Tier-1 Application form – English version

APPLICATIONS ARE PREFERABLY DRAWN UP IN ENGLISH. AN ENGLISH TRANSLATION HAS TO BE ENCLOSED WITH APPLICATIONS SUBMITTED IN DUTCH.

The application form is available in English on the website
https://www.vscentrum.be/en/access-and-infrastructure/project-access-tier1

Title of the application:

Investigating the low-frequency vibrational fingerprint of flexible metal-organic frameworks

Name and first name of the applicant:

Hoffman Alexander

Institution:

Ghent University

Research group / department:

Center for Molecular Modeling

Title / position:

ir. / PhD fellow

Email address:

Alexander.Hoffman@UGent.be

Total computing time that is needed, in node days:

1062

Total scratch disk storage that is applied for (in GiB) and number of files:

260.8 GiB + 21436 files on Tier-1 SCRATCH / 40.36 GiB + 9320 files on Tier-2 DATA/HOME

The total number of pages in this application should not exceed 18, excluding possible appendices (confirmation letter of financing institution,

software license, etc.) which may be taken into account by the Tier-1 Allocation Board.

1. Title of the research project (with IWETO or FRIS link if available) within the framework of which computing time is applied for:

Unraveling phase transformations and reactivity in nanoporous materials using computational vibrational spectroscopy

2. Describe your research project in short. Explicitly mention the scientific questions that you are planning to address and the overall scientific goals of the project. (max. 1 A4 in Arial 12):

Metal-organic frameworks (MOFs) are a class of porous crystalline materials consisting of metal-oxide clusters connected through organic linkers. [1,2] The structural and chemical versatility of MOFs allows for their use in many potential applications. [3] An important subclass of MOFs consists of flexible MOFs that can undergo phase transformations. [4] These so-called soft porous crystals display a bistable or multistable behavior with long-range structural order [5] and can undergo large volume changes upon exposure to external stimuli. [4] Within the past decade, an enormous endeavor has been undertaken to understand this experimentally observed phenomenon, [6] revealing various indications that terahertz vibrations (<100 cm^{-1}) may trigger the observed flexibility. [7] However, thus far, our understanding of this part of the frequency region is not yet well developed, due to the experimental difficulties to measure within the terahertz region on the one hand and the difficulty in overcoming computational hurdles to obtain accurate vibrational spectra for such low wavenumbers on the other hand.

Recently, we published an in-depth spectroscopic investigation on the prototype flexible MOF MIL-53(Al). [8] The vibrational spectra of the closed-pore (CP) and large-pore (LP) phases were measured experimentally and predicted theoretically by means of static and dynamic DFT calculations. We were able to identify several collective vibrations that are capable of triggering the phase transformation. This study can be regarded as a proof of concept for the identification of lattice modes that give rise to flexibility in MOFs via vibrational spectroscopy in the far-IR range (see Figure 1).

In this project, we will determine the Raman spectrum of MIL-53(Al) via static calculations. Furthermore, we will use the same methodology as in our previous study on MIL-53(Al) [8] to investigate the low-frequency spectrum of another flexible MOF, namely DUT-8(Ni). The underlying objective is to discover specific collective vibrations that induce phase transformations in this type of materials, for which the obtained computational results will be compared with experimental data in the far-IR range [7].

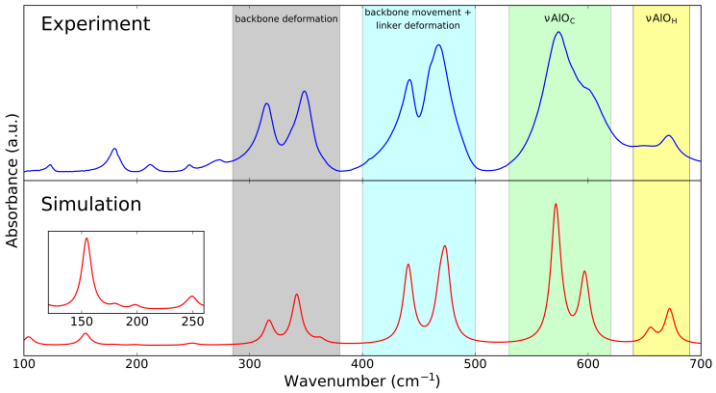

Figure 1. Experimental and theoretical far-IR spectrum of MIL-53(Al). [8]

3. Provide an engaging abstract (10 lines) for scientific communication on the website in layman's terms. Should this application be bound by a confidentiality agreement (see also item 12 of this application form), provide more details about the specific nature of the confidentiality and indicate why an abstract may not be published.

Metal-organic frameworks (MOFs) are a class of porous crystalline materials consisting of metal-oxide clusters connected through organic linkers. The structural and chemical versatility of MOFs allows for their use in many potential applications. An important subclass of MOFs are the soft porous crystals displaying a bistable or multistable behavior with long-range structural order, which can undergo large volume changes upon exposure to external stimuli. Various indications have been reported that terahertz vibrations may trigger the observed flexibility. Therefore, this project aims at an in-depth investigation of phase transformations in soft porous crystals via vibrational spectroscopy.

4. Financing institution or channel, financing the research project in full or in part (FWO, BOF, IWT, EU, etc.): Please attach the confirmation letter as enclosure. In case the project has not gone through a scientific approval process attach a letter of approval of your own institute.

FWO (confirmation letter at the end of this document)

5. Name and email address of the promoter(s) of the research project:

Prof. Dr. ir. Veronique Van Speybroeck

Veronique.VanSpeybroeck@UGent.be

6. Persons mandated by the Applicant to compute on the Tier-1 within the framework of the present project: Please provide for every person:
   - name and first name
   - institution

- research group / department
- title / position
- experience of using HPC resources (Tier-0/Tier-1/Tier-2 infrastructure in Belgium and abroad) in the past (specify names and when)

Hoffman Alexander

- Ghent University
- EA17 / Center for Molecular Modeling
- PhD student
- Experience with local CMM clusters, HPC UGent Tier-2 clusters (vsc41949)

Lamaire Aran

- Ghent University
- EA17 / Center for Molecular Modeling
- FWO doctoral fellow
- Experience with local CMM clusters, HPC UGent Tier-2 clusters and the Tier-1 machine breniac (vsc41948)

Wieme Jelle

- Ghent University
- EA17 / Center for Molecular Modeling
- FWO doctoral fellow
- Experience with local CMM clusters, HPC UGent Tier-2 clusters and the Tier-1 machines muk and breniac and Tier-0 PRACE (Marconi KNL, Marconi Broadwell, Curie TN) (vsc40944)

Van Speybroeck Veronique

- Ghent University
- EA17 / Center for Molecular Modeling
- Professor
- Experience with local CMM clusters, HPC UGent Tier-2 clusters and the Tier-1 machines muk and breniac (vsc40021)

7. Explain why this project needs to run on a Tier-1 system, why the machine you have requested is suitable for the project and how the use of the system will enable the science proposed (max. ½ A4 in Arial 12).

There are three main reasons why this project needs to be run on a Tier-1 system:

1) First of all, the calculation of vibrational spectra of molecular systems requires the determination of a Hessian, which is the matrix of

second-order derivatives of the energy with respect to the atomic positions. These calculations have to be performed in a single run, as they cannot be restarted, requiring a large number of nodes for simulations of complex structures, such as DUT-8(Ni).

2) The second reason concerns the calculation of Raman intensities. These intensities are third-order derivatives of the energy with respect to the atomic positions and the electric field. Using a finite difference approach, this requires a large amount of single-point calculations, which is not feasible on Tier-2 machines.
3) A third reason to run this project on Tier-1 is the fast access to computing resources. As our study on the static IR spectra of MIL-53(Al) has been recently published, we want to perform the complementary calculations of the Raman spectra as soon as possible.

8. Justify the number of node days requested. This should include information such as: number and nature of computing tasks, software used, and the sequence in which they will be performed.
   Indicate for each typical computing task the required resources:

   - wall clock time (note that 3 days is the maximal wall clock time for any job;)
   - memory (maximum 128 GiB/node; 256 GiB/node is available upon motivated request)
   - number of nodes (could be fractional)
   - number of CPU cores
   - disk space (estimated volume in GiB and the total number of files); make a clear distinction between usage of Tier-2 DATA/HOME partitions and the Tier-1 SCRATCH partition
   - number of tasks, and an indication of how many such tasks would be submitted concurrently.

