced about 5 TB of output. Here only parallel computing in combination with smart algorithms for speed-up yield scientifically useful statistics.

A particular challenge is to simulate Cherenkov emission of charged particles in the air shower. This is typically needed for imaging air Cherenkov telescopes as CTA [5] and increases the computing effort by about 30% and the output storage requirements by a factor of up to 10. Since the statistics required by the experiments can be very large, often the output is not even stored but directly piped into the subsequent workflow steps.

In recent years also the research of radio emission of air showers in the magnetic field of the Earth has become a very important field of research. The simulation of this in CORSIKA doubles the computational effort, however, it does not significantly affect the storage needs.

3 Open questions

Even after decades of progress in air shower physics there are still remaining tensions between observational data and simulations. This is even more pronounced in the era of precise large-scale experiments like the Pierre Auger Observatory [6] or IceCube [7]. It remains a critical aspect of air shower simulations to investigate those tensions and search for new solutions.

4 New developments

It has become evident that future challenges and special needs from the community can no longer be accommodated for by continued efforts spent on extending the 30 year old CORSIKA code written in FORTRAN 77. Thus, development on the entirely new CORSIKA 8 project has started to develop an up-to-date replacement based on modern C++ [8]. With the current state of the framework it is already possible to perform simulations limited to the hadronic and muonic components of air showers and first results are overall comparable to existing codes [9].

We use the resources of the BinAC cluster both for development purposes, and for running larger number of simulations for statistical meaningful studies: Due to large fluctuations among individual showers only values averaged over typically $\mathcal{O}(100)$ showers can be reasonably well compared. Each shower requires up to one day of runtime, depending mainly on the energy of the primary particle as well as the lower energy cutoffs (typically a few GeV) below which particles are discarded. The memory footprint is very small and typically does not exceed a few hundred MB in our setup.

Running these simulations on the BinAC cluster turned out to be very convenient. C++17-compliant compilers are available via the module system, which is a crucial requirement of our project, together with CMake. Furthermore, our jobs are usually scheduled quickly, allowing us to obtain the results in a timeframe limited only by the runtime of the job itself. The possibility to submit multi-core or GPU jobs also plays an important role.

5 Summary

Simulating extensive air showers is a key task in astroparticle physics. The operation of large scale experimental infrastructure and the analysis of collected data depend in a fundamental way on this. To identify the origin of the remaining tensions between data and simulations may point to so far unknown aspects of particle physics.

References

- Alexander Aab et al. Combined fit of spectrum and composition data as measured by the Pierre Auger Observatory. *JCAP*, 1704(04):038, 2017.
 [Erratum: JCAP1803,no.03,E02(2018)].
- [2] Pedro Abreu et al. Bounds on the density of sources of ultra-high energy cosmic rays from the Pierre Auger Observatory. JCAP, 1305(05):009, 2013.
- [3] Pedro Abreu et al. Measurement of the proton-air cross-section at $\sqrt{s} = 57$ TeV with the Pierre Auger Observatory. *Phys. Rev. Lett.*, 109:062002, 2012.
- [4] D. Heck, J. Knapp, J. N. Capdevielle, G. Schatz, and T. Thouw. COR-SIKA: A Monte Carlo code to simulate extensive air showers. 1998.
- [5] M. Actis et al. Design concepts for the Cherenkov Telescope Array CTA: An advanced facility for ground-based high-energy gamma-ray astronomy. *Exper. Astron.*, 32:193–316, 2011.
- [6] Alexander Aab et al. The Pierre Auger Cosmic Ray Observatory. Nucl. Instrum. Meth., A798:172–213, 2015.
- [7] M. G. Aartsen et al. The IceCube Neutrino Observatory: Instrumentation and Online Systems. JINST, 12(03):P03012, 2017.
- [8] Ralph Engel, Dieter Heck, Tim Huege, Tanguy Pierog, Maximilian Reininghaus, Felix Riehn, Ralf Ulrich, Michael Unger, and Darko Veberič. Towards a Next Generation of CORSIKA: A Framework for the Simulation of Particle Cascades in Astroparticle Physics. *Comput. Softw. Big Sci.*, 3(1):2, 2019.
- [9] Diego Melo, Maximilian Reininghaus, Felix Riehn, and Ralf Ulrich. First results of the CORSIKA 8 air shower simulation framework. *Proceedings* of Science, ICRC2019:399, 2019.



Land-atmosphere coupling strength in dependence of the landcover in European climate simulations with WRF

Lisa Jach^{*}, Kirsten Warrach-Sagi^{*} and Volker Wulfmeyer^{*}

*University of Hohenheim, Institute for Physics and Meteorology

1. Introduction

The land surface is an important driver for the climate. Land use and land cover changes (LULCC) modify the surface properties and in consequence may change feedbacks between the land surface and the atmosphere. Besides influencing the amount of carbon stored in soils and biomass, changes in the vegetation type modify the biogeophysical properties of the land surface – namely the leaf area, stomatal conductance and the albedo. In addition, the atmospheric response to modifications at the land surface – meaning a notable change in temperature or precipitation patterns here – depends on the atmospheric conditions themselves. The identification of regions with strong land-atmosphere feedbacks, so called coupling hotspots, improves the assessment of the potential of LULCC to mitigate climate change impacts. Three climate simulations were set-up with differing vegetation in Europe. The aim was to identify regions of strong land cover-precipitation coupling in the summer months and examine the sensitivity of land-atmosphere feedbacks to extreme LULCC.

2. Model and Analyses Methods

The numerical weather prediction model WRF in Version 3.8.1 (Skamarock et al., 2008) coupled to the NOAH-MP land surface model (Niu et al., 2011) was used to simulate the European climate. The three simulations were forced with ECMWF ERA-Interim Reanalysis data (Dee et al., 2011) and run on 0.44° grid resolution. They covered the period 1986-2015. An experiment with the realistic land use map of 2006 served as a baseline to assess the impacts of either extreme afforestation or extreme deforestation on the coupling of the land surface with precipitation in Europe. The coupled land-atmosphere model makes use of the pre-installed netcdf libraries, the intel compiler and OpenMPI. The multinode queue was used for running the experiments usually applying 3 nodes with 72 cores and MPI only. This was done in consideration of the problem size of 130*130*40 cells and the experience that each core should cover about 15*15 cells for the best performance. The singlenode queue was used for the post-processing usually applying one or two cores.

The analyses followed a two-step approach. Firstly, land-atmosphere coupling hotspots were investigated using the 'Convective Triggering Potential' – 'low-level Humidity Index' (CTP-HI_{low}) framework (Findell & Eltahir, 2003). It classifies early morning atmospheric profiles based on the prevailing relative humidity and stability.



