Enclosure 1. 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:

Full quantum mechanical study of the influence of functionalization and temperature on phase transformations in metal-organic frameworks

Name and first name of the applicant:

ROGGE Sven

Institution:

Ghent University

Research group / department:

EA17 / Center for Molecular Modeling

Title / position:

ir. / PhD-FWO Fellow

e-mail address:

Sven.Rogge@UGent.be

Total computing time that is needed, in node days:

2450

Total disk storage that is applied for (in GiB):

245 GiB

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

Advanced molecular dynamics simulations for a full characterization of mechanical properties of metal-organic frameworks (ir. Sven Rogge, FWO application number 11ZW417N).

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 crystalline, nanoporous materials consisting of inorganic metal nodes connected through organic linkers, and showing a lot of potential in a wide range of applications. Here, we will investigate the MIL-53 family, since these MOFs may exhibit two (meta)stable phases for which the stability is not only affected by external stimuli such as temperature and pressure, but also by the possible substituents on the organic linker. As a result, transitions between these two stable states ('breathing' behavior) occur at pressures and temperatures that can be engineered for specific applications by varying the nature and composition of the substituents on the organic linkers (see Fig. 1).

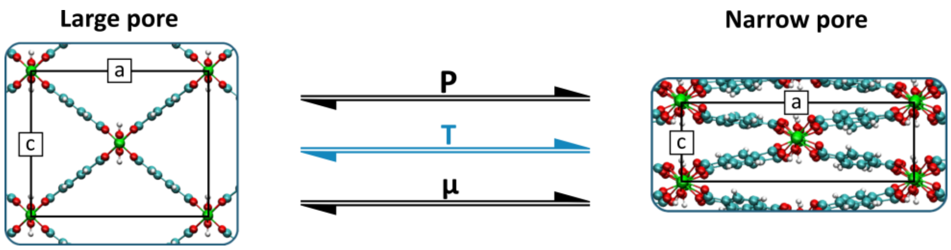

Fig. 1: Breathing of the nonfunctionalized MIL-53(Al) under external stimuli.

Several computational investigations based on a classical description of the potential energy surface have already been performed on the MIL-53 family. However, the number of studies with a fully quantum mechanical treatment of these materials remains scarce. Recently, we pioneered in this respect by constructing MIL-53(Al) free energy profiles using a full quantum mechanical description of the potential energy profile. In this proposal, we will investigate for the first time the peculiar breathing behavior of the MIL-53 family using molecular dynamics simulations based on umbrella sampling to explore the free energy space in terms of both temperature and linker functionalization, following a protocol similar to earlier, force-field based work (DOI: 10.1021/acs.jctc.5b00748, 10.1021/acs.jpcc.6b04422). Specifically, four combinations of the chemically very different -NO_2 (electron-withdrawing) and -NH_2 (electron-donating) groups will be considered as linker substituents to study the importance of the distribution of these functionalizations, and three temperatures will be considered. The effect of temperature and linker functionalization on the breathing behavior of the MIL-53 family will help to gain microscopic insight in the breathing behavior and to identify which effects play a dominant role in this breathing, such that guidelines can be formulated that will steer the ongoing search for future MOF applications as nanoshock absorbers and nanodampers.

3. Provide an engaging abstract (10 lines) for scientific communication on the website in layman's terms. See also item 12 of this application form.

Metal-organic frameworks (MOFs) are a new class of materials which show a great promise as, for instance, nanoscale sensors, nanodampers, and nanoshock absorbers. We will examine a particular subset of these MOFs: the MIL-53 series with different linker functionalizations. A characteristic feature of these MOFs is that external stimuli such as temperature can trigger reversible phase changes. In a previous study, we pioneered the study of this breathing behavior by constructing free energy profiles with the aid of a fully quantum mechanical description. In this proposal, we aim to further expand our knowledge of the breathing behavior by studying the effect of temperature and linker functionalization. A full understanding of these effects on the free energy profiles are essential to determine the interesting materials for future practical applications.

