The NOMAD Center of Excellence (CoE): Tackling Exascale Computing and Extreme-Scale Data

The NOMAD Team*

NOMAD
NOVEL MATERIALS DISCOVERY

Part 1

Mission of the NOMAD CoE

With 10^{18} floating point operations per second, exascale computers will be able to simulate and analyze complex processes more realistically than what is possible today. This will allow to study systems of higher complexity (in space and time) and to achieve higher accuracy and precision. In close contact with academia and industry, the NOMAD CoE will develop exascale algorithms to build reliable models for real-world, industrially relevant, complex materials. We will provide ***exascale software*** and ***artificial intelligence*** (AI) methods that provide more efficient workflows and advanced modeling. Finally, we will test and demonstrate our developments on ***urgent problems*** in the areas of ***energy and environment***. This, together with our extreme-scale ***data services*** and ***training*** of the next generation of HPC students reaches out to society.

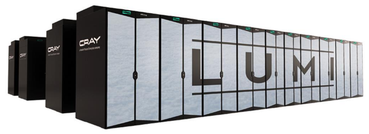

Fig 1.: Example for a supercomputer: LUMI at the CSC Data Center in Finland. HPE Cray EX cabinets, (copyright: Hewlett Packard Enterprise).

Structure of the NOMAD CoE

In the NOMAD CoE we prepare for the next HPC generation by

- advancing ***ab initio*** computational materials science for entire code families, e.g. pseudopotential and all-electron methods (**work packages (WPs) 1-3**),
- developing exascale workflows for high-throughput studies (**WP 4 and 5**),
- developing AI tools to near real-time performance (**WP 6**),
- developing a data infrastructure (**WP 7**) to support WP 1-6,
- co-designing soft- and hardware developments to achieve optimum performance (**WP 8**),
- testing and demonstrating our developments by two use cases, addressing urgent challenges for energy and environment needs (e.g. photocatalytic water splitting and waste-heat recovery, **WP 9**).
- **WP 10** ensures that our the developments benefit industry and society.
- **WP11** focuses on the dissemination of our concepts, results and software to academia, industry and society.
- Our training program in **WP12**, will expand the knowledge in exascale ***ab initio*** computational materials science in Europe.
- **WP13** manages and administrates all CoE activities, also ensuring the efficient collaboration of the team.

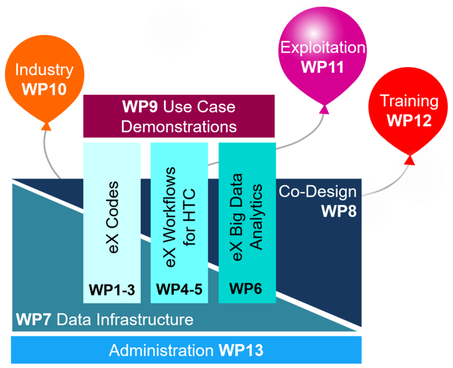

Fig. 2: The NOMAD CoE consists of 3 pillars and 13 work packages.

Cutting-Edge Electronic-Structure Methods

The NOMAD CoE will push the boundaries of electronic-structure theory to previously inaccessible system sizes, material complexity, accuracy, and precision. Higher-level methods, ranging from advanced exchange-correlation (XC) functionals to coupled-cluster (CC) theory, are essential to model new materials for modern technologies. Figure 3 displays the scaling of the computational costs with system size.

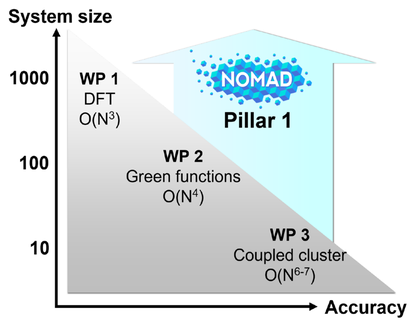

Fig. 3: Scope of applicability and accuracy for density-functional theory (DFT), Green-function, and coupled-cluster methods. The system size is measured in number of atoms per simulation cell. The formal scaling of (semi)local DFT functionals is cubic, *GW* quartic, CC even higher. Advanced functionals will be in the focus of WP1, Green-function based methods of WP2, and coupled-cluster method of WP3. The gray area indicates today's state-of-the-art. NOMAD will increase the accessible system-size (*y*-axis) by several orders of magnitude.

The NOMAD CoE Team

To achieve high efficiency in addressing the needs of academia, industry, and society, the NOMAD CoE is organized in two "shells". There is a ***core team**** (project beneficiaries, leading code developers) that is responsible for the research and development of the project. And there is the **second shell** including renowned experts from academia and industry who will provide advice and work with the NOMAD CoE team on the various developments. They also contribute to the setting of goals and priorities.