   This information should take the form of a table (an example is provided as Table 1 on the next page). Provide additional descriptions of the computing tasks and comments as needed and clearly relate the described tasks to the tasks in the table. Resource estimates should be preferably based on the results of actual calculations on Tier-1 (via, e.g., a Starting Grant) for system/problem sizes that are on par with those of the intended computing tasks (e.g., same mesh sizes, actual molecular system, …). If not, provide the name, architecture, #cores, memory, etc. of the machine that was used to obtain these results and explain how you have calculated/rescaled the wall clock times, number of cores, etc.

Please do take into account that after half of the allocation time has passed, the applicant loses 30% of the initially granted node days, if that 30% has not been used. (cf. regulations, § 9)

(typically, up to 2 A4 Arial 12).

In the first part of this project, we will calculate the Raman spectra of the closed-pore (CP) and large-pore (LP) phase of MIL-53(Al) (type 1 simulations). This requires the calculation of Raman intensities, which follows from mixed third-order derivatives of the energy with respect to the electric field and the atomic positions. Using the VASP software, the polarizability tensor (second-order derivative of the energy with respect to the electric field) can be calculated efficiently (simulations of +/- 12 h). The remaining derivative of the energy with respect to the atomic positions has to be calculated using a finite difference approach, which requires two separate simulations for every vibrational mode in MIL-53(Al). The total number of normal modes is 3N, with N=76 for each phase of MIL-53(Al). As such, the Raman spectrum of a single phase of MIL-53(Al) requires 2x228=456 evaluations of the polarizability tensor. For both phases of MIL-53(Al), this gives a total of 912 simulations.

The limited computational time per calculation is the result of previous calculations performed on Tier-1 and Tier-2. These allow us to start from optimized structures, so that only cheap DFPT simulations are needed for the calculation of the polarizability tensor.

The second part of the project entails the calculation of the Hessian of both the CP and LP phase of DUT-8(Ni). For the Hessian to be positive semidefinite, which is necessary to have only nonnegative frequencies, the CP and LP phase structures have to be brought to equilibrium, which requires a geometry optimization (type 2 simulation). The first step in the optimization procedure consists of a volume relaxation, where both the volume, cell shape, and ionic positions are allowed to relax. Once an estimated equilibrium volume is obtained, 7 additional optimizations will be performed at fixed volumes near the optimal volume (-6%, -4%, -2%, +0%, +2%, +4%, and +6%), in which the cell shape and ionic positons are still allowed to relax. A Rose-Vinet profile is then fitted to the resulting energy values to obtain the actual optimal volume. To conclude, a final optimization step is performed at this volume whereby the cell shape and ionic positions are allowed to relax. Both phases thus require 9 optimization steps, which gives a total of 18 optimization steps.

Thus far, we have not performed any calculations on the DUT-8(Ni) structure yet, but since its topology is similar to the one of DMOF-1(Cu), replacing BDC linkers by NDC linkers, we refer to benchmarking calculations obtained for DMOF-1(Cu). Scaling tests on DMOF-1(Cu) show a nearly perfect scaling when using 3 node calculations. One optimization

step takes about 9 days and needs to be restarted twice. This means that we will perform 18 x 3 = 54 calculations, an estimate which is based on our experience on similar systems and calculations run on the Tier-1 infrastructure within the framework of a previous proposal. Hence, we will use 3 nodes (x 28 cores), as this results in a nearly perfect scaling, while losing 15-30% efficiency when using more nodes (9).

Once the equilibrium structure of the CP and LP phases has been determined, we can start calculating the Hessian of both structures (type 3 simulations), simulations which have to be performed in a single run. Based on our experience with similar systems, we expect 20 nodes to be sufficient to complete the task within 3 days.

This project can hence be subdivided into four parts: two calculations of Raman intensities and two Hessian calculations. The aim is to spread these simulations as well as possible: one month for the calculation of a single Raman spectrum and two months for a geometry optimization and the subsequent Hessian calculation. In this way, a constant monthly usage is targeted.

In total, 932 independent VASP calculations need to be performed, which will be run in different directories, where each VASP simulation creates 25 files. After the VASP calculation is finished, some of the files will immediately be removed (CHGCAR + CHG = 2 GiB), which will be stated in the jobscript. After a successful calculation, we will therefore have 23 files which take about 0.175 GiB for the simulations on MIL-53(Al) and respectively 4.9 GiB and 6.5 GiB for the geometry optimization steps and Hessian calculation of DUT-8(Ni), where the latter results have been checked during the scaling tests on DMOF-1(Cu). For the geometry optimization, the calculation is restarted in the same directory, overwriting the existing files.

After the simulations, the data need to be copied back to the Tier-2 DATA/HOME (or to our DATA_VO). We will only copy back 0.03 GiB, 0.5 GiB, and 2 GiB per type 1, 2, and 3 simulation respectively, retaining only 10 files (OUTCAR, CONTCAR, POSCAR, KPOINTS, INCAR, tempout, vasprun.xml, POTCAR, and error/output files). This will take in total 40.36 GiB of permanent storage until the project has been published. The other files are not useful to keep, and the most memory consuming file (WAVECAR = 5 GiB) can be safely removed on the Tier-1 SCRATCH. This is also summarized in Table 1.

Table 1

	Node day calculation							Storage volume estimate	
Computational task	# of such tasks	Wall clock time (days) per task	# Tier-1 nodes per task	# total node days task	# CPU cores per task	Memory usage (GiB) / node per task	OpenMP / MPI / OpenMP + MPI (hybrid)	Tier-2 DATA/HOME volume (GiB) + number of files	Tier-1 SCRATCH volume (GiB) + number of files
1. Raman intensities MIL-53(Al)	2 x 456 = 912	0.5	1	456	28	18 GiB / node	MPI	0.03 GiB x 912 = 27.36 GiB + 10 files	0.175 GiB x 912 = 159.6 GiB + 23 files
2. Geometry optimization DUT-8(Ni)	2 x 9 = 18	3 x 3 = 9 (two restarts required)	3	486	84	37 GiB / node	MPI	0.5 GiB x 18 = 9 GiB + 10 files	4.9 GiB x 18 = 88.2 GiB + 23 files
3. Hessian DUT-8(Ni)	2	3	20	120	560	37 GiB / node	MPI	2 GiB x 2 = 4 GiB + 10 files	6.5 GiB x 2 = 13 GiB + 23 files

9. Describe the software required to perform the computing task(s). Please clearly provide the following per item in this regard:

   - a reference to the software's web page
   - the software license system (open source, GPL, etc.)
   - if there is no free academic use of the software, state which license makes the installation and the use valid on the Tier-1 by the Applicant (+ add a copy of the signed license)
   - if need be, which license server will be used (name + IP address)
   - whether the software is already available on the Tier-1 and, if this is not the case, compilation and installation instructions (possibly with reference to existing Tier-2 installation)

VASP

http://www.vasp.at/
non-exclusive academic license (see attachment)
available on BrENIAC (version 5.4.1-intel-2016a)

Provide the results of efficiency tests that were conducted with this software, preferably on the current VSC Tier-1 (using, e.g., a Starting Grant) for system/problem sizes that are on par with those of the intended computing tasks (e.g., same mesh sizes, actual molecular system, …). If not run on the current VSC Tier-1, provide the name, architecture, #cores, memory, etc. of the machine that was used to obtain these results and how you think this compares to the current VSC Tier-1. If a different system/problem size is used provide some guidance how it relates to the problem size in the application.
Provide both a table and efficiency plot such as Table 2 and Plot 1 below (log scale X-axis) (typically up to 3 A4 in Arial 12).

Important: the baseline configuration (first line in the table) should be the *minimal* configuration on which your computation could be carried out, where the number of cores for this baseline configuration could be 1, 2, …, 28, 56, ... *Table and plot should be created w.r.t. the baseline configuration.*

Discuss the plot and table (explain any anomalies). Explain why you make a specific choice for a given number of nodes (in case of multi node jobs).

In case your tasks are single node jobs (HTC), describe how you will manage your jobs, e.g. will you make use of the Worker Framework[1] or something similar?

If the simulation contains a significant amount of I/O, then mention its impact on the calculations.

VASP was tested extensively during the pilot projects on BrENIAC for various systems. The results of these tests are added in attachment. Moreover, in previous projects of some of the applicants of the Center for Molecular Modelling, several comparable systems were evaluated. The conclusions of the corresponding scaling tests are repeated here.

For DUT-8(Ni), no earlier scaling tests are performed, but this system is topologically similar to DMOF-1(Cu) for which scaling tests are already available. The main difference between both structures is the linker (NDC for DUT-8(Ni) and BDC for DMOF-1(Cu)), but the metals are both magnetic and thus comparable.