4. Financing institution or channel, financing the research project in full or in part (FWO, BOF, IWT, EU, ...): 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-PhD fellowship 11ZW417N (ir. Sven Rogge; see attachment)

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 in the past (Tier-0/Tier-1/Tier-2 infrastructure in Belgium and abroad)

ROGGE Sven

Ghent University
EA17 / Center for Molecular Modeling
ir. / PhD-FWO Fellow
vsc40686
Experience with the local CMM clusters and the HPC UGent Tier-2 clusters

DEMUYNCK Ruben

Ghent University
EA17 / Center for Molecular Modeling
ir. / PhD fellow
vsc40923
Experience with the local CMM clusters, the HPC UGent Tier-2 clusters, and the Tier-1 machines Muk and BrENIAC

VANDENBRANDE Steven

Ghent University
EA17 / Center for Molecular Modeling
ir. / PhD-FWO fellow
vsc40685
Experience with the local CMM clusters, the HPC UGent Tier-2 clusters, and the Tier-1 machines Muk and BrENIAC

VAN SPEYBROECK Veronique

Ghent University
EA17 / Center for Molecular Modeling
prof. dr. ir. / full professor
vsc40021
Experience with the local CMM clusters, the HPC UGent Tier-2 clusters, and the Tier-1 machines Muk and BrENIAC

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).

The strategy of computing free energy profiles consists of running a series of molecular dynamics (MD) simulations in parallel. For example, for the well-documented MIL-53(Al), one requires to run approximately 35 independent simulations in parallel to obtain a free energy profiles as a function of the volume. Moreover, the MD simulations must be sufficiently long to obtain an accurate free energy profile, requiring about 50 ps (10^5 MD steps). Each MD step consists of running an *ab initio* single-point calculation. Taking into account the size of the MIL-53-type unit cells (76 atoms), such an *ab initio* single-point calculation becomes very demanding. Moreover, such a series of about 35 independent MD simulations need to be carried out for each temperature and each linker functionalization.

The time and length scales of the aforementioned simulations, combined with the fact that a large number of these simulations need to be performed, require the use of many multinode jobs with a long walltime, for which the Tier-1 infrastructure is best suited. When only the Tier-2 infrastructure can be used, this project will not be able to compete with other outstanding research groups working on similar projects.

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
- 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.
(max. 1 A4 Arial 12).

The proposal can be divided into two main tasks:

1. Umbrella sampling for functionalized linkers

The general idea of constructing a free energy profile is to run several molecular dynamics simulations independently, where all the simulations should cover a certain volume range of the entire profile. To that end, we require initial structures along the profile. Therefore, we start with optimizing the geometry of the input structure at several volumes along the trajectory. Based on the force field benchmark simulations, we were able to conclude that 35 initial simulation points yield an accurate free energy profile. A procedure of fixed volume relaxations starting from the experimental structure was recently proposed to obtain the correct structure of other flexible nanoporous materials. These calculations have to be performed at highly accurate settings to capture all fine details that govern the peculiar behavior of MIL-53-type materials as was explained previously. Per task, 2 nodes and 5 days will be needed to reach the required accuracy (using restart files to go beyond the maximum walltime of 3 days), which need to be multiplied by the number of simulation points per free energy profile (35) and the number of functionalization distributions (4: one pure NO_2, one pure NH_2, and two different distributions of the NH_2/NO_2 combination).

2. Umbrella sampling for a set of temperatures

Due to the high computational cost, we can only construct free energy profiles for MIL-53(Al) at a limited number of temperatures. Liu *et al.* observed a transition between the large-pore and the closed-pore structure for MIL-53(Al) around 125 K, while the opposite transition was observed at 375 K (DOI: 10.1021/ja803669w). We choose to construct free energy profiles at temperatures of 100 K, 300 K, and 400 K. Hence, we expect to characterize a transformation to the closed-pore structure at a 100 K simulation, a transformation to the large-pore phase at 400 K and a bistable structure at 300 K. The free energy profiles will be constructed using umbrella sampling, requiring 35 independent MD simulations per profile. The simulation time, and hence the required resources, are the same as above. Per task, 2 nodes and 5 days will be needed to reach the required accuracy (using restart files to go beyond the maximum walltime of 3 days).