Four groups were distinguished: atmospheric conditions allowing for (1) high probability for positive feedbacks (wet soil advantage), (2) high probability for negative feedbacks (dry soil advantage), (3) high probability for feedbacks without a distinct sign (transition zone) and (4) atmospheric conditions barely allowing for feedbacks (atmospheric control). Secondly, the connection between the turbulent surface fluxes and cloud and precipitation patterns was examined using a modified linear regression approach.

3. Scientific Inside

In this study on land-atmosphere coupling strength, one land-atmosphere coupling hotspot was identified in the north-east of Europe. The region of strong coupling as well as the predominantly reinforcing nature of the feedbacks appeared insensitive to LULCC. However, the magnitude of feedbacks was reduced by afforestation as well as deforestation compared to the baseline. Afforestation caused a net drying of the atmosphere, what is reflected in an average increase of the humidity deficit in the lowlevel atmosphere (Hl_{low} +1.0°C) in the coupling hotspot. Less frequent favoring conditions for feedbacks weaken the coupling strength and decrease the probability for deep convection triggering by the land surface. Conversely, deforestation initiated a more humid climate (Hlow -0.75°C). The higher frequency of days with near surface humidity close to saturation reduced the probability for convection triggering by the land surface. This is because the probability for precipitation is generally high with high surface humidity and the transport of moisture and energy barely impacts the formation of precipitation. Thus, the results indicate that, especially in strongly coupled regions, LULCC initiate changes in the atmospheric background conditions which modify the strength on land-atmosphere feedbacks in the long run.

The conducted simulations are additionally analyzed in the scope of model intercomparison studies within the Flagship Pilot Study (FPS) 'Land-Use and Climate Across Scales' (LUCAS)(Davin et al., 2019). The FPS was launched with the aim of improving the understanding of biogeophysical LULCC impacts on the climate at local and regional scales, in 2017. It is hosted by the European branch of the Coordinated Regional Climate Downscaling Experiments (EURO-CORDEX) community and endorsed by the World Climate Research Program.

4. Acknowledgements

The research of this study was funded by the Anton and Petra Ehrmann-Stiftung Research Training Group "Water-People-Agriculture".

References

- [1] Davin, E. L., Rechid, D., Breil, M., Cardoso, R. M., Coppola, E., Hoffmann, P., et al. (2019). Biogeophysical impacts of forestation in Europe: First results from the LUCAS Regional Climate Model intercomparison. *Earth System Dynamics Discussions*, 1–31. https://doi.org/10.5194/esd-2019-4
- [2] Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., et al. (2011). The ERA-Interim reanalysis: configuration and performance of the data assimilation system. *Quarterly Journal of the Royal Meteorological Society*, *137*(656), 553–597. https://doi.org/10.1002/qj.828
- [3] Findell, K. L., & Eltahir, E. A. B. (2003). Atmospheric Controls on Soil Moisture–Boundary Layer Interactions. Part I: Framework Development. *Journal of Hydrometeorology*, 4(3), 552–569. https://doi.org/10.1175/1525-7541(2003)004<0552:ACOSML>2.0.CO;2
- [4] Niu, G.-Y., Yang, Z.-L., Mitchell, K. E., Chen, F., Ek, M. B., Barlage, M., et al. (2011). The community Noah land surface model with multiparameterization options (Noah-MP):
 1. Model description and evaluation with local-scale measurements. *Journal of Geophysical Research*, *116*(D12). https://doi.org/10.1029/2010JD015139
- [5] Skamarock, W., Klemp, J., Dudhia, J., Gill, D., Barker, D., Wang, W., et al. (2008). A Description of the Advanced Research WRF Version 3. UCAR/NCAR. https://doi.org/10.5065/D68S4MVH

A coupled bidomain solver with optimal memory usage for the numerical simulation of cardiac reentry

Nagaiah Chamakuri & Philipp Kügler Institute of Applied Mathematics and Statistics, University of Hohenheim, Stuttgart

September 23, 2019

Abstract

The bidomain model describes the electrical activity of the heart. Due to the sheer size of state variables, the numerical approaches require enormous memory requirements and computational times in 2D or 3D. In this talk, we will present the techniques to tackle these challenges and suggest a novel computational strategy, that exploits the sparsity of local matrices in the assembly of global FEM matrices. We demonstrate the practicability of our coupled approach by employing three popular physiological cell models and compare it with a commonly used decoupled strategy.

1 Introduction

The electrical behavior of the cardiac tissue is described by a system consisting of partial differential equations coupled with ordinary differential equations which model the ionic currents associated with the reaction terms. These are the so called bidomain model equations [8, 9]. The governing equations of the bidomain model are represented by the following equations.

$$0 = \nabla \cdot (\bar{\sigma}_i + \bar{\sigma}_e) \nabla u + \nabla \cdot \bar{\sigma}_i \nabla v \quad \text{in } Q \tag{1}$$

$$\frac{\partial v}{\partial t} = \nabla \cdot \bar{\sigma}_i \nabla v + \nabla \cdot \bar{\sigma}_i \nabla u - I_{ion}(v, \vec{w}, \vec{c}) + I_{tr}(x, t) \quad \text{in } Q \qquad (2)$$

$$\frac{\partial c}{\partial t} = S(v, \vec{w}, \vec{c}) \quad \text{in } Q, \tag{3}$$

$$\frac{\partial \vec{w}}{\partial t} = G(v, \vec{w}) \quad \text{in } Q, \tag{4}$$

Here $u: Q \to \mathbb{R}$ is the extracellular potential, $v: Q \to \mathbb{R}$ is the transmembrane voltage, $\vec{w} \colon Q \to \mathbb{R}^n$ represents the gating variables, $\vec{c} \colon Q \to \mathbb{R}^m$ represents the ionic concentration variables, $\bar{\sigma}_i, \bar{\sigma}_e \colon \Omega \to \mathbb{R}^{2 \times 2}$ are respectively the intracellular and extracellular conductivity tensors. The term I_{tr} is the transmembrane current density stimulus as delivered by the intracellular electrode. The I_{ion} is the current density flowing through the ionic channels. The function S and G determines the evolution of the concentrations and gating variables respectively which are determined by an electrophysiological cell model, see [1] for more description on these models. In our code development, several ionic models were implemented from simple phenomenological models to real physiological models. In this work, we will present results on Luo-Rudy phase-I (LRI) cell model [6] which describes the membrane action potential of the mammalian ventricular cell and O'Hara-Rudy (ORd) [7] model which is based on data from human ventricular myocyte experiments. The above mentioned Eq. (1) is an elliptic type equation, Eq. (2) is a parabolic type equation and Eq. (4) is a set of ordinary differential equations which needs to be resolved in each spatial point of the computational domain. We refer the readers to see [4] for more details of boundary and initial conditions.