The VASP parallelization settings include KPAR – the number of k-points that are to be treated in parallel – and NCORE, which determines how many cores work on one orbital. A single point calculation of a fixed number of 20 electronic steps was performed. The other VASP settings used in the test on DMOF-1(Cu) correspond with production settings that have the required accuracy and that are aimed to be used in this proposal. They include ENCUT = 600 eV, PREC = ACCURATE, LREAL = AUTO, EDIFF = 1.0E-8, EDIFFG = 1.0E-7, and a k-point mesh of 3x2x3. This k-point mesh leads to 9 irreducible k-points, which indicates that this is the maximum KPAR setting we can use.

An overview of all data at different parallelization settings is given in Extra Table 1 and 2 below and is summarized in Table 2 and Plot 1. The latter sketches the speedup as a function of the number of cores. From the plot, it is clear that for simulations with 3 nodes, we get nearly perfect scaling (95%), while losing 15%-30% efficiency when using more nodes. Moreover,

[1] https://www.vscentrum.be/cluster-doc/running-jobs/worker-framework

from the Extra Tables, we see that the KPAR=3 and NCORE=28 settings are most efficient.

Table 2: Summary of benchmark results for the magnetic DMOF-1(Cu)

# nodes	# cores	Wall clock time (s)	Speedup (w.r.t. respect to 1 node)	Efficiency
1	28	5603,260	1,00	1,000
3	84	2008,306	2,79	0,930
9	252	868,551	6,45	0,717
18	504	578,005	9,69	0,538

Plot 1

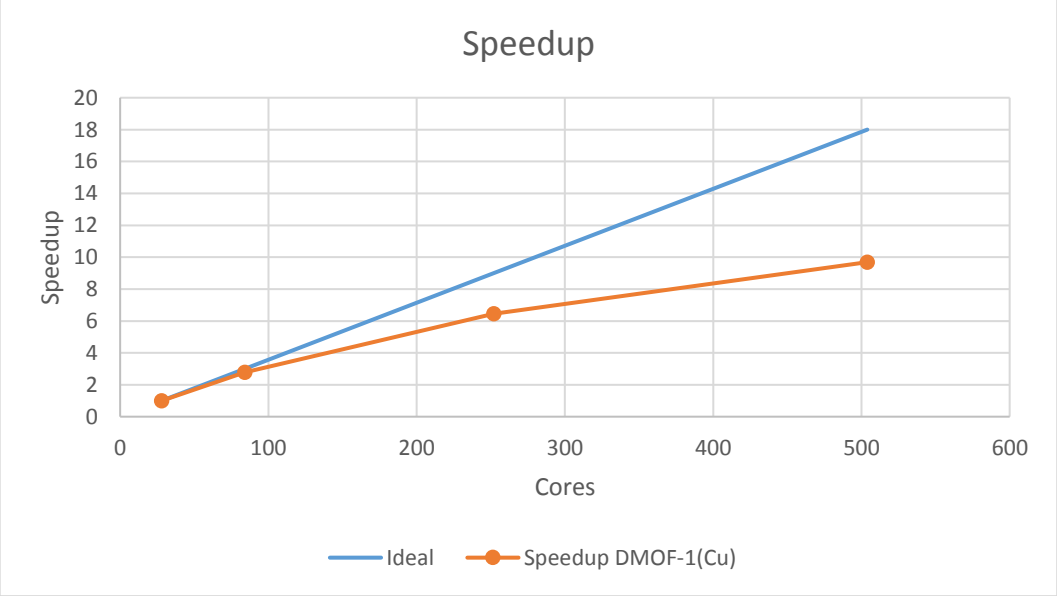

Extra Table 1: Overview of all scaling tests carried out for magnetic DMOF-1(Cu). The table elements are the time of the calculation in seconds. The two important parallellization parameters for VASP (NCORE and KPAR) are varied, together with the number of nodes. The numbers indicated in green are the fastest simulations at a specific number of nodes and are also summarized in Table 2.

	DMOF-1(Cu)		MAGNETIC		
KPAR=1		Number of nodes			
	NCORE	1	3	9	18
	1	7110,528	2493,731		
	7	6131,798	2375,737	1511,171	927,22
	14	5808,558	2142,882	1031,631	855,376
	28	5603,26	2008,306	1017,934	824,245
KPAR=3		Number of nodes			
	NCORE	1	3	9	18
	1		3857,041	1584,09	947,829
	7		3246,061	999,909	670,379
	14		2710,283	929,737	700,739
	28		2124,847	868,551	636,409
KPAR=9		Number of nodes			
	NCORE	1	3	9	18
	1			1447,924	907,861
	7			1878,574	622,478
	14			1219,593	776,179
	28			950,342	578,005

Extra Table 2: Overview of all scaling tests carried out for magnetic DMOF-1(Cu). The table elements are the memory usage per core in kB. The two important parallellization parameters for VASP (NCORE and KPAR) are varied, together with the number of nodes. The results show that the memory usage increases with decreasing NCORE and increasing KPAR.

	DMOF-1(Cu)				
KPAR=1		Number of nodes			
	NCORE	1	3	9	18
	1	3509500	1940384		
	7	3157928	1476336	950092	893396
	14	3110000	1421168	874784	795492
	28	3120732	1395020	856784	753900
KPAR=3		Number of nodes			
	NCORE	1	3	9	18
	1		3516812	1945100	1726132
	7		3167564	1477168	1056712
	14		3123516	1416680	1020064
	28		3127504	1397632	997740
KPAR=9		Number of nodes			
	NCORE	1	3	9	18
	1			3518384	2341952
	7			3168888	1873000
	14			3125984	1834960
	28			3121024	1826984

The calculations of the Raman intensities of MIL-53(Al) will be performed on 1 node. VASP scaling tests on this material can be found in Table 3, where the optimal combination of k-parallelization over electronic bands is determined. As was the case in previous benchmarks, a high k-point parallelization is favorable.

Table 3: Overview of the scaling tests carried out for MIL-53(Al). The two important parallellization parameters for VASP (NCORE and KPAR) are varied, together with the number of nodes.

		NCORE			
		1	7	14	28
KPAR	1	59,980	46,482	46,085	49,954
	4	56,843	45,510		

10. Describe how you will manage the resources requested in the period during which the task is to be performed. What usage pattern do you anticipate (similar usage on monthly basis, bursts, …)?

    In case the disk storage or the number of files that is applied for exceed the defaults of 2 TiB and 200.000 files (cf. regulations §4, (5)), describe how the transfer of files to/from Tier-1 will be managed, if data reduction and/or compression of files will be performed, and if you can provide information about IOPS.

This project can be subdivided in four parts: two calculations of Raman intensities and two Hessian calculations. The aim is to spread these simulations as well as possible: one month for the calculation of a single Raman spectrum and two months for the geometry optimization and the subsequent Hessian calculation of a single phase. In this way, a constant monthly usage is targeted, which also allows to gradually transfer back data from the separate Raman intensity calculations from Tier-1 to Tier-2, restricting the size of the data transfer between Tier-1 and Tier-2. See section 8 for more details on the planned calculations.

11. List the granted computing time allocations to the promoter(s) of this research project, on the Flemish Tier-1 systems, as well as other Tier-1 and Tier-0 systems. Also, describe the scientific output obtained within the framework of computing time that was granted during the past two years on the Flemish Tier-1 systems or on other Tier-1 or Tier-0 supercomputers. DOI links are sufficient.