Table 1

	Node day calculation							Storage volume estimate	
Computational task	# of such tasks	Wall clock time (days) per task	# Tier-1 nodes per task	# node days per task	# CPU cores per task	Memory usage (GiB) / node per task	OpenMP / MPI / hybrid / vSMP	Tier-2 DATA/HOME volume (GiB) + number of files	Tier-1 SCRATCH volume (GiB) + number of files
Umbrella sampling for functionalized linkers	4 x 35	5	2	1400	56	5	Hybrid OpenMP/MPI	1 GiB x 140 (600 files)	1 GiB x 140 (600 files)
Umbrella sampling for a set of temperatures	3 x 35	5	2	1050	56	5	Hybrid OpenMP/MPI	1 GiB x 105 (450 files)	1 GiB x 105 (450 files)
TOTAL				**2450**				**245 GiB**	**245 GiB**

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)

   Provide the results of scaling 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 scaling plot such as table 2 and plot 1 below (max. 2 A4 in Arial 12).

VASP

http://www.vasp.at/
non-exclusive academic license (see attachment)
Available on BrENIAC

ASE

http://wiki.fysik.dtu.dk/ase/
GNU Lesser General Public License v2.1
Available on BrENIAC

YAFF

http://github.com/molmod/yaff
GNU General Public License
Available on BrENIAC

The main computational cost of this project are the VASP calculations. YAFF and ASE solely serve to automate the VASP calculations in an intelligent fashion. Hence, the overhead cost of this software is negligible and we limit ourselves to providing the scaling properties of VASP.

We did not perform new multinode scaling tests for the current proposal, since our extended benchmark studies (see attachment and Table 2) show a close similarity in scaling behavior for all systems despite their large differences, *i.e.* both for large and smaller systems, and both for metals and nonmetals. The multinode scaling test (Table 2) for the current proposal are based on the scaling test of UiO-66, another material in the MOF class of materials. This test shows that for 2 or 4 nodes, the intended parallelization, scaling is very good.

The results of some VASP scaling tests, performed on BrENIAC, are available in Table 3 and Plot 1. Tests on 1 node (Table 3) were performed specifically with MIL-53(Al), the material under study in this work, employing the same settings as will be done during the project. We tested the optimal combination of parallelization over k-points or parallelization over electronic bands (Table 3), as our extended benchmark (see attachment) showed the optimal parallelization settings to be highly system-dependent. In line with previous benchmarks, which showed a high k-point parallelization to be favorable, we will perform all calculations with NCORE = 7 and KPAR as high as possible.

Table 2

# nodes	# cores	absolute timing (s)	speedup	efficiency
1	28	3810	1.00	1
2	56	2061	1.85	0.924
4	112	1176	3.24	0.810

8	196	720	5.29	0.661
16	448	517	7.36	0.460
20	560	478	7.97	0.399
35	980	432	8.82	0.252

Plot 1

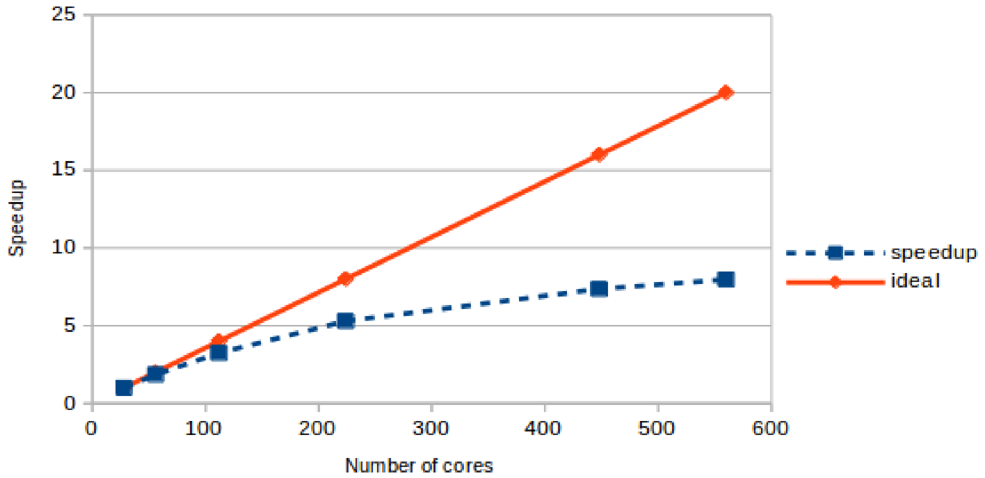

Table 3

absolute timing on 1 node with KPAR = 1 [s]			
NCORE = 1	NCORE = 7	NCORE = 14	NCORE = 28
59.980	46.482	46.085	49.954
absolute timing on 1 node with KPAR = 4 [s]			
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, …)?