2 Methods and Discussion

We used a piecewise linear conforming finite element method for the spatial discretization of the bidomain equations and the time discretization is performed using Rosenbrock method, *ROS3PL* method [5] which has four internal stages to solve in each time step. For further details we refer to [4]. In this work, we suggest a novel coupled strategy for solving the bidomain model equations that is based on an efficient assembly of the global compressed row storage (CRS) matrix of the FE discretization. The main idea is to exploit the sparse coupling of the cellular state variables and to implement the local block matrices appearing at each CRS element as compile-time sparse matrices (CTSM) rather than dense matrix (DM). In particular, this enables the reduction of computer memory as well as the accurate and efficient simulation of pathological behaviour such as cardiac reentry. In case of LRI model, with 8 state variables, we realize that the storage of the local block matrix entries only requires one dimensional array of 24 doubles instead of 64. The sparsity of the local block matrix is 62%. Furthermore, only 24 floating point operations are needed at each nodal point during the matrix-vector product operations of the Krylov solvers. For the ORd model, with its 41 state variables, the sparsity of the local block matrix is 86%. Moreover, the global FE matrix for the monodomain equation with the ORd model defined on a 2D slab geometry with 848421 FE cells requires about 7.5 times less memory if build according to the CTSM approach, see [4] for more details.

We employed a BiCGSTAB method with block Jacobi preconditioning to solve the system of equations which appears after the time discretization of parabolic-ODE equations. The obtained algebraic system after full spatial discretization of elliptic (1) is solved by the conjugate gradient (CG) method with AMG preconditioner [2]. For our numerical experiments, we have developed the user module which is based on the public domain software package DUNE [2] and Dune-PDELab [3]. We use a non-overlapping decomposition for the



mesh partitioning, based on the internal YASP Figure 1: Weak scaling grid in Dune [2], in our implementation. The of LRI model. computations were performed on the bwUniClus-

ter (located at KIT, Germany) using up to 64 Intel Xeon E5-2660 v4 (Broadwell) computing nodes, where each node consists of 28 cores and 128 GB main memory. In our numerical experiments for the LRI model, the matrix assembly with CTSM is approximately 1.20 times faster than with DM across different core numbers, and the total computation time is about 1.50 shorter. In case of ORd model, the total CPU gain is 3.7. Furthermore, good parallel efficiency was observed up to 1792 cores, see Figure 1.

Acknowledgments

We greatly acknowledge the support for high-performance computing time at the bwUniCluster which is funded by the ministry of science, research and arts of the state of Baden-Württemberg, Germany.

References

- [1] CellML Model Repository. http://models.cellml.org/cellml.
- [2] P. Bastian, M. Blatt, A. Dedner, C. Engwer, R. Klöfkorn, R. Kornhuber, M. Ohlberger, and O. Sander. A generic grid interface for parallel and adaptive scientific computing. Part II: implementation and tests in DUNE. *Computing*, 82(2):121–138, July 2008.
- [3] P. Bastian, F. Heimann, and S. Marnach. Generic implementation of finite element methods in the distributed and unified numerics environment (DUNE). *Kybernetika*, 46(2):294–315, 2010.
- [4] N. Chamakuri and P. Kügler. A coupled bidomain solver with optimal memory usage for the simulation of cardiac reentry. submitted.
- [5] J. Lang and D. Teleaga. Towards a fully space-time adaptive FEM for magnetoquasistatics. *IEEE Transactions on Magnetics*, 44(6):1238–1241, 2008.
- [6] C. Luo and Y. Rudy. A model of the ventricular cardiac action potential: Depolarization, repolarization, and their interaction. *Circ. Res.*, 68:1501– 1526, 1991.
- [7] T. O'Hara, L. Virág, A. Varró, and Y. Rudy. Simulation of the undiseased human cardiac ventricular action potential: Model formulation and experimental validation. *PLOS Computational Biology*, 7(5):1–29, 05 2011.
- [8] R. Plonsey. Bioelectric sources arising in excitable fibers (ALZA lecture). Ann Biomed Eng, 16(6):519–46, 1988.
- [9] L. Tung. A bi-domain model for describing ischemic myocardial DC potentials. PhD thesis, MIT, Cambridge, MA, 1978.

Atomistic simulation of surface and interface energies in a nickelbased SOFC anode

Anika Marusczyk*, Matthias Kappeler*, Senja Ramakers*, Patricia Haremski*, Matthias Wieler*, Piero Lupetin*

*Robert Bosch GmbH, Corporate Sector Research and Advance Engineering, Robert-Bosch-Campus 1, 71272 Renningen

1. Introduction

The degradation mechanism of Ni-based solid oxide fuel cell (SOFC) anodes under operating conditions is dominated by an unfavorable microstructural evolution that involves the coarsening of nickel grains. The resulting reduction of the active triplephase boundary leads to a significant loss in performance. Within the publicly funded project Kersolife100, we follow a combined simulative and experimental multi-scale approach to predict the microstructural evolution, and therefore degradation, of the anode while avoiding time-consuming long-term experiments. To generate a fundamental understanding of the relevant mechanisms, all applicable material parameters, such as interfacial energies and diffusion constants, are determined and exploited in large-scale 3D phase-field simulations. In this work, we report on the determination of surface and interface energies of the nickel-based SOFC anode using atomistic simulation methods.

2. Computational Methods

All of our density functional theory calculations were performed using the Vienna Ab initio Simulation Package (VASP) [1-4] as compiled on the ForHLR I and conducted on the multinode queue. To model the effects of exchange and correlation, we used three different exchange-correlation functionals based on the local density approximation (LDA) [5], the generalized gradient approximation (GGA) under the Perdew-Burke-Ernzerhof parameterization (PBE) [6, 7] as well as the Perdew-Burke-Ernzerhof parameterization for solids and surfaces (PBEsol) [8]. To ensure high accuracy, the relevant convergence parameters were carefully investigated for the fcc nickel surface energies with low Miller index orientations as well as low Σ grain boundary structures.