Aran Lamaire:

Tier-1:

Understanding the high-pressure behavior of a flexible nanoporous material (Node days: 2430)

Benchmarking elastic properties of a metal-organic framework (Node days: 1296)

Jelle Wieme:

Tier-1:

Investigating the phases of MIL-53-type materials (Node days: 4644)

Computational exploration of the free energy profile of guest-free M(bdp) (M=Co,Fe) (bdp2-=1,4-benzenedipyrazolate) (Node days: 2940)

Construction of ab initio free energy profile for MIL-53-type materials (Node days: 4200)

Assessing the accuracy of hybrid functionals for the relative stability of a flexible MOF (Node days: 3420)

Understanding the high-pressure behavior of a flexible nanoporous material (Node days: 2430)

Benchmarking elastic properties of a metal-organic framework (Node days: 1296)

Tier-0:

Elucidating the interplay between structure and catalytic activity in nanoporous materials (Preparatory access on Marconi KNL, Marconi Broadwell and Curie TN to benchmark)

Veronique Van Speybroeck:

Tier-1:

Assessing the accuracy of hybrid functionals for the relative stability of a flexible MOF (Node days: 3420)

Full quantum mechanical study of the influence of functionalization and temperature on phase transformations in metal-organic frameworks (Node days: 2450)

Construction of ab initio free energy profile for MIL-53-type materials (Node days: 4200)

Assessing the accuracy of an efficient meta-GGA functional for property predictions of elemental solids (Node days: 640)

Benchmark study of ab initio molecular dynamics simulations for the methylation of HMB (Node days: 4812)

High-throughput screening of an unknown quaternary crystal space (Node days: 4056)

Computational exploration of the free energy profile of guest-free M(bdp) (M=Co,Fe) (bdp2-=1,4-benzenedipyrazolate) (Node days: 2940)

Ab initio umbrella sampling simulations of dehydration of UiO-66 (Node days: 3972)

Electronic properties of 3D nitrogen-containing covalent organic frameworks from first-principles simulations (Node days: 4140)

Structural transformations during dehydroxylation reactions of UiO-66 type metal-organic frameworks; an extension with normal mode analysis (Node days: 2726)

Exploring the kinetics and selectivity of butene cracking using molecular dynamics simulations (Node days: 4864)

Characterizing adsorption properties of C4-C6 alkenes on H-ZSM-5 using molecular dynamics simulations (Node days: 4260)

Unraveling dehydroxylation pathways on UiO-66 type systems with metadynamics (Node days: 3304)

Dynamical first principle benchmark studies on alkene methylation in H-ZSM-5 (Node days: 1400)

Investigating the phases of MIL-53-type materials (Node days: 4644)

Investigating active sites in hydroxylated and dehydroxylated UiO-66 for catalysis of Oppenauer-type oxidation (Node days: 2110)

Shape tuning of CdSe nanostructures by ab initio determination of the anisotropic growth mechanism (Node days: 4752)

Ab initio molecular dynamics study on the role of water in the reaction mechanism during methanol conversion in H-SAPO-34 (Node days: 4880)

Defect engineering in UiO-66: How linker defects affect the electronic structure (Node days: 4116)

DFT study of reaction paths in zeolite-catalyzed 2-hexene cracking (Node days: 4536)

Investigating active sites in hydroxylated and dehydroxylated UiO-66 for catalysis of Oppenauer-type oxidation (Node days: 2110)

Dynamical first principle modelling of zeolite dealumination in H-SSZ-13 (Node days: 3624)

http://dx.doi.org/10.1039/C4RA16800C

http://dx.doi.org/10.1039/C4CE01672F

http://dx.doi.org/10.1002/chem.201500473

http://dx.doi.org/10.1016/j.jcat.2015.01.013

http://dx.doi.org/10.1002/cctc.201402146

http://dx.doi.org/10.1021/cs400706e

http://dx.doi.org/10.1039/c4mh00127c

http://dx.doi.org/10.1039/C3CP54132K

http://dx.doi.org/10.1021/acs.jpcc.5b06809

http://dx.doi.org/10.1016/j.jcat.2015.08.015

http://dx.doi.org/10.1039/C7TA01559C

12. Are the applicants of this application bound by a confidentiality agreement? If so, the abstract of this application will not be published on the website of the FWO / Flemish Supercomputer Center, only the title.

    No

References

[1] H. Furukawa, K. E. Cordova, M. O'Keeffe, O. M. Yaghi, *The chemistry and applications of metal-organic frameworks*, Science **2013**, 341, 1230444

[2] G. Maurin, C. Serre, A. Cooper, G. Férey, *The new age of MOFs and of their porous-related solids*, *Chem. Soc. Rev.* **2017**, 50, 514-516

[3] H.-C. Zhou, S. Kitagawa, *Metal-Organic Frameworks*, Chem. Soc. Rev. **2014**, 43, 5415-5418

[4] A. Schneemann, V. Bon, I. Schwedler, I. Senkovska, S. Kaskel, R. A. Fischer, *Flexible metal-organic frameworks*, Chem. Soc. Rev. **2014**, 43, 6062-6096

[5] S. Horike, S. Shimomura, S. Kitagawa, *Soft porous crystals*, Nat. Chem. **2009**, 1, 695-704

[6] L. Vanduyfhuys, S. M. J. Rogge, J. Wieme, S. Vandenbrande, G. Maurin, M. Waroquier, V. Van Speybroeck, *Thermodynamic insight into stimuli-responsive behaviour of soft porous crystals*, Nat. Commun. **2018**, 9, 204

[7] A. Krylov, A. Vtyurin, P. Petkov, I. Senkovska, M. Maliuta, V. Bon, T. Heine, S. Kaskel, E. Slyusareva, *Raman spectroscopy studies of the terahertz vibrational modes of a DUT-8(Ni) metal-organic framework,* Phys. Chem. Chem. Phys. **2017**, 19, 32099-32104

[8] A. E. J. Hoffman, L. Vanduyfhuys, I. Nevjestić, J. Wieme, S. M. J. Rogge, H. Depauw, P. Van Der Voort, H. Vrielinck, V. Van Speybroeck, *Elucidating the vibrational fingerprint of the flexible metal-organic framework MIL-53(Al) using a combined experimental/computational approach*, J. Phys. Chem. C. **2018**, published online, 10.1021/acs.jpcc.7b11031

Should you have any questions or encounter any difficulties during the electronic submission of an Application, please contact by email:
Associatie KU Leuven: hpcinfo@kuleuven.be
Associatie Universiteit Gent: hpc@ugent.be
Associatie Universiteit Hogescholen Antwerpen: hpc@uantwerpen.be
Associatie Universiteit Hogescholen Limburg: geertjan.bex@uhasselt.be
Universitaire Associatie Brussel: hpc@vub.ac.be
For the other institutions: caroline.volckaert@fwo.be

Fonds Wetenschappelijk Onderzoek
Research Foundation – Flanders

Prof. Karen Hemelsoet
Universiteit Gent
Moleculaire modellering
Technologiepark 903
9052 GENT

uw kenmerk	ons kenmerk	contact	datum
	56225 G063814N	bio@fwo.be cult@fwo.be wt@fwo.be gm@fwo.be med@fwo.be interdisciplinair@fwo.be	13 november 2013

Betreft: **Toekenning project G063814N**

Geachte professor,

De Raad van Bestuur van het FWO heeft uw bovenvermelde aanvraag voor een onderzoeksproject goedgekeurd. De bijgevoegde overeenkomst werd opgesteld volgens het reglement van de Onderzoeksprojecten van het FWO.

Graag wil ik u hierbij feliciteren met de toekenning van het door u aangevraagde onderzoeksproject.

Mag ik u vragen alle exemplaren te ondertekenen, één kopie voor u te bewaren en mij het origineel samen met de overige kopie(ën) per kerende post terug te sturen.

De toegekende toelage is beschikbaar voor zover de financierende overheden hiervoor de nodige middelen vrijmaken. Onderzoekers in de medische of biomedische wetenschappen dienen, voor wat betreft de ethische problemen inzake de betrokkenheid van proefpersonen en/of proefdieren, het advies van de lokale ethische commissie in te winnen.

Slechts uitzonderlijk kan, met akkoord van het FWO, de oorspronkelijke bestemming van dit krediet worden gewijzigd.

Tenslotte vraag ik u met aandrang om op uw publicaties steeds te vermelden dat ze tot stand kwamen met de financiële steun van het FWO. Vacatures gekoppeld aan onderzoeksprojecten kunnen bekend worden gemaakt via de FWO-website. U kan het vacaturebericht aan het FWO bezorgen via communicatie@fwo.be. Aanvullend kan het nuttig zijn om dergelijke berichten ook te publiceren via de Euraxess website van de Europese Commissie: http://ec.europa.eu/euraxess/ .

./..

EGMONTSTRAAT 5
1000 BRUSSEL
TEL +32 2 512 91 10
FAX +32 2 512 58 90
E-MAIL: POST@FWO.BE
WWW.FWO.BE

FWO
VLAANDEREN
Fonds Wetenschappelijk Onderzoek
Research Foundation – Flanders

../..

Voor bijkomende inlichtingen kan u steeds terecht bij uw dossierbeheerder. De e-mailadressen zijn voor de Biologische Wetenschappen: bio@fwo.be, voor de Cultuurwetenschappen: cult@fwo.be, voor de Gedrags- en Maatschappijwetenschappen: gm@fwo.be, voor Wetenschap en Technologie: wt@fwo.be, voor de Medische Wetenschappen: med@fwo.be en voor het Interdisciplinair onderzoek: interdisciplinair@fwo.be.

