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:

Advanced electronic structure calculations of a metal-organic framework

Name and first name of the applicant:

Lejaeghere Kurt

Institution:

Ghent University

Research group / department:

Center for Molecular Modeling / EA17

Title / position:

FWO postdoctoral researcher

Email address:

Kurt.Lejaeghere@UGent.be

Total computing time that is needed, in node days:

1022 node days

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

1514 GiB (3090 files)

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:

Unravelling energetic and mechanical properties of framework materials by means of advanced electronic-structure methods

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) show unrivalled potential for many specialized applications, varying from gas adsorption to luminescence and sensing. They possess a modular make-up, with metal nodes interconnected by inorganic linkers, which provides them with a high tunability. This makes it possible to adapt MOF properties for almost any desired purpose.

Because of the huge variation offered by MOFs, rational design remains a significant challenge. Computational research is able to mitigate this problem, as it offers insight on interactions within the material down to the atomic scale. However, current computational techniques fall short and fail to describe the delicate interplay between various electronic energy contributions. Especially London dispersion is found to strongly affect a particular subclass of MOFs: the breathing frameworks. These materials possess two stable polymorphs, a small-volume or "narrow-pore" phase, and a large-volume or "large-pore" phase. The relative stability of these two crystal structures is often strongly dependent on the dispersion interaction between the linkers, which is much stronger in the narrow-pore phase.

The current proposal aims to provide unprecedented understanding in this complex interplay between dispersion and other electronic effects by applying the random-phase approximation (RPA) to the correlation energy. In particular, we will determine the RPA energy of MIL-53(Al), the prototype breathing MOF, as a function of volume. RPA is an advanced many-electron technique, which seamlessly integrates short-range interaction with long-range dispersion. It has never been applied to breathing MOFs before, but comes at a computationally very high cost. This Tier1-proposal therefore goes beyond the state of the art for the description of MOFs with the prospect of having a tremendous impact on the field.

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) show unrivalled potential for many specialized applications. Breathing MOFs, for example, display a stimulusresponsive phase transition between a small- and large-volume structure while retaining their crystalline topology. This makes them promising for use in dedicated sensors or actuators. However, computational design of these materials is hindered because current techniques are unable to fully capture all contributions to the electronic energy sufficiently accurately. This project aims to apply advanced electronic-structure methods to resolve the discussion on the relative phase stability of breathing metal-organic frameworks.

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:

Lejaeghere Kurt

- Ghent University
- Center for Molecular Modeling / EA17
- FWO postdoctoral fellow
- Experience with local CMM clusters (2009-current), local clusters of the Computational Physics Group of the University of Vienna (2015), HPC UGent Tier2 clusters (2010-current), Tier1 machines muk (2012-2016) and BrENIAC (2016-current), and Tier0 PRACE infrastructure (Curie thin nodes, Marconi Broadwell, 2017). VSC number: vsc40323.

Wieme Jelle

- Ghent University
- Center for Molecular Modeling / EA17
- FWO PhD fellow
- Experience with local CMM clusters (2014-current), HPC UGent Tier2 clusters (2014-current), Tier1 machines muk (2015-2016) and BrENIAC (2016-current), and Tier0 PRACE infrastructure (Curie thin nodes, 2017). VSC number: vsc40944.

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 envisioned calculations go beyond the state of the art: we are not aware that RPA calculations have ever been applied to systems as large and sensitive to small energy differences as flexible MOFs. However, the size of the considered systems requires us to go to a computational infrastructure that is not routinely accessible. Especially the memory requirements are staggering, as was found using a previous BrENIAC pilot project. For the heaviest calculation that is envisioned, the large-pore structure of MIL-53(Al) at numerically converged settings (see Table 1, task 3), we have thus far only succeeded in performing an RPA calculation on 90 BrENIAC nodes, reaching about 8.5 GiB per core or 237 GiB memory usage per node. Fewer nodes lead to excessive memory requirements. The proposed workload for the challenging MIL-53(Al) system cannot feasibly run on high-memory Tier-2 level infrastructure (e.g. cluster phanpy in Ghent, Cerebro in Leuven) available within the VSC. Simply not enough nodes are available in these setups. Indeed, using fewer nodes increases the memory requirements even further (see point 9). Only BrENIAC can offer the right mix of a considerable number of readily available nodes with sufficient memory per node.

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.