In addition to this, we investigated the applicability of several nickel/nickel-hydrogen empirical potentials that are freely available in the OpenKim database [9]. To this end, we compared the results generated by LAMMPS [10] to our accurate density functional theory values for each material property and chose the most appropriate empirical potential, if possible. Given the drastically decreased limitations on the number of atoms compared to density functional theory, we were then able to confirm cell size convergence of our previous results and to calculate the energies of higher Σ grain boundaries. To search for the global minimum of the grain boundary

structures, we modeled annealing by first equilibrating to 800K, then annealing to 0.1K and then using a conjugate gradient algorithm to find the energy minimum.

For the job management, the Automated Interactive Infrastructure and Database for Computational Science (Aiida) [11] software was used in combination with Custodian [12]. The crystal structures were generated with the Atomic Simulation Environment (ASE) [13] as well as the Python Materials Genomics (pymatgen) [12] and Vesta [14] was used for the visualization.

3. Summary of Results and Conclusions

In this work, we present the computation of the relevant surface energies of differently oriented nickel surfaces up to Miller index three using density functional theory with three different exchange-correlation functionals. All of our results are in agreement with previously computed values derived from LDA and PBE functionals and lie within the expected experimental scatter. PBEsol was chosen as the most accurate functional for the calculation of surface energies. Among our investigated empirical potentials, some were able to satisfactorily reproduce our density functional theory weighted surface energies, but not the anisotropy.

As another relevant mechanism in an SOFC anode, we investigated the adsorption of nickel as well as hydrogen onto differently oriented nickel surfaces and derived adsorption energies that are in agreement with known literature values when available. Encouragingly, nickel adsorption energies are well reproduced by our employed empirical potentials, which can therefore be used for more extensive investigations. There were, however, significant discrepancies between our density functional theory results and those derived from the used nickel-hydrogen empirical potential, which will therefore not be further considered for the simulation of hydrogen adsorption. The derived models will be used for the simulation of diffusion on the nickel surface, which is an ongoing work package.

In addition to this, we determined nickel-nickel grain boundary energies using density functional theory for the small \sum structures, which were then used to benchmark several empirical potentials. As the agreement was highly satisfactory, we were able to conduct an extensive study on larger \sum structures and determine the relevant grain boundary energies. We then began the investigation of the relevant nickel-electrolyte interface energies, which will be presented on the next occasion.

Typical run times for the described density functional theory calculations can be found in Tab. 1. As these are strongly dependent on the used convergence parameters as well as cell size, they should be understood as an approximated average.



Sixth bwHPC Symposium, 30th September 2019, Karlsruhe

Simulation	CPU-time, h	Number of	bwHPC Cluster
		cores	
Surface energy (low Miller)	18	80-120	ForHLR I
Surface energy (high Miller)	140	120	ForHLR I
Nickel adsorption	120	80-120	ForHLR I
Hydrogen adsorption	70	80-120	ForHLR I
Grain boundary energy	50-100	80-140	ForHLR I

Table 1. Approximate average use of computational resources for the indicated work

 package

4. Acknowledgements

This work was performed on the computational resource ForHLR I funded by the Ministry of Science, Research and the Arts Baden-Württemberg and DFG ("Deutsche Forschungsgemeinschaft") and conducted within the project "Kersolife100", funded by the Federal Ministry for Economic Affairs and Energy (03ET6101A).

References

[1] G. Kress and J. Furthmüller. Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set. Computational Materials Science, 6(1):15-50, 1996

[2] G. Kresse and J. Furthmüller. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. Phys. Rev. B, 54:11169-11186, 1996

[3] G. Kresse and J. Hafner. Ab initio molecular-dynamics simulation for liquid metals. Phys. Rev. B, 47:558-561, 1993

[4] G. Kresse and J. Hafner. Ab initio molecular-dynamics simulation of the liquidmetal-amorphous-semiconductor transition in germanium. Phys. Rev. B, 49:14251-14269, 1994

[5] J. P. Perdew and A. Zunger. Self-interaction correction to density-functional approximations for many-electron systems. Phys. Rev. B, 59:1758. 1999

[6] J. P Perdew, K. Burke and M. Ernzerhof. Generalized Gradient Approximation Made Simple. Phys. Rev. Lett. 77, 3865. 1996

[7] J. P. Pewdew, K. Burke and M. Ernzerhof. Erratum: Generalized gradient approximation made simple. Phys. Rev. Lett. 78, 1396. 1997

[8] J. P Perdew, A. Ruzsinszky, G. I. Csonka, O. A. Vydrov, G. E. Scuseria, L. A. Constantin, X. Zhou and K Burke. Restoring the density-gradient expansion for exchange in solids and surfaces. Phys. Rev. Lett. 100:136406. 2008



[9] E. B. Tadmor, R. S. Elliott, J. P. Sethna, R. E. Miller and C. A. Becker, "The Potential of Atomistic Simulations and the Knowledgebase of Interatomic Models" *JOM*, **63**, 17 (2011)

[10] Steve Plimpton. Fast parallel algorithms for short-range molecular dynamics. Journal of Computational Physics, 117(1):1-19, (1995).

[11] G. Pizzi, A. Cepellotti, R. Sabatini, N. Marzari and B. Kozinsky. Aiida: automated interactive infrastructure and database for computational science. Computational Materials Science, 111:218-230, 2016

[12] S. P. Ong, W. Davidson Richards, A. Jain, G. Hautier, M. Kocher, S. Cholia, D. Gunter, V. L. Chevrier, K. A. Persson and G. Ceder. Python materials genomics (pymatgen): A robust, open-source python library for materials analysis. Computional Materials Science, 68:314-319, 2013

[13] A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Duak, J. Friis, M. n. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. B. Jensen, J. Kermode, J. R. Kitchin, E. L. Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J. Bermann Maronsson, R. Maxon, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schitz, O. Schtt, M. Strange, K. S. Thygesen, T. Vegee, L. Vilhelmsen, M. Walter, Z. Zheng and K. W. Jacobsen. The atomic simulation environment python library for working with atoms. Journal of Physics: Condensed Matter, 29(27):273002, 2017

[14] K. Momma and F. Izumi. VESTA3 for three-dimensional visualization of crystal, volumetric and morphology data. Journal of Applied Crystallography, 44(6):1272-1276. 2011.