Hoogachtend,

dr.ir. Elisabeth Monard
secretaris-generaal

Bijlage: overeenkomst

Kopie: Prof. Veronique Van Speybroeck

EGMONTSTRAAT 5
1000 BRUSSEL
TEL +32 2 512 91 10
FAX +32 2 512 58 90
E-MAIL: POST@FWO.BE
WWW.FWO.BE

FWO - Stichting van Openbaar Nut – Ondernemingsnummer 0880.212.840

OVEREENKOMST FWO - ONDERZOEKSPROJECT

TUSSEN

Het Fonds Wetenschappelijk Onderzoek – Vlaanderen, Egmontstraat 5, 1000 Brussel, vertegenwoordigd door dr. ir. Elisabeth Monard, secretaris-generaal, hierna het FWO genoemd, enerzijds,

de onthaalinstelling(en)

Universiteit Gent

EN

de verantwoordelijk woordvoerder van het onderzoeksproject van het FWO met dossiernummer **G.0638.14N**

Karen Hemelsoet
Universiteit Gent

anderzijds,

WORDT HET VOLGENDE OVEREENGEKOMEN

Artikel 1 - De volgende bijlagen maken integraal deel uit van deze overeenkomst:

Bijlage 1: Projectomschrijving – onderzoeksproject G.0638.14N
Bijlage 2: Het FWO - Reglement inzake de onderzoeksprojecten
Bijlage 3: Het FWO - Algemeen reglement

UITVOERING

Artikel 2 - De vermelde promotor en copromotoren verbinden er zich toe het onderzoeksproject, zoals omschreven in de oorspronkelijke kredietaanvraag, uit te voeren.

Artikel 3 - het FWO verbindt er zich toe de subsidies vermeld in bijlage 1 voor de duur van de overeenkomst en overeenkomstig de modaliteiten van het FWO-reglement inzake onderzoeksprojecten, uit te betalen.

Artikel 4 - De onthaalinstelling(en), vertegenwoordigd door hun rector of directeur, bieden de promotor of de copromotor(en) de mogelijkheid om in hun laboratoria of diensten het in deze overeenkomst vermelde onderzoek door te voeren.

DUUR

Artikel 5 - Het onderzoek dat het voorwerp van deze overeenkomst uitmaakt zal worden uitgevoerd gedurende de periode zoals bepaald in bijlage 1 en voor zover de bevoegde instanties het FWO de hiertoe nodige middelen verschaffen.

BEHEER VAN DE MIDDELEN

Artikel 6 – Het financieel opvolgen van de toegekende kredieten wordt toevertrouwd aan de financiële diensten van de universiteit of de wetenschappelijke instelling waaraan de promotoren en copromotoren verbonden zijn.

Artikel 7 - De promotor treedt op als verantwoordelijk woordvoerder t.o.v. de administratie van het FWO en beheert de middelen.

VERANTWOORDELIJKHEID

Artikel 8 - De uitvoering van deze overeenkomst stelt in geen geval het FWO aansprakelijk voor schade aan personen of goederen, die rechtstreeks of onrechtstreeks het gevolg is van het gesubsidieerd onderzoek.

Artikel 9 – Indien het onderzoek experimenten op de menselijke persoon omvat, die vallen onder de bepalingen van de Wet van 7 mei 2004 betreffende de experimenten op de menselijke persoon, dan dient voorafgaandelijk een verzekering afgesloten te worden door de instelling die optreedt als opdrachtgever (universiteit of ziekenhuis).

EINDE VAN DE OVEREENKOMST

Artikel 10 - Elk der partijen kan, mits een opzeg van minimum zes maanden, de overeenkomst verbreken, die bijgevolg voor alle betrokken partijen vervalt.

Artikel 11 - Elke wijziging van het voornaamste onderdeel van het onderzoeksprogramma heeft automatisch tot gevolg dat de overeenkomst binnen de zes maanden wordt stopgezet.

Het origineel is beschikbaar bij het FWO. De overige partijen verklaren een eensluidend afschrift van de overeenkomst en van het reglement inzake Onderzoeksprojecten te hebben ontvangen.

Brussel, 13 november 2013

voor het FWO

Elisabeth Monard
secretaris-generaal

voor de Universiteit Gent

Anne De Paepe
rector

de verantwoordelijk woordvoerder

Karen Hemelsoet
Promotor

BIJLAGE 1: PROJECTOMSCHRIJVING

VOORWERP

I) Nummer van het onderzoeksproject : G.0638.14N

II) Duur van de overeenkomst: 1/01/2014 - 31/12/2017

III) Promotor(en) van het onderzoeksproject :

De promotor, **verantwoordelijk woordvoerder** tegenover het FWO krachtens artikel 9 van het reglement inzake Onderzoeksprojecten van het FWO:

Karen Hemelsoet
Universiteit Gent

Copromotor(en) :

Veronique Van Speybroeck
Universiteit Gent

IV) Omschrijving van het onderzoeksproject :

Berekening van elektronische excitatie- en emissiespectra via geavanceerde moleculaire dynamicasimulaties gecombineerd met TD-DFT.

V) Onthaalinstelling(en) :

Universiteit Gent

VI) Toegekende kredieten :

Universiteit Gent

Personeel:

1 equivalent halftijds wetenschappelijk medewerker(s) van 01/01/2014 tot 31/12/2017
Het voor de bezoldiging van het personeel in 2014 te besteden krediet bedraagt: 30.000 €

Uitrustingskrediet :

2014 : 0 €

Werkingskrediet :

2014: 10.000 €
2015: 10.000 €
2016: 10.000 €
2017: 10.000 €

SOFTWARE LICENSE AGREEMENT FOR THE USE OF VASP5.2 BY ACADEMIC INSTITUTIONS

The Universität Wien, Austria (UW in the following) and
Ghent University, Belgium (UG in the following) [1]
conclude the following agreement:

(1) The UG acquires a non-exclusive academic license for the use of the software-package VASP (Vienna ab-initio simulationprogram) for ab-initio local-density-functional total-energy and molecular-dynamics calculations, versions VASP5.2 and VASP4.6, by the research group Functional Nanomaterials (FUNNANO)[2]. Under this licence the use of the software is restricted to a maximum of six researchers or students, all belonging to this research group and to the same organisatorial unit and working at the same location. The licence does not cover the use of VASP by external collaborators working at other institutions.

(2) The license covers access to the source-code, the program documentation and to the data-base for ultrasoft pseudopotentials and PAW-potentials. UW reserves the exclusive property of the software. It declines any liability for the software and any responsibility for the results of calculations produced with the program. The license does not cover any maintenance service for the software or support for its implementation.

(3) The license is not transferable to another research group of UG without the written agreement of UW. UW reserves the right to refuse authorization of such a transfer. A transfer to a research group not belonging to UG is excluded.

(4)The UG guarantees that the software or parts thereof shall not be made accessible to third parties without the explicit written consent of UW. Access to the code and to the data-base shall be made available through an account of the UW. The UG guarantees that the password for this account will be known only to one contact-person and shall not be communicated to temporary co-workers or guests. All installations of the source code, the executable or the data-base must be copy-protected and accessible only to the authorized users.

[1] Please insert here the name of the institution concluding this agreement with UW. This institution must be a legal person and the agreement must be signed by an authorized representative of this institution. Define the acronym (replacing .....) under which this institution is referred to in the text of the agreement.

[2] Please insert here the name and affiliation research group for which the license is acquired

SOFTWARE LICENSE AGREEMENT FOR THE USE OF VASP5.2 BY ACADEMIC INSTITUTIONS

(5) If VASP is used as the basis of further methodological or software-development, UG agrees to make these additions available to UW. UW will also be entitled to include these additions in further releases of VASP.

(6) In future publications of work performed using VASP, the use of the software shall be properly acknowledged, e.g. in the form
"The calculations have been performed using the ab-initio total-energy and molecular-dynamics program VASP (Vienna ab-initio simulation program) developed at the Institut für Materialphysik of the Universität Wien [1-3]."
[1] G. Kresse and J. Furthmüller, Phys. Rev. B **54**, 11 169 (1996).
If the PAW-version is used, reference will be made to
[2] G. Kresse and D. Joubert, Phys. Rev. **59**, 1758 (1999).
If special features implemented in VASP will have been used, reference should be made to the relevant publications as listed on the VASP home-page.