Fig. 4: The NOMAD CoE consortium.

Get in Touch with Us

- **General questions** concerning the CoE
  contact@nomad-coe.eu
- **NOMAD *ab initio* software**
  FHI-aims: sebastian.kokott@fhi.mpg.de
  exciting:
  claudia.draxl@physik.hu-berlin.de
  Abinit: xavier.gonze@uclouvain.be
  GPAW: thygesen@fysik.dtu.dk
- **Workflows**
  ASE/ASR: thygesen@fysik.dtu.dk,
  FireWorks:
  geoffroy.hautier@uclouvain.be
- **NOMAD data infrastructure**
  markus.scheidgen@physik.hu-berlin.de
- **Exascale HPC facilities**
  msmit1@bsc.es
- **Use-case demonstrators**
  christian.carbogno@fhi.mpg.de
  and thygesen@fysik.dtu.dk
- **Data analytics/ artificial intelligence**
  luca.ghiringhelli@physik.hu-berlin.de

Use-Case Demonstrators for Urgently Needed Technologies

To achieve the critical climate targets by, e.g., reducing CO_2 emission, we need improved or new technologies that enable a CO_2-neutral energy production. The NOMAD CoE will develop methods to identify new materials that i) serve as catalysts for the sustainable production of hydrogen as a renewable and clean-energy carrier. And ii) enable the recovery of energy that is presently wasted in the form of heat.

(1) PHOTOCATALYTIC WATER SPLITTING

The sustainable production of hydrogen is of paramount importance for the future synthesis of chemicals and fuels. Photocatalytic water splitting (PCWS) could become a key technology for hydrogen production, but so far it is not feasible due to limitations of the catalysts that are used. The discovery of the urgently needed catalysts faces a number of challenges arising from the complexity of the materials space and the difficulties in theoretically describing the interactions of water with the catalyst surface and the band gaps of photo- absorbers. DFT does not provide the required accuracy, and beyond DFT methods are currently computationally too intensive. Based on the various NOMAD developments, and in close collaboration with industrial partners, we will advance the field of PCWS by performing high-throughput machine-learning-based screening of novel PCWS materials and beyond DFT methods.

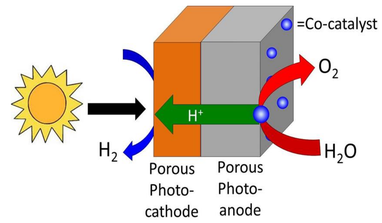

Fig. 5: Identification of a better catalyst for water splitting has enormous potential in carbon-dioxide management, in hydrogen-based energy, and more.

(2) WASTE-HEAT RECOVERY

Over 40% of our current energy consumption is wasted in the form of heat. Thermo-electric devices can convert heat into electricity, be it from industrial furnaces, chemical plants, exhaust pipes, or computers. They therefore have the potential to reduce the global energy consumption. More efficient and affordable thermoelectric materials are needed to make such devices economically viable. The thermoelectric efficiency depends on subtle details of charge and heat transport.

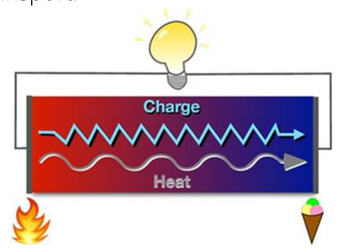

Fig. 6: In thermoelectric materials, a temperature gradient induces a charge flux that can be used to obtain a useful electrical current from otherwise wasted heat.

NOMAD's exascale developments will enable us to calculate reliable electronic structures and electron mobilities as well as thermal conductivities by molecular dynamics with an accuracy better than that of current DFT. Using our exascale workflows we will provide a comprehensive database of thermoelectric efficiencies with unsurpassed accuracy. Our powerful AI techniques then enable a systematic exploration and identification of high-performance thermoelectrics.

The NOMAD Center of Excellence (CoE): Tackling Exascale Computing and Extreme-Scale Data

The NOMAD Team*

NOMAD
NOVEL MATERIALS DISCOVERY

Part 2

First Achievements

1. Bringing the Eigensolver ELPA from the Peta- to Exascale

The computational cost of large-scale, semi-local DFT calculations is dominated by the solution of (generalized) complex eigenproblems. This scales as N^3 with the system size N, which severely limits the complexity of the scientific questions that can be investigated. To date, the open-source, highly-scalable eigensolver ELPA, originally developed in a collaboration involving the MPCDF and NOMAD Laboratory, is superior to all alternatives. As NERSC and Cray recently have demonstrated, it outperforms SCALAPACK on all architectures tested (see Fig. 7)[1]): ELPA is significantly faster and exhibits improved scalability to much larger number of cores. These advantages translate into immense computational savings in actual DFT calculations. The 2021 release of ELPA already works on current (pre)-exascale technologies. This includes support for GPUs [2,3] from Nvidia, Intel, and AMD as well as the ARM-based architecture used in Fugaku, one of the fastest supercomputers available to date.