A typical total energy calculation at the RPA level consists of three contributions: a part exact exchange (EXX), a part RPA correlation energy (RPA) and a contribution from single excitations (SE). However, before the RPA correlation energy can be obtained, one also needs to determine the exactly diagonalized wavefunction at the DFT level, from which the RPA calculation can be started in a non-self-consistent manner.

The computational tasks are therefore:

(i) exact diagonalization at the DFT level – task 1
(ii) RPA correlation calculation – task 2 or 3
(iii) SE calculation – task 4
(iv) EXX calculation – task 5.

Based on a previous pilot project on BrENIAC (lt_pilot_42), an estimate of the computational requirements is listed below in Table 1. The RPA calculation is typically the most intensive part, since the necessary RAM memory is substantial. There is however a large difference between calculations on structures which are more akin to the narrow-pore (np) phase of MIL-53(Al) and the large-pore (lp) kind. Lp structures are much heavier, because the large unit cell volume generates much larger numbers of basis functions, and hence strongly increases the memory usage of the calculation. We found that any calculation on a lp structure using fewer than 90 nodes (of 256 GiB RAM each) failed due to excessive memory requirements, and even the 90-node jobs are close to the 256 GiB limit. In contrast, np structures have a smaller volume, and we found calculations to complete successfully already from 30 nodes onwards. This distinction between lp and np structures is also made for the RPA tasks in Table 1 (tasks 2 and 3). The differences between np and lp MIL-53(Al) are not as striking for the other tasks, so an lp/np average value is displayed in Table 1 for the preceding exact diagonalization and the determination of SE and EXX contributions (tasks 1, 4 and 5).

The goal of this Tier1 proposal is to establish the RPA energy-versus-volume behaviour of MIL-53(Al) using two different series of input structures. Indeed, it is currently not feasible to also perform geometry optimizations at the RPA level. We have already optimized structures along the energy profile of MIL-53(Al) using two commonly used DFT methods, PBE+D3(BJ) and PBE+MBD/FI. There are 15 structures per profile, yielding in total 30 structures for which we will calculate the energy at the RPA level. 8 structures per profile need to be treated as a large-pore crystal (16 x task 3), while 7 can be considered narrow-pore-like (14 x task 2).

In terms of data management, the major load is caused by the WAVECAR file, which is used for restarting purposes. It is typically 51 GiB for narrow-pore-like structures and 137 GiB for large-pore ones. To minimize the storage required during the project, we will make sure that no more than a quarter of the 30 structures is on the SCRATCH partition at a time (by deleting the WAVECAR files after manual assertion of the quality of the output).

Table 1 – Justification of the requested computational resources

	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
Task1: Exact diagonalization	2 x 15	0.17	10	51	280	42	MPI	0.014 GiB / 25 files per task = 0.42 GiB / 750 files	51 GiB / 32 files per task = 1530 GiB / 960 files
Task2: RPA narrow pore	2 x 7	0.11	30	46	840	234	MPI	0.005 GiB / 21 files per task = 0.07 GiB / 294 files	51 GiB / 25 files per task = 714 GiB / 350 files
Task3: RPA large pore	2 x 8	0.49	90	706	2520	237	MPI	0.005 GiB / 21 files per task = 0.08 GiB / 336 files	137 GiB / 25 files per task = 2192 GiB / 400 files
Task4: Single excitations (SE)	2 x 15	0.06	10	18	280	37	MPI	0.005 GiB / 21 files per task = 0.15 GiB / 630 files	51 GiB / 21 files per task = 1530 GiB / 630 files
Task5: Exact exchange (EXX)	2 x 15	0.67	10	201	280	73	MPI	0.006 GiB / 23 files per task = 0.18 GiB / 690 files	3 GiB / 28 files per task = 90 GiB / 840 files
SUM				**1022 node days**				**0.9 GiB / 2700 files**	6056 GiB / 3180 files (in practice no more than ¼ of heavy-memory files kept at the same time: **1514 GiB / 3090 files**)

9. Describe the software required to perform the computing task(s). 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, …).