(7) The UG accepts to pay to UW a licence fee Euro 4.000,- (fourthousand Euro). The licence fee is strongly discounted and applies only to academic institutions with undergraduate teaching.

(8) The licensee will use VASP exclusively for non-profit research. If VASP is used in contractual research in cooperation with or for industry or for military institutions, the financial conditions will have to be re-negotiated.

(9) UW declares that it has the full power and authority to grant the rights granted in this agreement without the consent of any other person, and that the license and use of the software by the licensee will not in any way constitute an infringement or other violation of any copyright, proprietory right or any other rights of any third party.

(10) Any disputes arising from the license agreement are subject to the laws of the Republic of Austria.

(11) The terms of this agreement shall prevail any terms or conditions of the licensee.

SOFTWARE LICENSE AGREEMENT FOR THE USE OF VASP5.2 BY ACADEMIC INSTITUTIONS

For the Universität Wien:

Jürgen Hafner
Fakultät für Physik, Universität Wien
Sensengasse 8/12, A-1090 Wien, Austria

Date

For the UG

Name (in print): Michel Waroquier
Institution: Faculty of Sciences, Ghent University

Address: Technologiepark 903, BE-9052 Zwijnaarde, Belgium

Date: 26 January 2010

For the research group entitled to use VASP5.2:

Name (in print): Veronique Van Speybroeck (FUNNANO)

VASP benchmark on BrENIAC

Kurt Lejaeghere – Arthur De Vos – Sam De Waele

1. Background

BrENIAC contains 580 nodes with 28 cores each, which are of the Broadwell E5-2680v4 type. Each node has 128 or 256 GB RAM and consists of 2 NUMA regions of 14 cores. The network is connected through an Infiniband EDR 2:1 connection.

To benchmark the performance of VASP (module VASP/5.4.1-intel-2016a) on BrENIAC, three very different test systems were considered:

- A doubled $Fe_{16}N_2$ unit cell with one N atom removed
  (35 atoms, 224 bands, 196 irreducible k-points, vasp_std)
  designated by tag METAL

- a Ge semiconductor surface with Pt atoms adsorbed
  (100 atoms, 336 bands, 8 irreducible k-points, vasp_std)
  designated by tag SEMI

- the metal organic framework UiO-66 with two missing linker defects
  (420 atoms, 1120 bands, 1 irreducible k-point, vasp_gam)
  designated by tag PORE

2. Optimal parallelization on 1 node

VASP has the possibility to parallelize over k-points and, for a given k-point, over electronic bands. In general, parallelization over k-points is more efficient, since it requires almost no communication between subprocesses. However, it also substantially increases the memory requirements, since the calculation of the wavefunction at 1 k-point is based on knowledge of all energy levels at that k-point. The memory needed therefore increases when more k-points are computed simultaneously (KPAR). Analogously, parallelization within 1 band occurs by grouping blocks of plane waves in diagonalization routines and allows spreading the memory even thinner. It is more favourable for the memory requirements to devote more cores to a single electronic band (NCORE), equivalent with fewer bands per node, but this behaviour is less distinct.

Table I: Walltime of a calculation of METAL, SEMI and PORE on 1 node, depending on the parallelization settings (number of k-points treated simultaneously, KPAR, and number of cores per band, NCORE).

wall time METAL [s]	NCORE = 1	NCORE = 7	NCORE = 14	NCORE 28
KPAR = 1	9863	6772	6402	6924
KPAR = 2	8654	6515	5601	
KPAR = 4	8435	6369		

wall time SEMI [s]	NCORE = 1	NCORE = 7	NCORE = 14	NCORE 28
KPAR = 1	777	770	710	687
KPAR = 2	775	769	685	
KPAR = 4	720	729		

wall time PORE [s]	NCORE = 1	NCORE = 7	NCORE = 14	NCORE 28
KPAR = 1	4900	4091	4059	3809

Table II: Memory usage per core for a calculation of METAL, SEMI and PORE on 1 node, depending on the parallelization settings (number of k-points treated simultaneously, KPAR, and number of cores per band, NCORE).

mem METAL [MB]	NCORE = 1	NCORE = 7	NCORE = 14	NCORE 28
KPAR = 1	1499	933	883	896
KPAR = 2	2002	1486	1446	
KPAR = 4	3091	2601		

mem SEMI [MB]	NCORE = 1	NCORE = 7	NCORE = 14	NCORE 28
KPAR = 1	328	199	193	187
KPAR = 2	424	307	297	
KPAR = 4	643	525		

mem PORE [MB]	NCORE = 1	NCORE = 7	NCORE = 14	NCORE 28
KPAR = 1	736	406	369	352

In terms of timing, we see that a higher order of k-point parallelization reduces the required wall time. However, it is not immediately clear which band parallelization is optimal. Many systems benefit from NCORE = 7 or 14, in line with the use of 1 shared memory per band, but for specific systems and number of plane waves, NCORE = 1 may become most favourable (e.g. when increasing the number of plane waves for SEMI). We can only conclude that the *best tradeoff between k-point parallelization and band parallelization needs to be tested for the particular system at hand*. This can be done quite easily, using only a few test calculations (e.g. NCORE = 1, 7, 14 and 28 at KPAR = 1 on 1 node for a representative

system and cutoff energy) and for the optimal configuration taking *KPAR as high as possible*. In addition, the guidelines for memory should be taken into account as well, since *large systems or systems with many k-points (like METAL) may suffer from too high memory requirements*. Finally, the NSIM tag does not matter too much, but NSIM = 1 is strongly discouraged, as it drastically increases the computation time (default is NSIM = 4).

In comparison to Ghent clusters, the (empty) BrENIAC machine performs exceptionally well. For the SEMI system, timings are about two times as good as the best wall times ever achieved on Muk (1378 s in 2013). The same is true in comparison to golett, one of the most recent machines on the UGent HPC (1300 s in 2016). These numbers were scaled to be comparable to the 28 cores per node of BrENIAC. Note, however, that the wall time on golett was measured on the machine in full loading (whereas the BrENIAC machine was almost empty), which has a large impact on the speed of the calculations.

3. Intra- and multinode scaling

Figure 1: Intra- and multinode scaling of the wall time for the METAL system.