Optimization of the Charge Motion of a Lean Burn Gas Engine: Grid-Independence Study, Validation and Piston Geometry Optimization

Ghael, Kamlesh*; Rosenthal, Felix**; Kubach, Heiko*; Blesinger, Georg*; Koch, Thomas*

*Karlsruhe Institute of Technology, Institute of Internal Combustion Engines (IFKM)

**ex-Research Assistant at Karlsruhe Institute of Technology, Institute of Internal Combustion Engines

Abstract

In comparison to diesel or gasoline engines, gas engines for power generation or stationary applications can contribute to the goal to reduce specific CO₂ emissions due to higher H/C ratio. Nevertheless, strict emission limits regarding NO_x and formaldehyde emissions have to be met. Therefore gas engines have to be redesigned for lower equivalence ratios down to 0.6 ($\lambda \ge 1.7$). However, with leaner mixtures, combustion tends to be incomplete and slow, thus result in more CO, formaldehyde, and unburnt CH₄. But to effectively use the lean-burn gas engines, its design and the charge motion must be optimized. In this paper, a 3D-CFD-optimization study of the in-cylinder swirl and turbulent kinetic energy (TKE) is performed by varying geometry of the piston bowl. Optimization is conducted based on an initial grid independence study and model validation of the open cycle at motored operation for the reference geometry with experimental data. The result shows the increment of the swirl, squish and TKE in the cylinder at firing top dead center (fTDC). It is also investigated, that with the help of bwHPC cluster (ForHLR) the multinode simulation is possible which results in lesser simulation time.

Optimizing Universality for Air Shower Reconstructions

Johannes Hulsman Institut für Kernphysik Karlsruhe Institute of Technology

September 18, 2019

1 Scientific Motivation

Cosmic rays are carrier of information of our universe, continually bombarding the earths atmosphere since its discovery in 1912 by Victor Hess. At the top of the atmosphere, they are found to be mainly atomic nuclei and only about 2% electrons/positrons. Their composition, charge and energy provide invaluable information on the particle production, propagation and interaction on a galactic and extra-galactic scale.

The most abundant cosmic rays particles have a low energy and are thus easily affected by magnetic fields and more susceptible to interactions with the interstellar medium. A very small fraction, yet very interesting one, extend to ultra-relativistic energies of 10^{20} eV (several orders of magnitude above LHC energies). Their origin and acceleration source remain unknown -a fundamental question in the astrophysical community. As these particles are extremely rare, they must be studied by large air shower arrays, exposed for many years on the surface. Here, the atmosphere serves as a calorimeter where one infers the nature of the primary from its secondary particle cascade (commonly referred to as *extensive air showers*).

The Pierre Auger Observatory is the largest cosmic ray observatory in the world. It was conceived to measure and study the characteristics and interactions of extensive air showers for primary energies above 10^{17} eV with a hybrid detection system. It spans 3000 km^2 over a vast plane in Western Argentina. The principle design features of the Observatory is the more than 1600 water Cherenkov particle detectors (WCD) separated by 1500 m in a triangular grid array. It has been collecting data since 2004. As part of the upgrade initiative to the observatory, three additional detector systems are being developed; muon detectors (MD) and scintillating surface detectors (SSD). These detectors will be associated to each WCD (together referred to as a *station*).

The foundations of reconstructing extensive air shower are the fit to the geometry and lateral distribution function (LDF). Air shower universality is based on a phenomenological method which encapsulates the underlying shower physics from singal and time information and allows for a reconstruction based on mass-composition sensitive shower parameters -among those being the shower maximum X_{max} , maximum of muon production depth X^{μ}_{max} and relative muon content R_{μ} . It assumes that shower secondaries have the same fractional rate of change with increasing depth at the same shower age. In other words, the normalized shape of the longitudinal shower profile is expected to be equivalent for all showers.

2 Universality Framework

As the principles of air shower universality rely on the *average* signal [1] and time distribution [2], a special framework was developed to specially process the vast amount of secondary particles and condense with all the relevant physics. Extensive air showers are simulated with CORSIKA and then read by the Universality framework. In figure 1a a brief schematic can be found. It shows that the framework assumes a circular array with different radii. Every marker represents a *sampling area*, which accounts for a WCD, SSD and MD. Each detector stores information on shower geometry, shower physics, detector electronics, detector signal and trace. The latter is divided into four components to emphasize the underlying physics seen between muon, electromagnetic particles from muon decays, pure electromagnetic particles and electromagnetic particles from hadronic jets.



Figure 1: **a**) A brief schematic of the Universality framework. **b**) Schematic describing how the detector response is obtained. Individual detector responses for muons, electrons and photons (depending on their injected energy and angle) are simulated. These are stored in tables and called upon by particles from CORSIKA showers.

Conventionally, particles are injected into a WCD/SSD/MD to simulate a signal with GEANT4. This is a highly time-consuming strategy, requiring months (assuming a personal computer) to simulate a single low energy shower at 10¹⁷ eV. A special GEANT4-table library had been developed to pre-simulate the particle signal response in the detector. Signal responses of secondary particles from CORSIKA showers are interpolated with the table. This reduces the computation time significantly as particle of the same type, energy and direction do not need to be simulated more than once. However, this approach does not account for non-linearities in detector signals, which may introduce biases for detectors with large signals. The production of these GEANT4-tables could be completed at the IKP-computing cluster within two weeks (assuming low traffic). A schematic is shown in figure 1b.

3 Results and Discussion

Despite many optimizations, the Universality framework has proven to be very memory intensive as illustrated in figure 2a. The memory usage depends on the on the number of particles, which depends on particle type and energy. At 10^{17} eV, a typical air shower consumes 3GB of memory, whereas a shower at 10^{20} eV consumes 8GB. The latter is far above the 2GB limit imposed by the IKP-computing cluster (simulations above the limit "swap memory").

Using the HPC resource bwUniCluster, the model development for Universality could be completed. A sample shower is shown in 2b. The framework's expansion, detailed work in air shower physics and usage with the HPC has also paved the way towards a major upgrade -dubbed 'Universality v2'. The most significant feature of the new framework is its modular design, addressing the memory consumption seen in figure 2a.



Figure 2: a) Memory consumption of an extensive air shower at 10^{17} eV. b) Sample hybrid reconstruction of a $10^{19.5}$ eV proton air shower.

References

- M. Ave, R. Engel, M. Roth, and A. Schulz. A generalized description of the signal size in extensive air shower detectors and its applications. *Astropart. Phys.*, 87:23 – 39, 2017.
- [2] M. Ave, M. Roth, and A. Schulz. A generalized description of the time dependent signals in extensive air shower detectors and its applications. *Astropart. Phys.*, 88:46 – 59, 2017.