VASP

http://www.vasp.at/
non-exclusive academic licence (see attachment)
own compiled developer's version of VASP with RPA functionality
available on BrENIAC for personal use

In view of the tremendous computational load associated with the VASP RPA calculations, we are not able to perform scaling tests at production run settings. Instead, we have tested the performance of exact diagonalization (task 1), RPA (task 2/3) and EXX (tasks 4/5) for MIL-53(Al) at reduced computational settings (1x1x1 k-grid instead of 3x3x2/3x2x2 for narrow/large pore respectively) during a previous project on BrENIAC (lt_pilot_42). The walltime is reported in Table 2 and Plot 1a, while memory is listed in Table 3 and Plot 1b. In the latter case, the memory reduction factor is defined as the baseline memory usage divided by the memory usage on a higher number of nodes. The memory scaling efficiency is defined as (baseline memory * baseline number of nodes) / (actual memory * actual number of nodes).

By using reduced settings, the quality of the walltime scaling is underestimated: simulations at higher numerical settings (denser k-grids) show a much better scaling than suggested by Table 2, as VASP typically parallelizes very efficiently over k-points. Indeed, we have performed extensive benchmark studies for a wide range of different systems (see attachment), which all showed a close similarity in scaling behaviour. The results in attachment show that a 60-80% efficiency is feasible up to 10 nodes for common calculations like DFT and exact exchange ones, independent of the considered system. In view of memory requirements (exact diagonalization), walltime reduction (exact exchange) and internal consistency (single excitations), we will therefore perform all non-RPA calculations – tasks 1, 4 and 5 – on 10 nodes.

Table 3 shows that the memory load per node for the RPA calculations (and to a lesser extent the exact diagonalizations) decreases significantly as the number of nodes increases. This demonstrates that it is possible to successfully perform production run RPA calculations if sufficiently high numbers of nodes are used. At production run settings, we manually tested as from what number of nodes the memory usage dropped below the available 256 GiB RAM on BrENIAC. We found that we need at least 30 nodes for narrow-pore MIL-53(Al) and 90 nodes for large-pore MIL-53(Al), both yielding a memory usage per node close to 256 GiB. Smaller numbers of nodes lead to excessive memory requirements and causes the RPA calculations to fail. Further scaling could moreover not be tested due to the already heavy demands at this baseline level.

Table 2 – Walltime scaling per computational task

		Exact diagonalization (similar to Task 1)			RPA (similar to Tasks 2 and 3)			Exact exchange (similar to Tasks 4 and 5)		
# nodes	# cores	Wall clock time (s)	Speed-up (w.r.t. baseline)	Efficiency	Wall clock time (s)	Speed-up (w.r.t. baseline)	Efficiency	Wall clock time (s)	Speed-up (w.r.t. baseline)	Efficiency
1	28	3888.4	1,00	100 %	insufficient memory (> 256 GiB / node)			342.5	1,00	100 %
2	56	2227.8	1,75	87 %	2826.8	1,00	100 %	184.6	1,86	93 %
5	140	1079.7	3,60	72 %	1451.8	1,95	78 %	174.9	1,96	39 %
10	280	778.7	4,99	50 %	962.3	2,94	59 %	150.4	2,28	23 %
30	840	908.5	4,28	14 %	793.4	3,56	24 %	286.3	1,20	4 %

Table 3 – Memory scaling per computational task

		Exact diagonalization (similar to Task 1)			RPA (similar to Tasks 2 and 3)			Exact exchange (similar to Tasks 4 and 5)		
# nodes	# cores	Memory load (GiB/node)	Reduction factor (w.r.t. baseline)	Efficiency	Memory load (GiB/node)	Reduction factor (w.r.t. baseline)	Efficiency	Memory load (GiB/node)	Reduction factor (w.r.t. baseline)	Efficiency
1	28	56.6	1,00	100 %	insufficient memory (> 256 GiB / node)			13.2	1,00	100 %
2	56	32.9	1,72	86 %	198.6	1,00	100 %	21.4	0,62	31 %
5	140	18.2	3,11	62 %	84.8	2,34	94 %	23.6	0,56	11 %
10	280	14.1	4,02	40 %	48.3	4,11	82 %	24.2	0,55	5 %
30	840	19.8	2,85	10 %	40.4	4,92	33 %	26.3	0,50	2 %