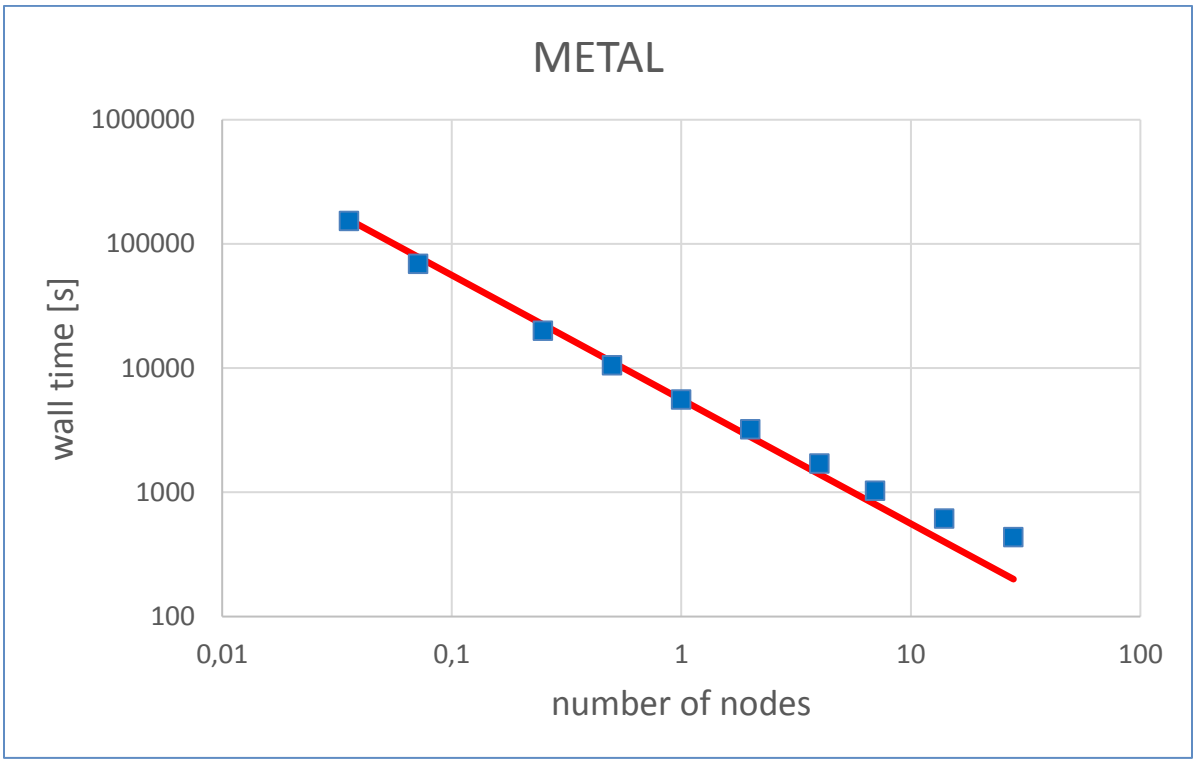

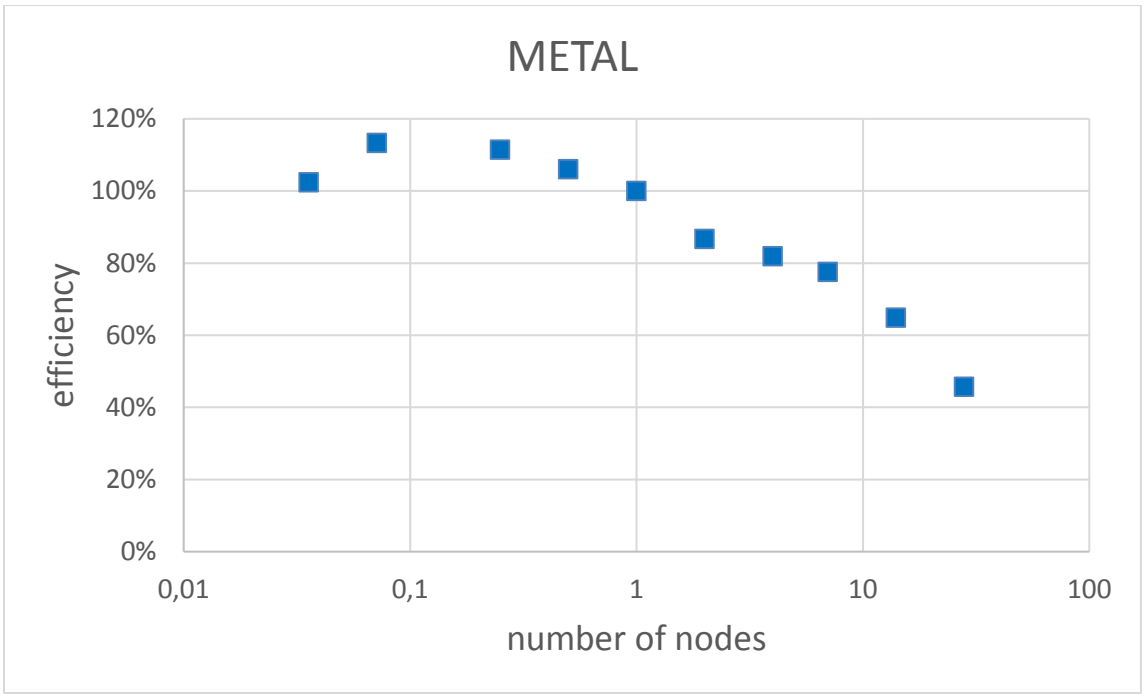

Figure 2: Intra- and multinode scaling of the wall time for the SEMI system (BrENIAC 2016 and Muk 2013). The red and orange lines denote the ideal scaling behaviour.

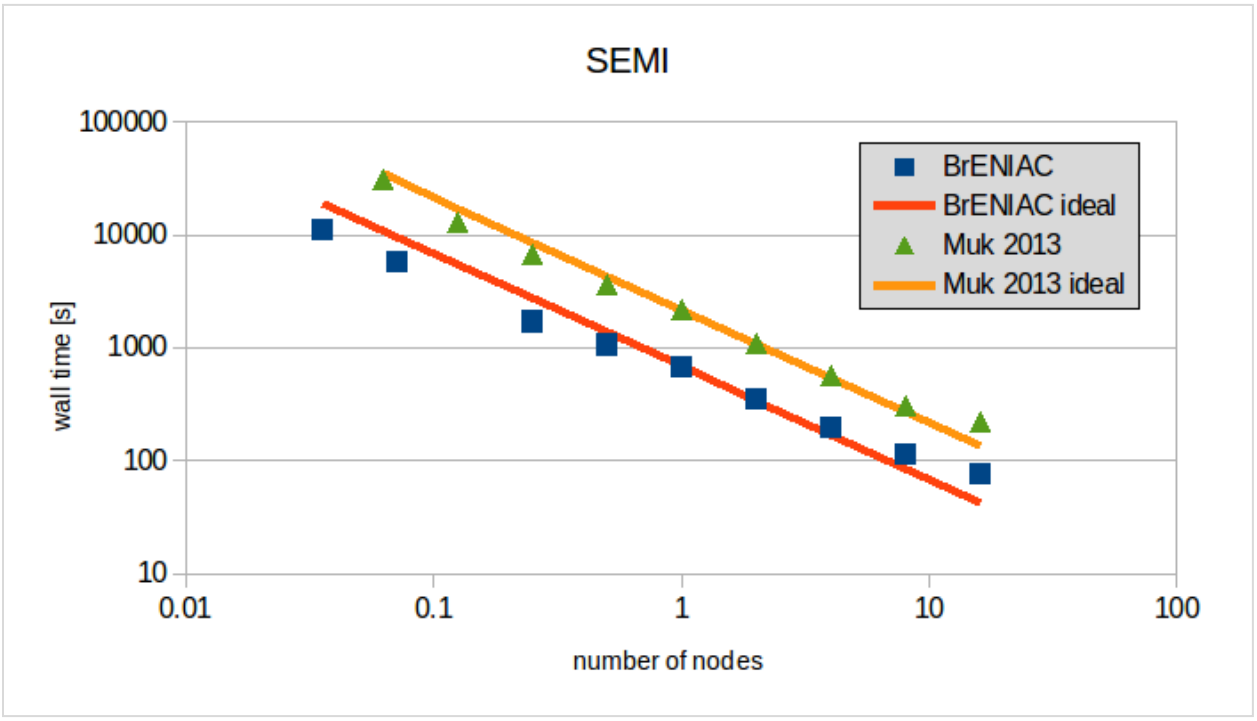

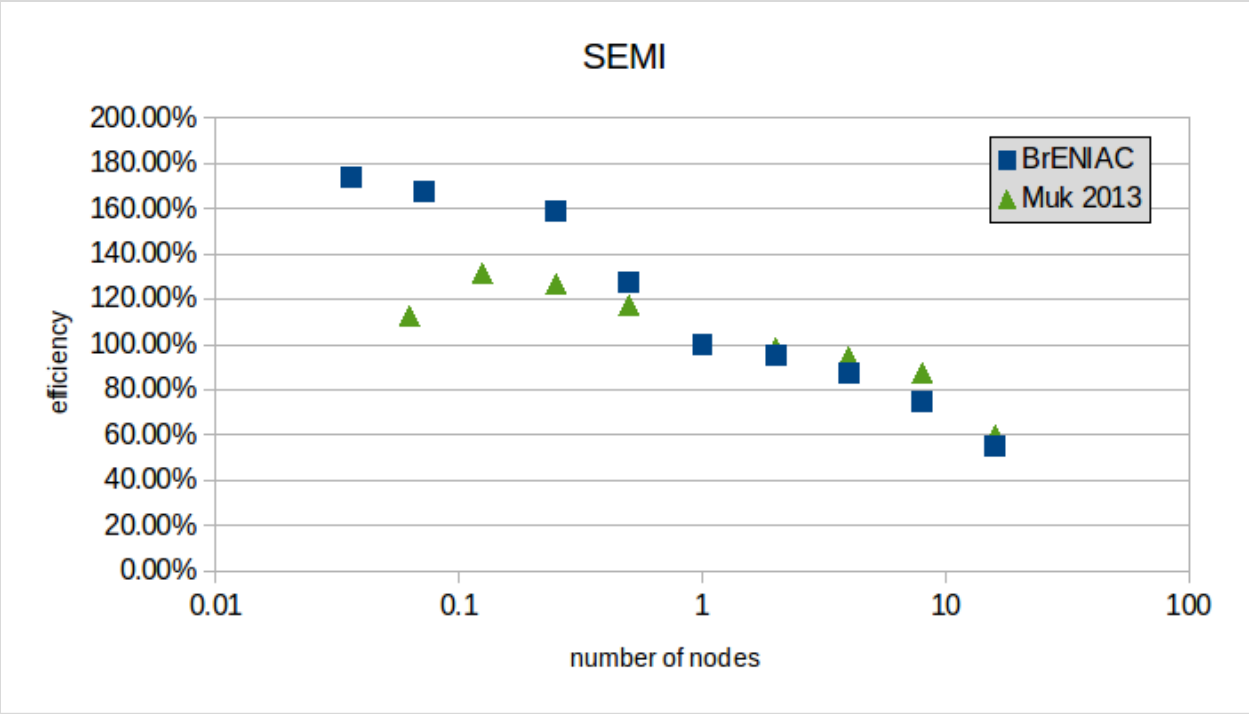

Figure 3: Intra- and multinode scaling of the wall time for the PORE system. The red line denotes the ideal scaling behaviour.