LARGE-EDDY SIMULATION OF TURBULENT HEAT TRANSFER IN A HELICALLY RIB-ROUGHENED PIPE

Kevin Akermann^{*} and Peter Renze^{*}

*University of Applied Sciences Ulm, Germany Institute for Energy and Drive Technology

Abstract

Numerical investigations of turbulent heat transfer in pipes with an inner helically ribbed surface have been carried out by performing Large-Eddy Simulations (LES) for Re = 8000, 16000, 32000 and Pr = 7. The simulation results were compared to measurement data of Wieland-Werke AG for validation purposes. The deviation between simulation and measurement lies within 8% in heat transfer (Nusselt number) and 10% in pressure loss.

1. Introduction

Pipes with a rough surface on the inside enhance the heat transfer due to an increased surface and flow separation. However, with internal roughness the pressure loss will be also increased as a consequence of the higher skin friction [4]. At the present time, the development of new structured pipes is mostly done by experimentally determined data due to the complexity of heat transfer in internally roughened pipes. Pressure loss as well as heat transfer strongly depend on geometrical parameters such as the roughness height *e*, the pitch of the roughness *p*, the helical angle α and the contact angle β . A sketch of a helically ribbed pipe and different profile shapes are illustrated in figure 1.



Fig. 1 Sketch of a helically ribbed pipe and profile shapes, cf. [2]

Detailed numerical simulations can be used, to support the development process of new, improved structured pipes. In order to perform detailed numerical simulations high resolution CFD-simulations and high computing power are required.



In the scope of the present work, Large-Eddy Simulations of turbulent flow in a pipe with a helically ribbed surface (p/e = 6, e/D = 0.028, $\alpha = 55^{\circ}$ and $\beta = 25^{\circ}$) were performed for several Reynolds numbers and Pr = 7. The chosen numerical approach was validated and verified by preliminary examinations of smooth pipe flow simulations. In order to validate the ribbed pipe flow simulations, measurements of heat transfer and pressure loss of Wieland-Werke AG have been used.

2. Numerical Method

The simulations were performed using the solver *buoyantPimpleFoam* of the opensource software library OpenFOAM®. Physical modelling is based on conservation equations of continuity, momentum and energy, which are derived in detail in Anderson [1]. To take turbulence into account, the LES method with the 0-equation WALE [3] model was used.

The computational domain has a length L of 4D, where D is the diameter. Periodic (cyclic) boundary conditions were applied in order to get a fully developed turbulent pipe flow. A schematic representation of the domain is shown in figure 2.





A fine resolution and a high number of mesh cells are required to ensure a well resolved LES and to build the structured surface properly. In order to avoid high numerical costs, the mesh remained unchanged for the investigated Reynolds numbers of 8000, 16000 and 32000, except the height and number of the layers were adapted to the Reynolds number. Thus, the first grid point was located far inside within the viscous sublayer for each LES, $y^+ \leq 0.1$. The simulations were performed on the bwUniCluster due to the high number of mesh cells. Summarized information about the mesh sizes and the CPU-usage are given in table 1.



Sixth bwHPC Symposium, 30th September 2019, Karlsruhe

Simulation	Mesh Cells ×10 ⁶	Number of cores	Simulated time	CPU-time, h
<i>Re</i> = 8000	23.9	224	~2 sec	~55000
<i>Re</i> = 16000	26.5	280	~1 sec	~64000
<i>Re</i> = 32000	28.1	280	~0.5 sec	~75000

Table 1 Information about the mesh sizes and HPC-usage

3. Results and Discussion

In figure 3 the instantaneous velocity and temperature field for Re = 8000 in the helically ribbed pipe are shown. The laminar and the thermal sublayer in near wall region as well as the vortical, coherent structures can be clearly seen in those snapshots. The fluid motion is from left to right.



Fig. 3 Instantaneous velocity (left) and temperature (right) field of the helically ribbed pipe, Re = 8000

The comparison of the Nusselt number (*Nu*) and the pressure loss per length ($\Delta p/L$) with the measurement data of Wieland-Werke AG are shown in figure 4a) and 4b). It can be seen, that the simulation results deviate from the measurements by less than 8% in heat transfer and 10% in pressure loss for all Reynolds numbers. Furthermore can be noticed, that the deviations in pressure loss increase from 2% with *Re* = 8000 up to 10% with *Re* = 32000. The deviation in heat transfer is 4% and 7.8% for *Re* = 8000 and 32000, respectively. This increase in the percentage deviations may occur due to the grid resolution, which might be to coarse for high Reynolds numbers.

4. Conclusions

Large-Eddy Simulations of turbulent heat transfer in a helically ribbed pipe were performed for several Reynolds numbers. In order to avoid prohibitive numerical costs, the mesh sizes almost remained unchanged. The comparison of the results and



Sixth bwHPC Symposium, 30th September 2019, Karlsruhe



Fig. 4 Comparison of the Results of the present LES and measurements of **a**) Nusselt number and **b**) pressure loss

measurement data showed, that the deviation of LES and measurements is in an acceptable tolerance for all investigated Reynolds numbers. However, the deviation increases with higher Reynolds number, which could be the result of the constant grid resolution.

The next steps will be to apply the chosen numerical approach to new, enhanced pipes and to improve and optimize the structured surface in terms of heat transfer and pressure loss.

5. Acknowledgements

The authors like to gratefully acknowledge support by the Steinbuch Center for Computing (SCC) at the Karlsruhe Institute of Technology (KIT), the project InnoSÜD for financial support and Wieland-Werke AG for providing geometry and measurement data.

References

[1] J. Anderson: Computational Fluid Dynamics, McGraw-Hill, Inc.: New York, USA, 1995

[2] Gee, D.; Webb, R.: Forced convection heat transfer in helically rib-roughened tubes, Int. J. Heat and Mass Transfer, 23:1127-1135, 1980

[3] Nicoud, F.; Ducros, F.: Subgrid-scale stress modelling based on the square of the velocity gradient tensor, Flow. Turb. and Comb., 62:183–200, 1999.

[4] Webb, R.; Kim, N.: Principles of enhanced heat transfer, 2nd edn., Chapter 9: Integral Roughness, Taylor and Francis Group, 2005

bw HPC

WORKFLOW MANAGEMENT FOR PARAMETRIC AFM STUDY ON SURFACTANT ADSORPTION FILM BY MOLECULAR DYNAMICS

Johannes Hörmann, Bryan Ramos and Lars Pastewka

University of Freiburg, Department of Microsystems Engineering (IMTEK)

Abstract

On the example of two model surfactants, anionic sodium dodecyl sulfate (SDS) and cationic cetyltrimethylammonium bromide (CTAB), at the idealized water–gold interface, our present simulation work illuminates the interdependence of adsorption film morphology and its mechanical properties on the nanoscopic scale. The workflow management framework *FireWorks* facilitates systematic exploration of a high-dimensional parametric space.