Plot 1 – Walltime and memory scaling per computational task

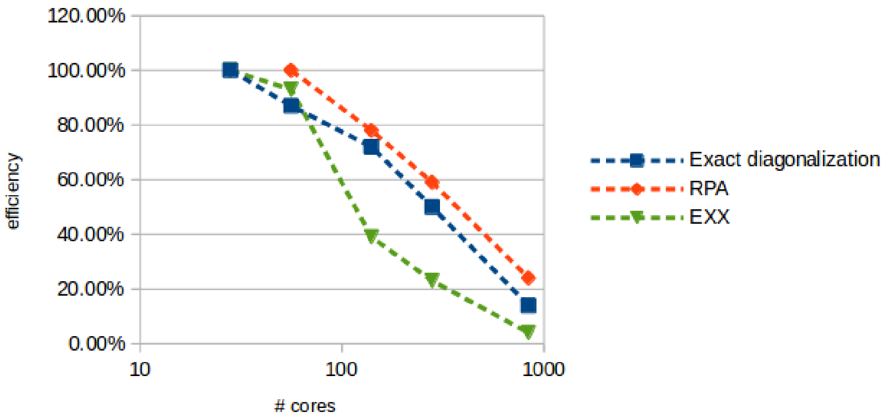

(a) Walltime scaling

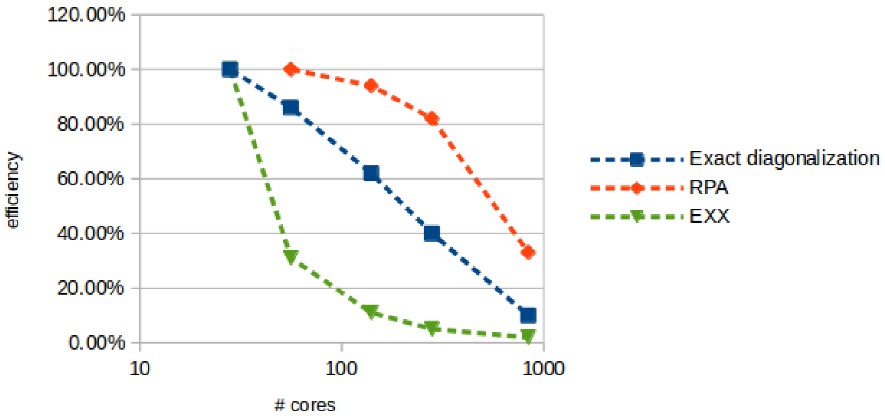

(b) Memory scaling

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 intention of this project is to maintain a constant flow of jobs. Only a limited number of jobs will be submitted at a time (1-5), to be able to keep a close eye on the jobs that require the highest numbers of nodes. As mentioned in point 8, only a quarter of all output will be kept on SCRATCH at most, limiting storage to 1.5 TiB. As soon as the calculations have been checked and found to have finished correctly, the heaviest output files will be removed manually. The remaining files are limited in both number and size and can be copied from Tier-1 very easily.

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.

Research projects

Kurt Lejaeghere

High-throughput screenen van ternaire wolfraamlegeringen met DFT (Muk, 4053 node days).

High-throughput screening of ternary tungsten alloys with DFT (Muk, 4000 node days).

Assessing the accuracy of a screened hybrid functional for property predictions of elemental solids (Muk, 1742 node days).

Error quantification for the adiabatic connection fluctuation-dissipation theorem in the random phase approximation (ACFDT-RPA) (BrENIAC, 1000 node days).

Assessing the accuracy of an efficient meta-GGA functional for property predictions of elemental solids (BrENIAC, 640 node days).

Elucidating the interplay between structure and catalytic activity in nanoporous materials (Tier0 preparatory access; Marconi Broadwell 50 000 core hours, Marconi KNL 100 000 core hours, Curie 50 000 core hours).

Jelle Wieme

Investigating the phases of MIL-53-type materials (Muk, 4644 node days)

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

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

Assessing the accuracy of hybrid functionals for the relative stability of a flexible MOF (BrENIAC, 3420 node days)

Understanding the high-pressure behavior of a flexible nanoporous material (BrENIAC, 2430 node days)

Benchmarking elastic properties of a metal-organic framework (BrENIAC, 1296 node days)

Elucidating the interplay between structure and catalytic activity in nanoporous materials (Tier0 preparatory access; Marconi Broadwell 50 000 core hours, Marconi KNL 100 000 core hours, Curie 50 000 core hours).