The calculations are intended to be performed in the timeslot February 2017 – September 2017 with a similar monthly usage.

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.

Sven Rogge

None

Ruben Demuynck

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

Steven Vandenbrande

Insight into the stacking of 2D COFs from ab initio and force-field calculations (Node days: 2200 days)

Several publications based on the results of these two Tier-1 projects are currently in preparation and will be submitted in 2017

Veronique Van Speybroeck

Computational exploration of the free energy profile of guest-free M(bdp) (M=Co,Fe) (bdp^{2-} = 1,4-benzenedipyrazolate) (Node days: 2940 days)

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

Dynamical kinetic study of zeolite-catalyzed reactions. (Node days: 4371 days)

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

Structural transformations during dehydroxylation reactions of UiO-66 type metal-organic frameworks. (Node days: 4720 days)

Modeling aldol condensations in metal-organic frameworks with hybrid functional calculations. (Node days: 2304 days)

Unraveling reaction pathways on UiO-66 type systems with metadynamics. (Node days: 4432 days)

The electronic and magnetic structure of breathing metal-organic frameworks. (Node days: 4725 days)

Molecular dynamics study of pentene in H-ZSM-5: towards a better estimate of adsorption enthalpies. (Node days: 1824 days)

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

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

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

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

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

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

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

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

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

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

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

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

Articles:

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

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.

    yes

Should you have any questions or encounter any difficulties during the electronic submission of an Application, please contact by e-mail:
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

fwo
Opening
new
horizons

De heer Sven Rogge
Raspaillebosstraat 54
9506 GERAARDSBERGEN

uw kenmerk

contact
wt@fwo.be

datum
27 juni 2016

ons kenmerk
117533
11ZW417N LG

betreft **Hernieuwing mandaat aspirant 2016-2018**

Geachte heer,

In opvolging van de toekenningbrief bezorg ik u in bijlage twee overeenkomsten houdende de toekenning van een beurs.

Indien u het mandaat aanvaardt, stuurt u één exemplaar aangevuld en ondertekend vóór **31 augustus 2016** terug naar het FWO.

Ik wens u veel succes toe in uw verdere onderzoeksloopbaan.

Met vriendelijke groeten,

Danny Huysmans
Directeur Intern Beheer

bijlagen: - 2 overeenkomsten

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.