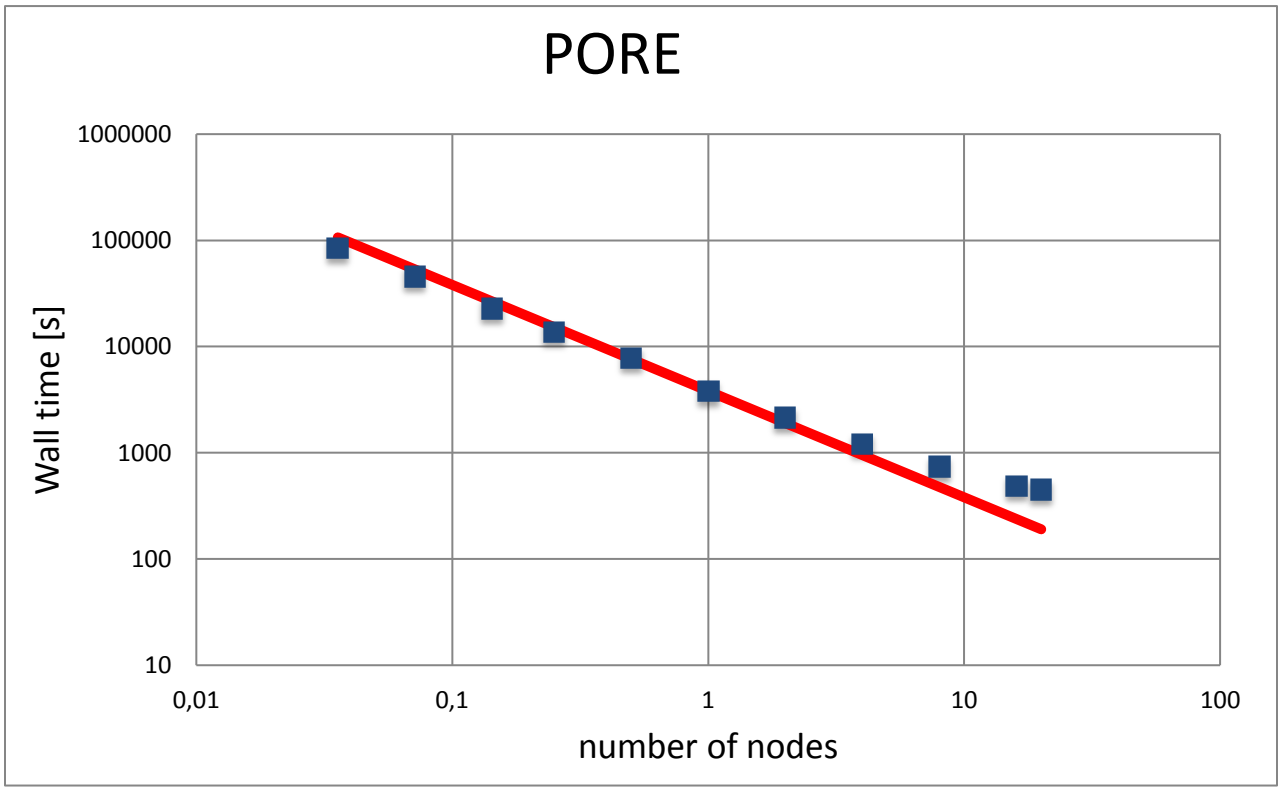

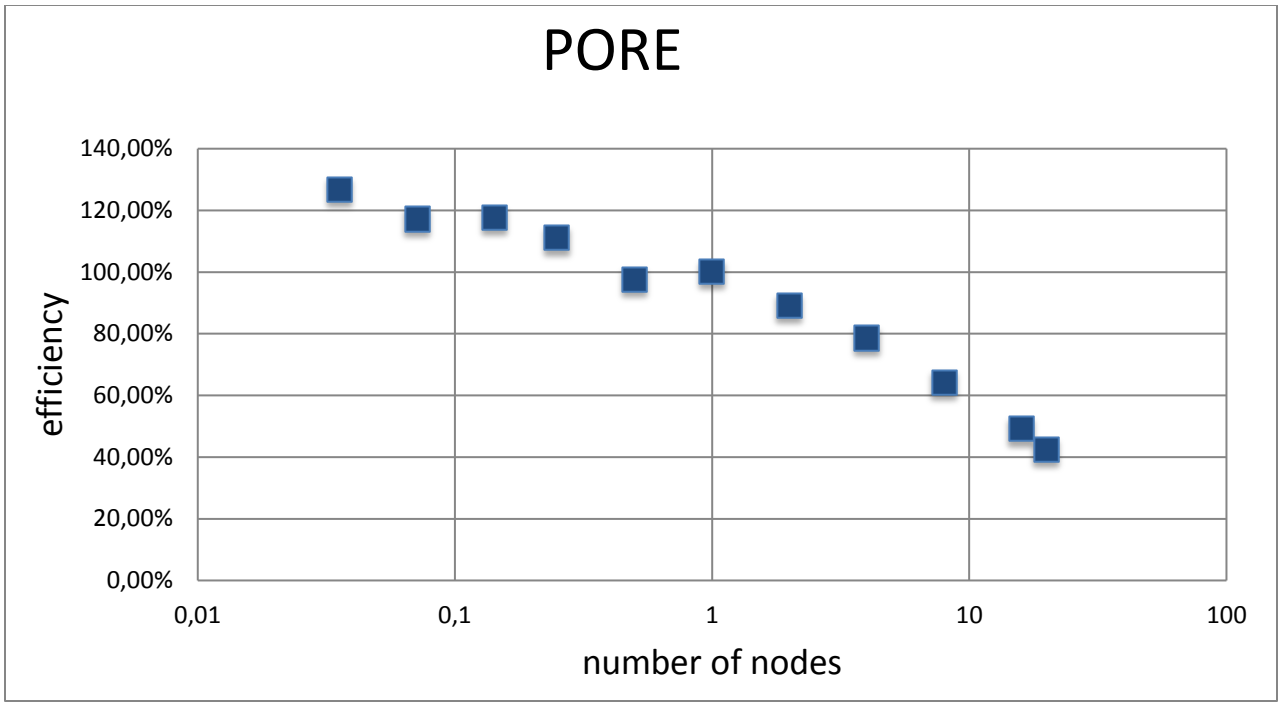

By performing the calculations on a few cores up to multiple nodes, we note that the computational efficiency proceeds in 2 steps. On the one hand, it remains most efficient to perform calculations on 1 or a few cores, and up to the use of an entire node, the efficiency steadily declines. This intranode scaling differs significantly for different systems, however, with poor scaling for SEMI and almost ideal scaling for METAL. *The multinode scaling, however, is quite efficient*, and parallelization over 8 nodes leads to wall times that are still 60-80 % of the efficiency of a single node. Beyond 16 nodes, efficiency drops below 50 %, and calculations are only advisable if they cannot be calculated within 72h on fewer nodes. This behaviour is similar for all tested systems, despite their large diversity, and in line with tests on Muk in 2013 (see Figure 2). We may therefore conclude that it is *not meaningful to perform such scaling tests time and again; only the optimal parallelization settings on 1 node need to be examined when considering a new system*.

As a final note concerning the parallelization settings in multinode calculations, we remark that it is best not to parallelize 1 k-point or 1 band over multiple nodes. Using KPAR equal to the number of nodes (or higher) decreases the computational load significantly, because k-point parallelization requires little communication. For the SEMI system on 2 nodes, for example, a k-point-parallelized calculation (KPAR = 4, NCORE = 14) takes 359 s, while a band-parallelized calculation (KPAR 1, NCORE = 14) takes 408 s. For the METAL system, the difference is huge: 14 002 s for KPAR = 4 and NCORE = 14, compared to 43 229 s for KPAR = 1 and NCORE = 14.