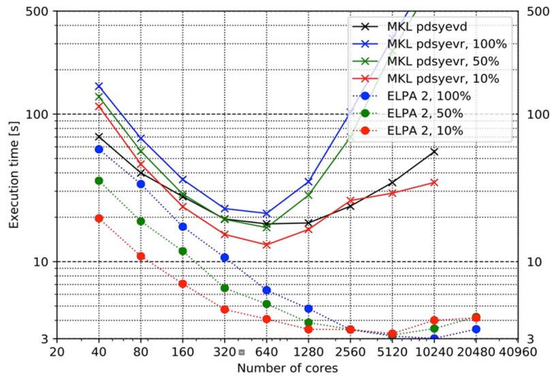

Fig. 7: Performance and scaling comparison of ELPA and the Intel Math Kernel Library (MKL) for the diagonalization of a double-precision, real matrix of size n = 20,000. The percentage of computed eigen-vectors (pdsyevd in %) is noted as well.

In addition, ELPA now supports mixed-precision calculations [4] as well as autotuning.[5] For the latter, it identifies the degree of GPU deployment, algorithmic parameters, CPU-specific vectorization kernels, etc. in an iterative fashion. With that, the ELPA 2021 release is the most performant, scalable, and user-friendly eigensolver for materials science available to date.

2. Workflow Recipes for Extensive *ab initio* Calculations

The recently developed Atomic Simulation Recipes (ASR) [6] allow for efficiently managing and performing advanced computational tasks that, e.g. involve several concurrent or sequential first-principles calculations.

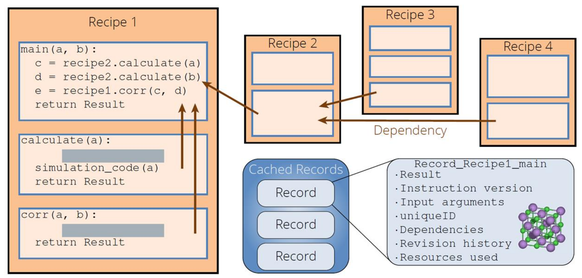

Fig. 8: ASR defines complex workflows from a set of smaller smaller tasks, i.e. recipes. Each result is cached as a record with all relevant metadata to preserve the provenance of each calculation.

Examples are the assessment of Raman spectra, electronic effective masses, or energy barriers. The central concept of ASR is a high-level python script used to preform a task while preserving the data provenance.

3. A Massive Parallel Framework for Symbolic Regression and Compressed Sensing

The SISSO++ package [7, 8] provides a massively parallel implementation of the sure-independence screening and sparsifying operator (SISSO) method. This AI technique combines symbolic regression and compressed sensing and provides interpretable descriptors and identifies "materials genes". So far, SISSO has proven successful in modeling, for example, the stability of crystal phases, catalytic reactivity and selectivity, glass transition temperatures, Stability of perovskites, and the topological character of materials. The new implementation also provides a python interface to facilitate the post processing and analysis of the discovered expressions. Detailed documentation and tutorials provide a good starting point to learn how to apply SISSO to your own applications.

4. Coupled-Cluster Theory for Periodic Solids

Coupled-cluster (CC) theory is considered the gold standard of quantum chemistry. It is the benchmark for approximate theories. The NOMAD team has developed an interface (CC-aims, [10]) to the CC-for-solids (cc4s) code [11]. With this interface, CC and MP2 calculations can now be performed with all-electron localized basis set codes, e.g. FHI-aims.

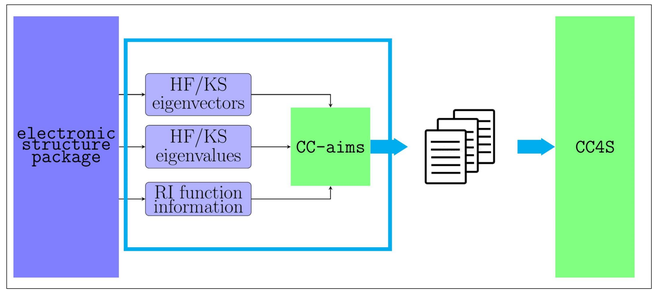

Fig. 9: CC-aims is an interface the to coupled-cluster-for-solids (CC4S) code. It was developed for all-electron codes.