Scientific output in the last two years

http://dx.doi.org/10.1103/PhysRevB.94.235418
http://dx.doi.org/10.1126/science.aad3000
http://dx.doi.org/10.1021/acs.chemmater.6b05444
http://dx.doi.org/10.1016/j.commatsci.2016.01.039
http://dx.doi.org/10.1039/C7TA01559C
http://dx.doi.org/10.1039/C7DT02752D
http://dx.doi.org/10.1038/s41467-017-02666-y

Moreover, results from the previous TIER1 work have been in the PhD thesis of Kurt Lejaeghere (ISBN 978-90-8578-690-0) and in posters and talks at conferences (ICAMM 2016 (Rennes, September 5-7 2016); EUROMAT 2017 (Thessaloniki, September 17-22 2017); …). Several additional papers are in preparation or have already been submitted.

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

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

fwo
Opening
new
horizons

De heer Kurt Lejaeghere
Beernemsteenweg 42
8750 WINGENE

uw kenmerk	contact	datum
	wt@fwo.be	22 juni 2016
ons kenmerk	02 550 15 88	
66991 12O0117N		

betreft **Toekenning postdoctoraal onderzoeker 2016-2019**

Geachte heer,

De raad van bestuur heeft u vandaag aangesteld als postdoctoraal onderzoeker van het FWO met ingang van 1 oktober 2016 tot en met 30 september 2019. Het reglement vindt u als bijlage.

Graag willen wij u feliciteren met deze aanstelling.

Voor de verdere administratieve opvolging zal het FWO u nog vóór 15 juli 2016 een arbeidsovereenkomst bezorgen.

Onderzoekers 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.

Voor bijkomende inlichtingen kan u via bovenvermeld e-mailadres steeds terecht bij de dossierbeheerder van uw wetenschapsgebied. Vóór 1 oktober 2016 zal u nog een gedetailleerde feedbackbrief ontvangen.

Tenslotte vragen wij u met aandrang om op uw publicaties steeds uw titel "postdoctoraal onderzoeker van het FWO" te vermelden.

Wij wensen u veel succes toe in uw verdere onderzoeksloopbaan.

Met vriendelijke groeten,

Waarnemend secretariaat-generaal,

Danny Huysmans
Directeur Intern Beheer

Hans Willems
Directeur Steun aan Onderzoekers

bijlagen: 1

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.