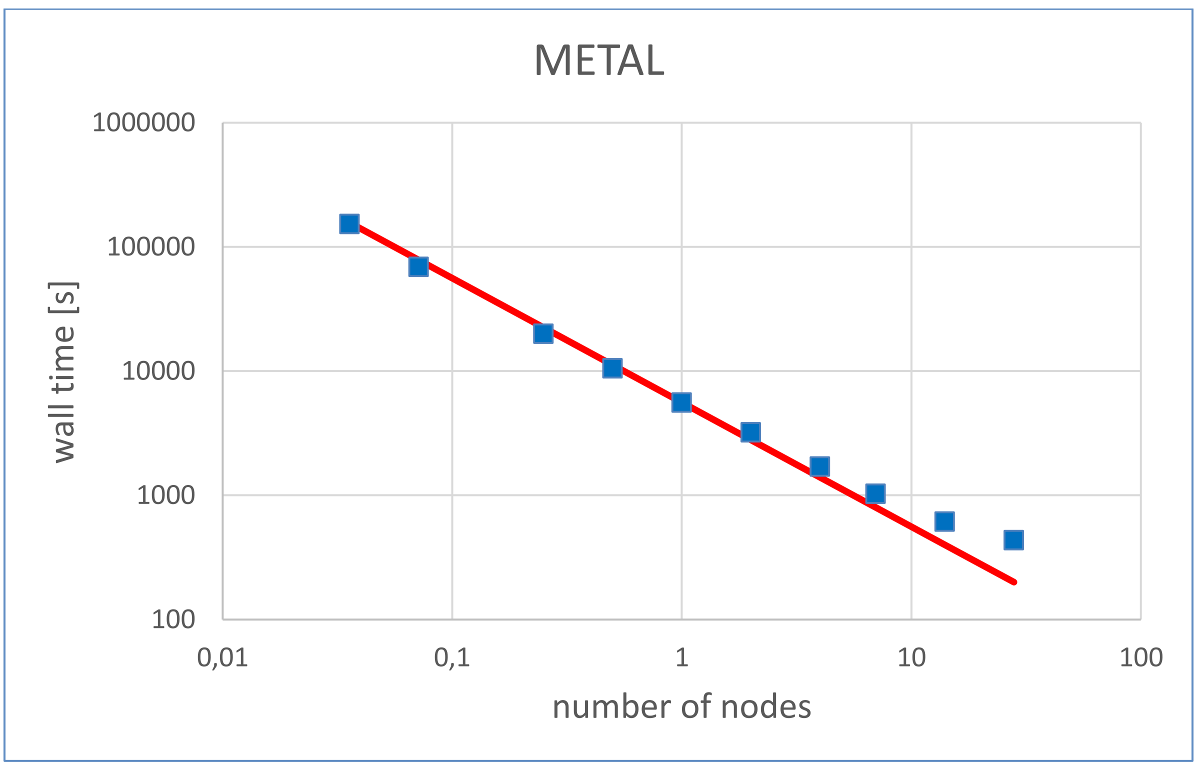

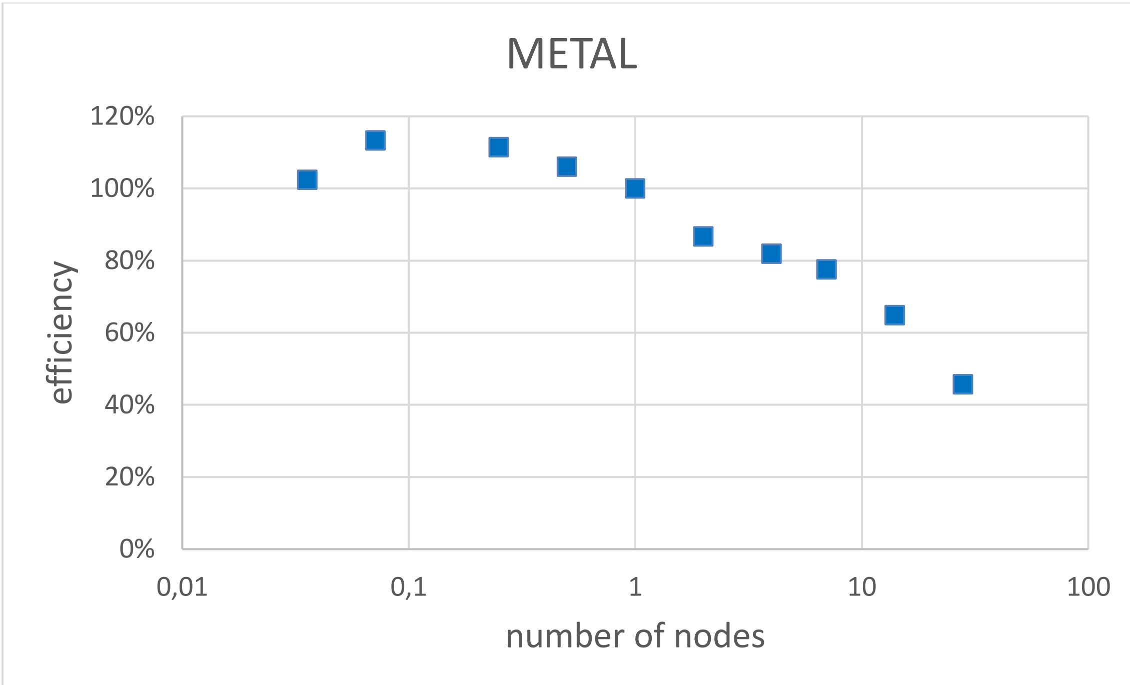

Figure 2: Intra- and multinode scaling of the wall time for the SEMI system (BrENIAC 2016 and Muk 2013). The red and orange lines denote the ideal scaling behaviour.