1. Introduction

In aqueous solution, amphiphilic surfactants adsorb on immersed surfaces. Resulting films exhibit morphology-dependent mechanical properties, such as nanoscopic indentation response as well as macroscopic lubrication performance [1]. Our work focuses on two model surfactants, SDS and CTAB, at the





H₂O - Au (111) interface. Both are known to form flat-lying monolayers at low concentrations. With increasing surface coverage, stripe-like aggregates assemble – hemicylinders in the case of SDS [2] and cylinders for CTAB [3]. We utilize existing SDS and CTAB parametrizations to simulate atomic force microscopy (AFM) and friction force microscopy (FFM) experiments by means of classical all-atom molecular dynamics. The workflow management framework *FireWorks* [4] allows systematic exploration of a parameter space spanned by surfactant type, concentration, film morphology, as well as AFM probe approach or sliding velocity and direction amongst other dimensions.



Sixth bwHPC Symposium, 30th September 2019, Karlsruhe

Figure 1 displays the experimental adsorption isotherms of SDS and CTAB on gold at open circuit potential. Measured adsorption masses have been converted to surface number concentrations on the left hand axis. Marked data points along the curves indicate our subset of representative configurations, labeled with corresponding numbers of adsorbed molecules at 15² nm² and 18.4² nm² Au(111) surfaces for SDS and CTAB respectively. Inlets illustrate experimentally expected film morphologies at selected data points and evolved snapshots of accordingly preassembled initial configurations after 10 ns in the isothermal-isobaric ensemble.

2. Method

Parameters primarily originate from CHARMM 36 [6] and INTERFACE FF [7]. Compatible bromide ion [8] and CTAB [9] models are utilized. An embedded atom method potential by Grochola [10] describes the Au-Au interaction. After minimization, NVT-, and NPT-equilibration, all systems evolve for 10 ns, with only the substrate Langevin-tempered in otherwise pressurized environment. Following the insertion of a 5 nm diameter AFM probe model and repeated minimization, NVTand NPT-equilibration, substrate and probe are tempered by a dissipative parti-





cle dynamics thermostat while holding the system volume fixed. After 0.5 ns, the probe model's frozen core of 2.4 nm diameter is instantaneously set into motion at prescribed velocity, approaching the substrate with its bottommost 1.4 nm layer immobile. Eventually, a discrete subset of evenly distanced snapshots along the approach trajectory is employed as initial configurations for lateral sliding. Preceded by another 200 ps equilibration, the probe now moves in tangential direction. Figure 2 illustrates this procedure, indicating the process pipeline's branching for different morphologies, velocities and distances. All calculations apply 3D-periodic boundary conditions and particle-particle particle-mesh Ewald summation treatment of long interactions.

While all molecular dynamics production runs are carried out via *LAMMPS*, a multitude of academic software is employed for intermediate processing. *FireWorks* [4] helps to integrate diverse steps into seamless work flow pipelines and administer these as well as their data throughput in a comfortably searchable NoSQL database, thereby tracking provenance. Depending on resource demand and availability, tasks are distributed across heterogeneous infrastructure at hand. Jobs of moderate to in-



tense resource requirements (signaled by yellow and red colors in Figure 2) are predominantly processed on *bwForCluster NEMO*, while other low cost tasks are dispatched to *bwCloud SCOPE*.

Previous benchmarks have shown *LAMMPS* to yield close to optimal parallel efficiency for at least 1000 atoms per core [11]. Our SDS and CTAB reference systems contain 900,000 and 550,000 atoms respectively, corresponding to 2800 and 1700 atoms per core when running on 320 cores. Both cases result in characteristic costs of 0.0075 core hours per atom per nanosecond on *NEMO*, which underlies Table 1.

Table 1. All runs on 320cores, MPI parallelized,*NEMO*. Multiplicity indi-cates runs due to unlistedparametric dimensions.

Туре	Sur-	Velocity	CPU-	run time	multi-	total
	factant	(m / s)	time (h)	(days)	plicity	cost (h)
normal	SDS	1.0	15,000	2	10	150,000
approach		0.1	150,000	20	3	450,000
	CTAB	1.0	43,000	6	11	473,000
		0.1	430,000	56	3	1,290,000
lateral	SDS	10.0	20,000	3	15	300,000
shear		1.0	65,000	9	15	975,000
				total		3,638,000

3. Results

Characteristic force response features are unidentifiable in the fast approach case of 10 m/s, as apparent in Figure 3. While two slower approaches still exhibit nonnegligible quantitative discrepancies, they produce comparable qualitative features, e.g. the transition from bi- to monomolecular layer



by squeeze-out and simultaneous onset of Au-Au attraction with surface-surface distance decreasing from 7 Å to 5 Å. Events at closer distances correspond to the probe's deformation around a few remaining surfactant molecules within the contact area. The table inlet shows diminishing Stokes drag within bulk water for slow velocities. Compromising between costs and accuracy, we pick 1 m/s as reference velocity.

Figure 4 (a) emphasizes distinguished force response features of SDS monolayers at increasing concentrations. Figure 4 (b) discriminates experimentally expected hemicylinders from the hypothetic monolayer for SDS surface number concentrations of 3 nm⁻². At comparable CTAB coverage, Figure 4 (c) displays the squeezing of a hypothetic monolayer against the approach above a bare substrate spot between globular aggregates. Results of lateral sliding tests on 3 nm⁻² SDS films are shown in Figure 4 (d) and (e). Former inspects normal and friction force at increasing distances above the monolayer, while latter contrasts sliding along and across experimentally expected hemicylinders against the monolayer case at a fixed distance of 5 Å.