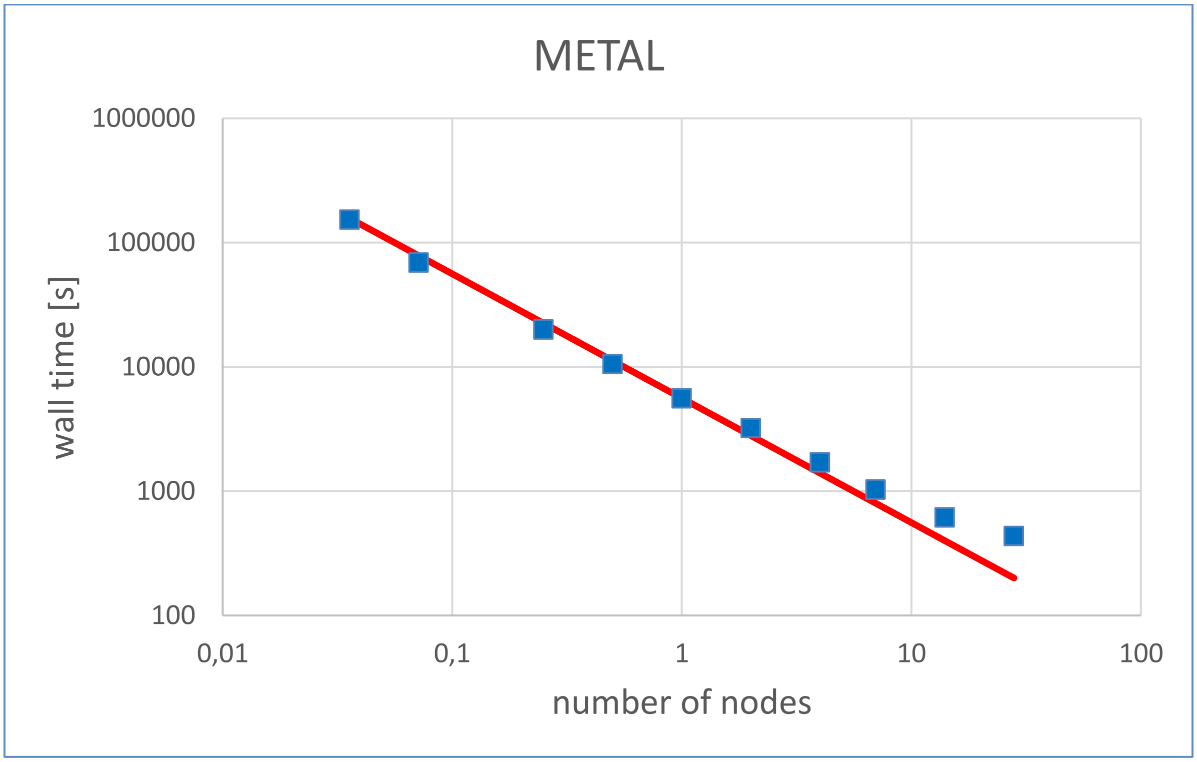

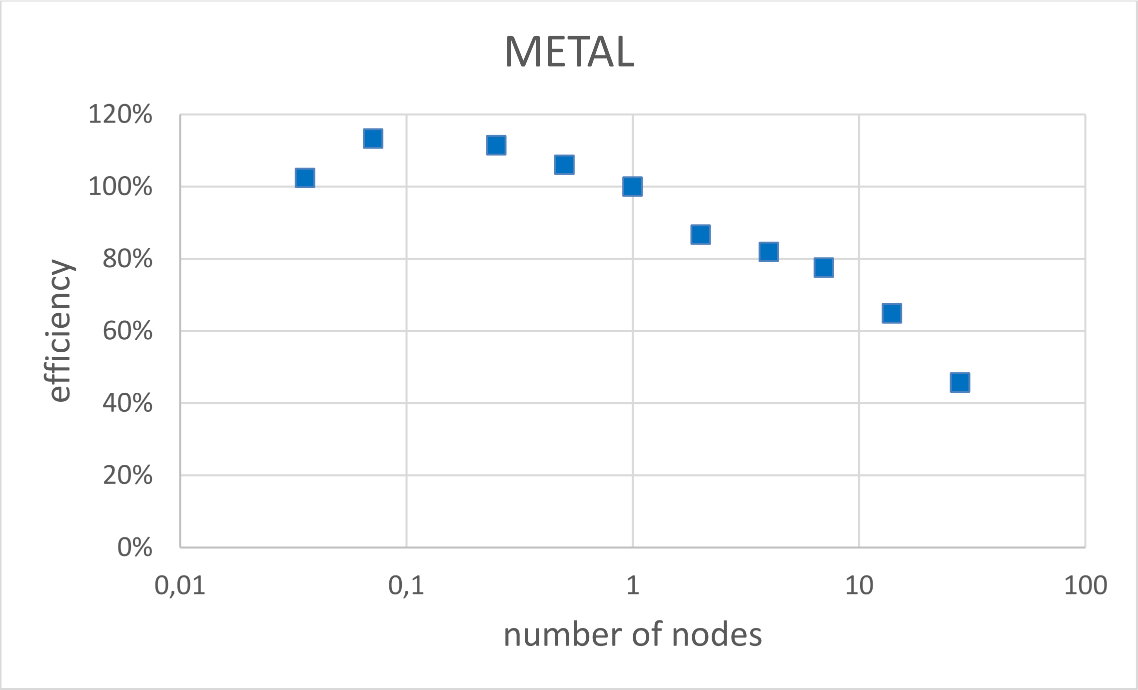

Figure 2: Intra- and multinode scaling of the wall time for the SEMI system (BrENIAC 2016 and Muk 2013). The red and orange lines denote the ideal scaling behaviour.