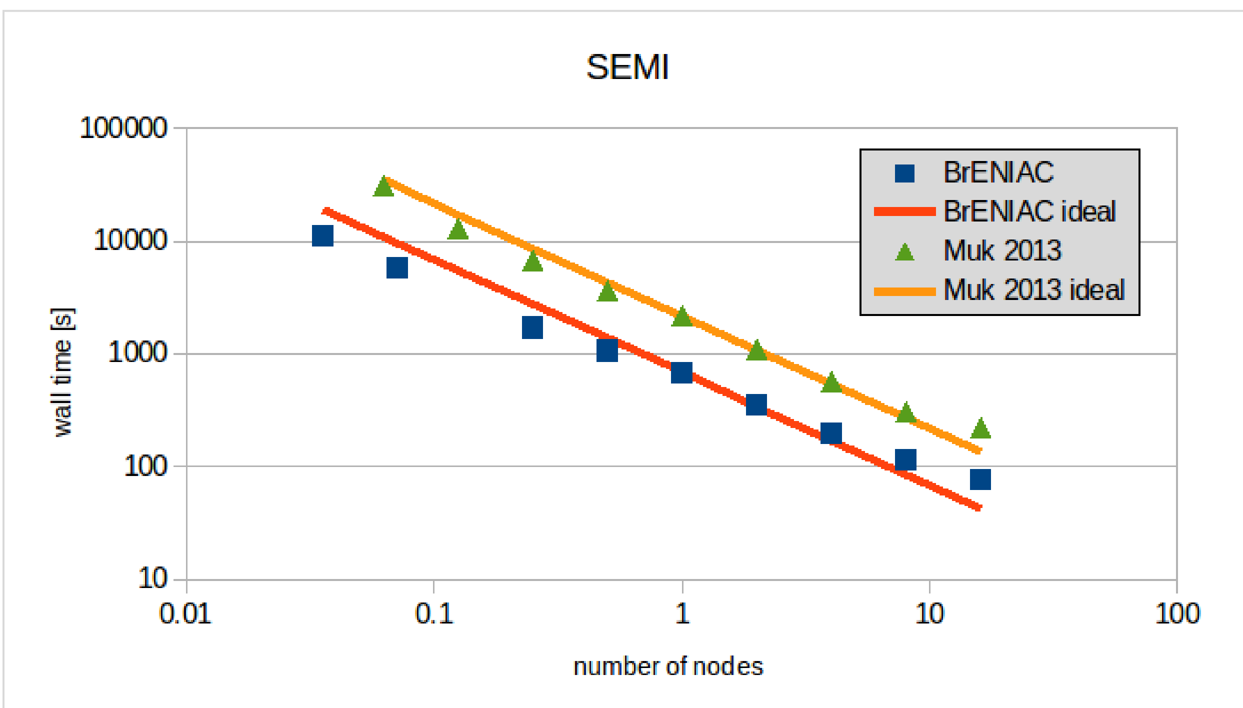

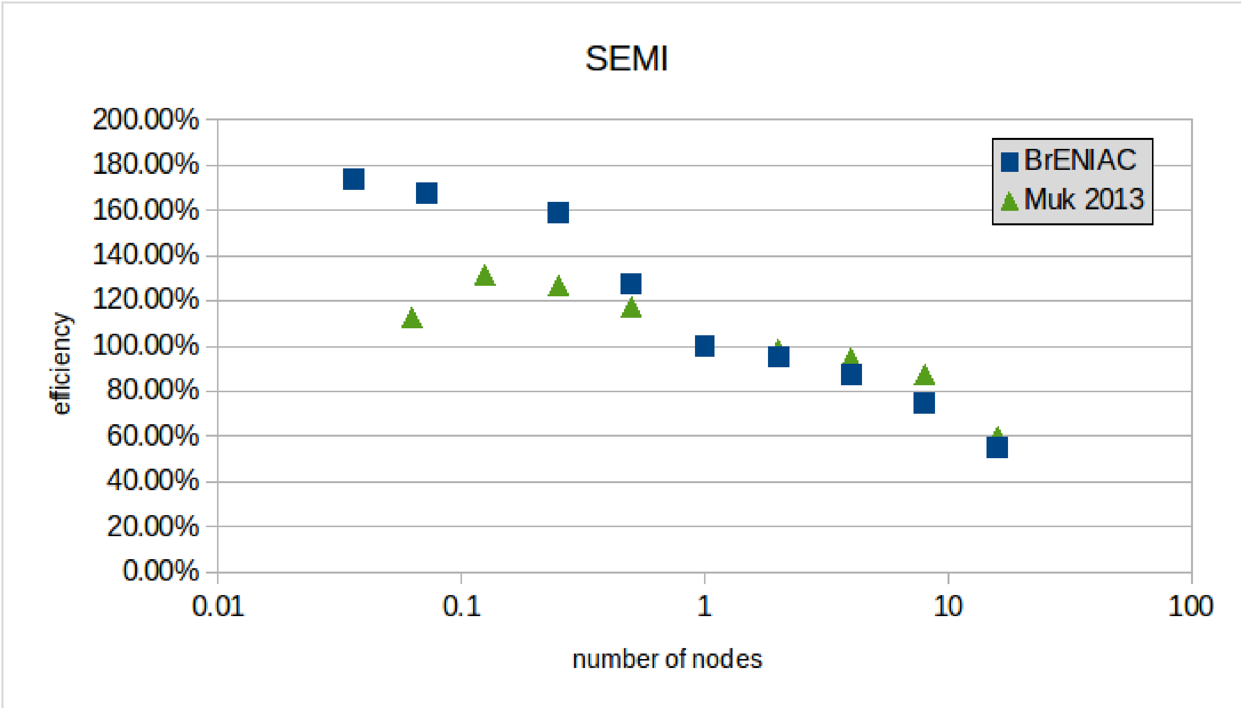

Figure 3: Intra- and multinode scaling of the wall time for the PORE system. The red line denotes the ideal scaling behaviour.