A recent advancement of cc4s exploits that spin and momentum conservation imply a significant degree of sparsity in central quantities of correlated methods. The resulting tensor contractions are efficiently executed by the Cyclops Tensor Framework (CTF), which is designed for modern super computer architectures. In combination this makes a fundamental step towards making CC methods accessible to exascale architectures.

5. High-throughput Calculations of Vibrational Properties

The newly developed python package FHI-vibes [11] facilitates high-throughput, first-principles calculations of vibrational properties of solids. It builds and interlinks established frameworks such as, e.g., the *Atomic Simulation Environment* ASE [12], the workflow management suite *Fireworks* [13], and the harmonic phonon code *phonopy* [14]. By this means, it seamlessly bridges between perturbative methods based on the harmonic approximation and fully anharmonic, *ab initio* molecular dynamics simulations. For this reasons, *FHI-vibes* is the key piece of software that has enabled the high-throughput search for thermal insulators via *ab initio* Green-Kubo calculations, [15].

A FAIR Research-Data Infrastructure for Materials Science

The FAIRmat project – FAIR Data Infrastructure for Condensed-Matter Physics and the Chemical Physics of Solids – is an off-spring and advancement of the original NOMAD CoE. It expands the computational materials science focus by including experiments and materials synthesis. FAIRmat is being funded as a consortium of the German National Research-data infrastructure initiative (NFDI), starting October 1st, 2021. The initial runtime is 5 years (possible prolongation by another 5 years) and the funding is 3.3 Mio Euro/year.

Fig. 10: General structure of the FAIRmat project

For more information please visit
https://fair-di.eu/fairmat/

References

[*] The PIs of the NOMAD CoE are: M. Scheffler, L. Ghiringhelli, C. Carbogno, (NOMAD Laboratory at the Fritz Haber Institute, Berlin), E. Laure, H. Lederer, A. Marek, M. Rampp (Max Planck Computing and Data Facility, Garching), C. Draxl, M. Scheidgen (Humboldt-Universität zu Berlin), P. Rinke (Aalto University), J.-M. Cela, M. Smith (Barcelona Supercomputing Center), C. Menache, M. Torrent (Commissariat à l'Énergie Atomique et aux Énergies Alternatives, Bruyeres-le-Chatel), K. Koski, J. Heikonen, J. Enkovaara (CSC–IT Center for Science, Espoo), K. S. Thygesen (Technical University of Denmark, Lyngby), A. Grüneis (Technische Universität Wien), X. Gonze, G.-M. Rignanese, G. Hautier (Université Catholique de Louvain), G. Csányi (University of Cambridge), A. Gulans (University of Latvia, Riga), J. Kermode (University of Warwick) and many associated researchers from academia and industry, see: https://nomad-coe.eu/about/principal-investigators

[1] B. Cook, et al., Concurrency and Computation: Practice and Experience **180**, e4997 (2018).

[2] Hans-Joachim Bungartz, et al., Parallel Computing: Technology Trends **36**, 647 (2019).

[3] V. W.-Z. Yu, et al., Comp. Phys. Comm. **262**, 107808 (2021).

[4] A. Alvermann, et al., Japan Journal of Industrial and Applied Mathematics **37**, 1 (2019).

[5] P. Kůs, et al., Parallel Computing **85**, 167 (2019).

[6] M. Gjerding et al., Comp. Mat. Sci. **199**, 110731 (2021).

[7] https://sissopp_developers.gitlab.io/sissopp

[8] T. A. R. Purcell, M Scheffler. C. Carbogno, LM Ghiringhelli, J. of Open Source Software, to be published.

[11] E. Mörmann et al.,https://joss.theoj.org/papers/72b4c8d40d9a0401039ade39a7801759

[10] F. Hummel, T. Tsatsoulis, A. Gallo, A. Irmler, T. Schäfer, and A. Grüneis: to be published

[11] F. Knoop, T. A. R. Purcell, M. Scheffler, C. Carbogno, J. of Open Source Software **5**, 2671 (2020).

[12] A.H. Larsen, et al., Journal of Physics: Condensed Matter **29**, 273002 (2017).

[13] A. Jain, et al., Concurrency Comput Pract Exp **27**, 5037 (2015).

[14] A. Togo and I. Tanaka, Scripta Materialia **108**, 1 (2015).

[15] C. Carbogno, R. Ramprasad, and M. Scheffler, Physical Review Letters **118**, 175901 (2017).

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 951786.