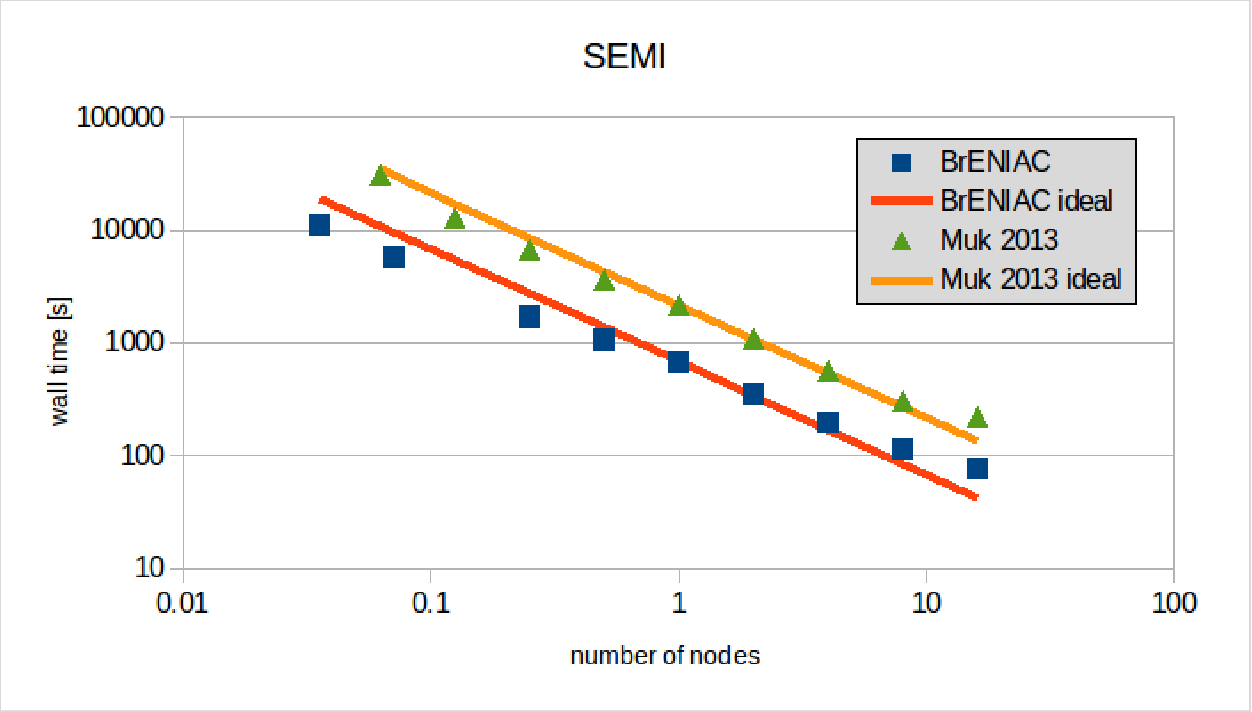

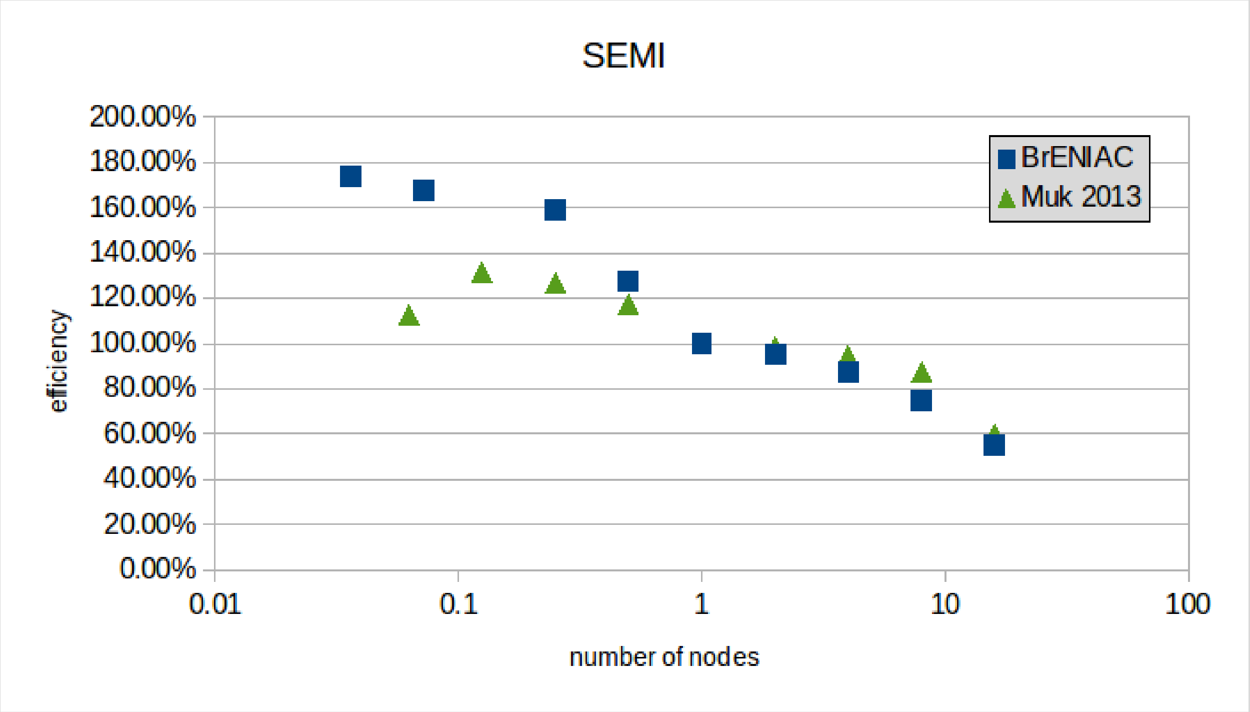

Figure 3: Intra- and multinode scaling of