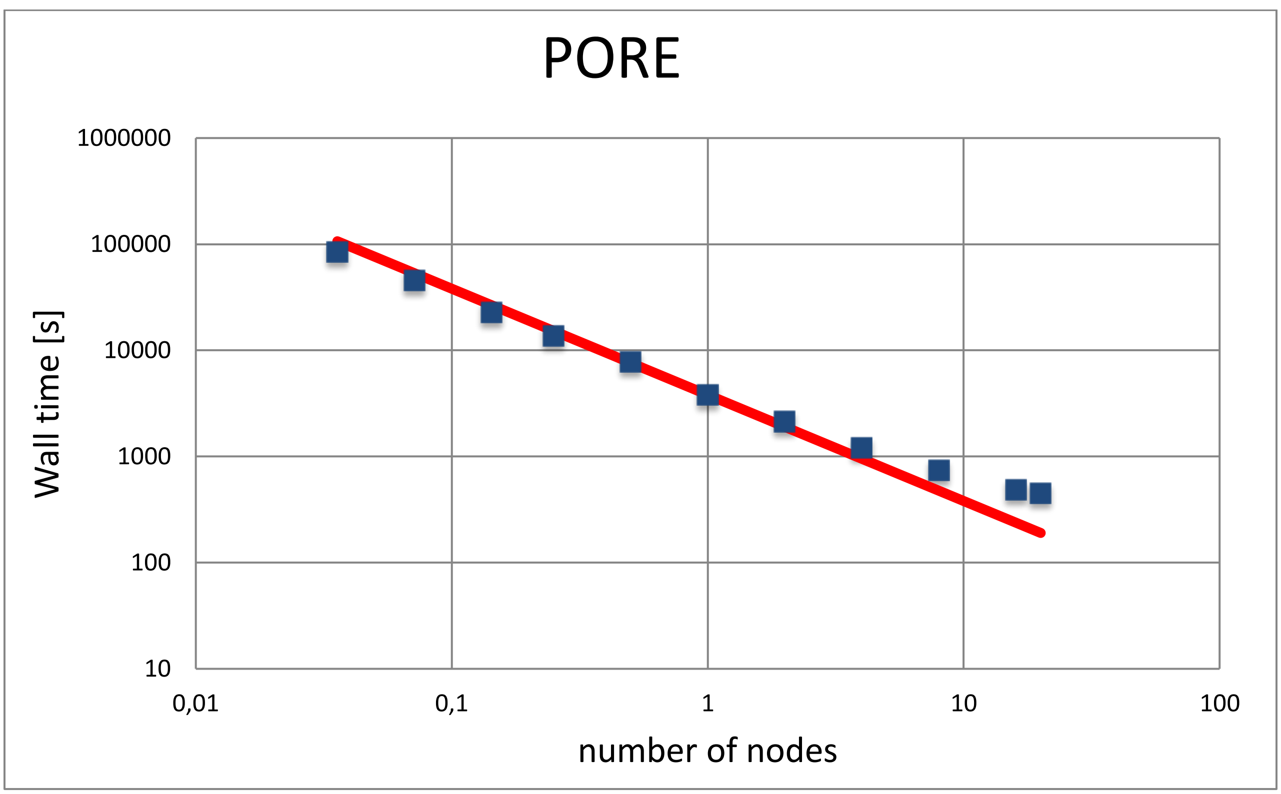

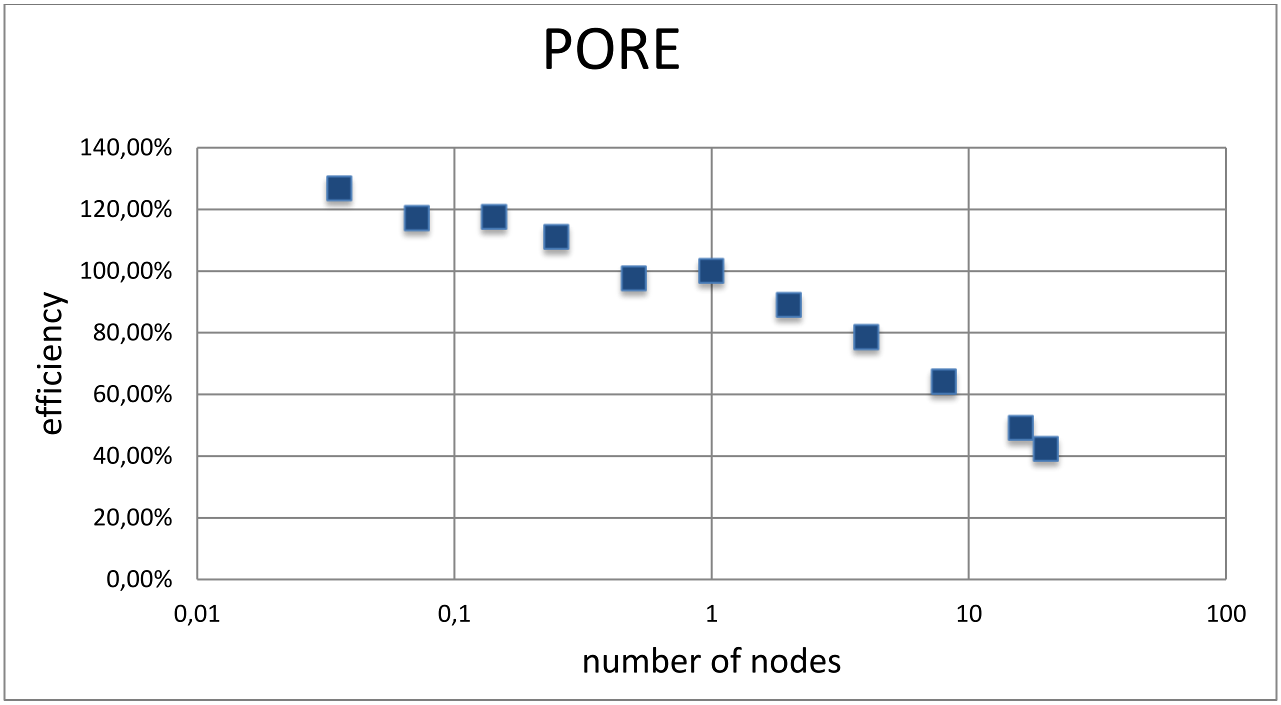

By performing the calculations on a few cores up to multiple nodes, we note that the computational efficiency proceeds in 2 steps. On the one hand, it remains most efficient to perform calculations on 1 or a few cores, and up to the use of an entire node, the efficiency steadily declines. This intranode 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.