Sixth bwHPC Symposium, 30th September 2019, Karlsruhe



References

- [1] J. Zhang and Y. Meng, Tribol. Lett. 56, 543 (2014).
- [2] M. Chen, I. Burgess, and J. Lipkowski, Surf. Sci. 603, 1878 (2009).
- [3] M. Jaschke, H.-J. Butt, H. E. Gaub, and S. Manne, Langmuir 13, 1381 (1997).
- [4] A. Jain, S. P. Ong, W. Chen, B. Medasani, X. Qu, M. Kocher, M. Brafman, G. Petretto, G. M. Rignanese, G. Hautier, D. Gunter, and K. A. Persson, Concurr. Comp.-Pract. E. 27, 5037 (2015).
- [5] H. Kawasaki, K. Nishimura, and R. Arakawa, J. Phys. Chem. C 111, 2683 (2007).
- [6] R. B. Best, X. Zhu, J. Shim, P. E. M. Lopes, J. Mittal, M. Feig, and A. D. MacKerell, J. Chem. Theory Comput. **8**, 3257 (2012).
- [7] H. Heinz, T. J. Lin, R. Kishore Mishra, and F. S. Emami, Langmuir 29, 1754 (2013).
- [8] D. Horinek, S. I. Mamatkulov, and R. R. Netz, J. Chem. Phys. **1301**, 124507 (2009).
- [9] S. Storm, S. Jakobtorweihen, I. Smirnova, and A. Z. Panagiotopoulos, Langmuir **29**, 11582 (2013).
- [10] G. Grochola, S. P. Russo, and I. K. Snook, J. Chem. Phys. 123, 204719 (2005).
- [11] S. Mashayak, M. Brown, C. Ponder, C. Trott, P. Crozier, F. Reid, and C. Vaughan, accessed 17 September 2019, https://lammps.sandia.gov/bench.html.


OpenCarme: - A Softwarestack for Interactive Machine Learning on HPC Systems

Dominik Straßel*, Franz-Josef Pfreundt* and Janis Keuper**

*Competence Center HPC, Fraunhofer ITWM, Kaiserslautern **Institute for Machine Learning and Analytics, Offenburg University

Abstract

We present our approach towards a customized HPC software-stack which combines high performance and usability for machine learning projects on HPC systems. Build on an opensource tool chain, we start at the system level and build an optimized stack which addresses and optimizes all aspects of ML projects: from storage, data I/O, support of various ML frameworks up to multi-GPU and distributed GPU training. Featuring a container based multi-user support which allows scheduling interactive container jobs alongside traditional batch processes. Finally, a convenient web-interface allows a user friendly integration of typical ML workflows on the HPC stack.

Project website: www.open-carme.org

Acknowledgements

The development of Carme has been supported by the BMBF project "AppliedML" and the state of Reinland-Pfalz.

SDS@hd - Scientific Data Storage

Martin Baumann, Alexander Balz, Oliver Mattes, Sabine Richling, Sven Siebler University Computing Centre (URZ), Heidelberg University, Germany

September 19, 2019

Abstract

SDS@hd is a central storage service for hot large-scale scientific data that can be used by researchers from all universities in Baden-Württemberg. It offers fast and secure file system storage capabilities to individuals or groups, e.g. in the context of cooperative projects. Fast access is possible from data generating facilities like microscopes as well as from data analysis systems like HPC systems. Data protection requirements can be fulfilled by data encryption and secure data transfer protocols. The service is operated by the Heidelberg University Computing Centre.

1 Introduction

 $SDS@hd^1$ is a service that provides large-scale storage for scientific data and is meant to be used for "hot data" – data that is frequently accessed and worked with. The service can be used by researchers at most public higher education institutions in Baden-Württemberg. User authentication and authorization are implemented in terms of the federated identity management in Baden-Württemberg bwIDM² allowing researchers to use the existing ID of their home institution transparently for this service. The SDS@hd service

¹https://sds-hd.urz.uni-heidelberg.de

²http://bwidm.de



Figure 1: SDS@hd can be accessed from HPC clusters in Baden-Württemberg.

website provides further information about the technical and institutional requirements and the registration process.

2 Features of SDS@hd

Tailored to the safety of research data Research data is a precious resource and should be stored using a trustworthy service to keep it safe from prying eyes. Data handled via SDS@hd is stored in an appropriate environment at the Heidelberg University Computing Centre (URZ). It is protected by state-of-the-art technologies, encryption as well as restrictive access and data policies.

Ideal for collaborations SDS@hd is useful for researchers from different departments or institutions who want to work together. They can join a collaborative storage project and store their research data at a single spot. Using a web interface, the storage project owner can manage user groups and user roles and can thus determine who is allowed to access which parts of the data storage.

Storing HPC and visualisation data The SDS@hd service is connected to a range of data-intensive computing resources also available through the Heidelberg University Computing Centre. These resources – such as the bwForCluster MLS&WISO or the bwVisu visualisation service – can automatically obtain or save data to and from the SDS@hd storage infrastructure. General access to SDS@hd, e.g. from the bwUniCluster and other bwFor-Clusters or from your notebook, is possible via the protocols SMB (2.x/3.x), NFSv4 (Kerberos) or SFTP (see Fig. 1). Users no longer need to find their own large-scale storage solutions.

3 Hardware

The data is stored on the Large Scale Data Facility (LSDF2), a state-ofthe-art storage system located at the URZ's main server rooms in Heidelberg. The system is associated to a project between URZ and the Steinbuch Centre for Computing (SCC) at Karlsruhe Institute of Technology (KIT). The project aims to provide modern and valuable storage services for researchers.

The system is gradually extended and will provide up to 25 petabytes of storage space in 2020. A range of the newest HDDs interspersed with several solid-state drives (SSDs) guarantees high access speeds. Even large numbers of smaller files – usually a time-consuming challenge for storage systems – can be swiftly processed.

LSDF2 is made for high availability: An efficient RAID concept prevents any data loss, and a backup power system weathers any potential power outages. There is no single point of failure in the system. Additionally a backup concept allows to access the data, even in a disaster situation without any time consuming data restore needed.

Acknowledgements

The LSDF2 infrastructure was built as part of the bwDATA initiative that fosters collaboration in the field of data-intensive computing between higher education institutions in Baden-Württemberg. The system is funded by the Ministry of Science, Research and the Arts Baden-Württemberg (MWK) and the German Research Foundation (DFG) through grant INST 35/1314-1 FUGG. Title: Preservation of Scientific Experiments in High Performance Computing – how could it be done?

Volodymyr Kushnarenko

Abstract:

Amount of the research experiments using High Performance Computing (HPC) is continuously growing. For the scientists and for the HPC providers it means, that not only a performance, but also a management of research data play an important role for any computations. One of the goals here is a preservation of research experiments for the long term. It can move the research process to a different level allowing storage, sharing, publishing and further reuse of the experimental results.

In this poster we want to discuss an idea of a preservation of HPC computations for the long term with a possibility of further reproduction. Starting with the HPC-related challenges we provide a prototypical realization of such a preservation system. The system uses a containerization technology (e.g. Docker, Singularity) as a preservation mechanism and has a client-server architecture. Also the prototype allows to replace some commercial software with the open source analogues to solve a legal issue which can appear by the reproduction. The poster presents the first experimental results and further investigation steps.