the wall time for the PORE system. The red line denotes the ideal scaling behaviour.

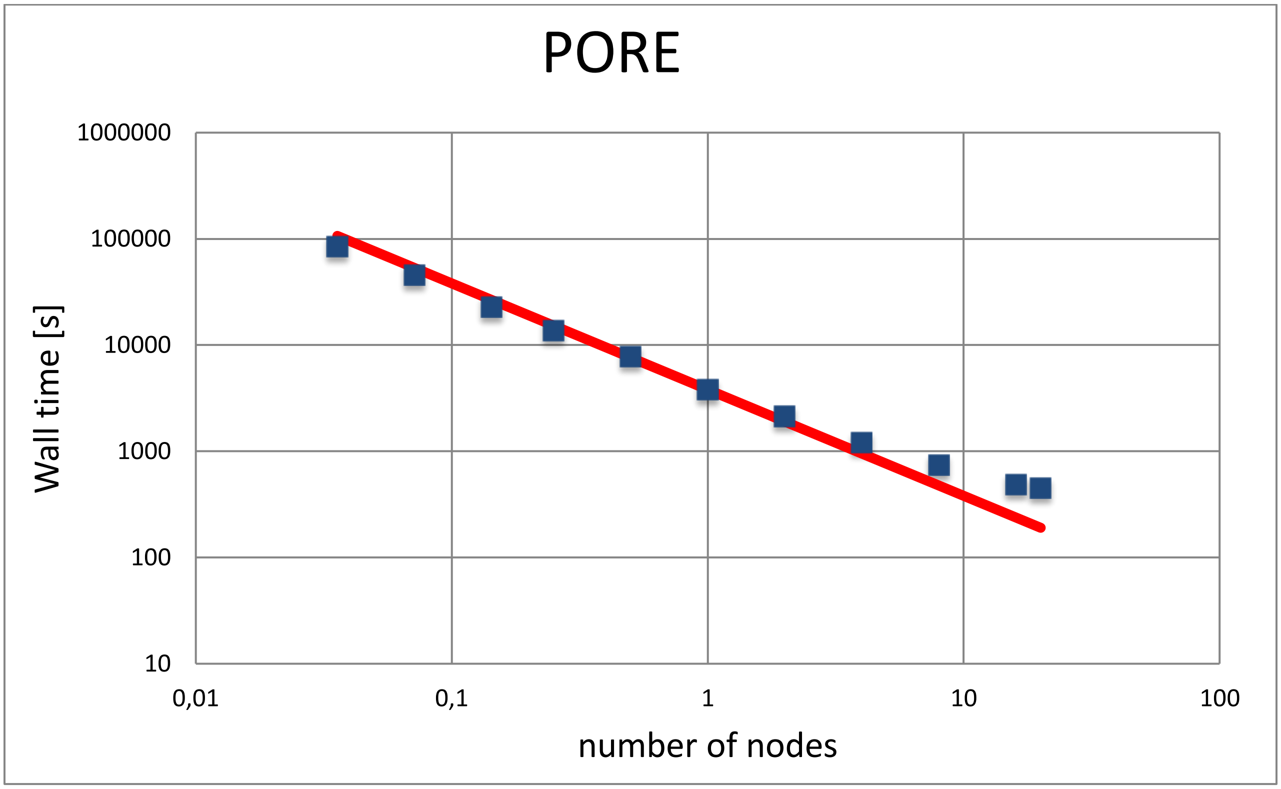

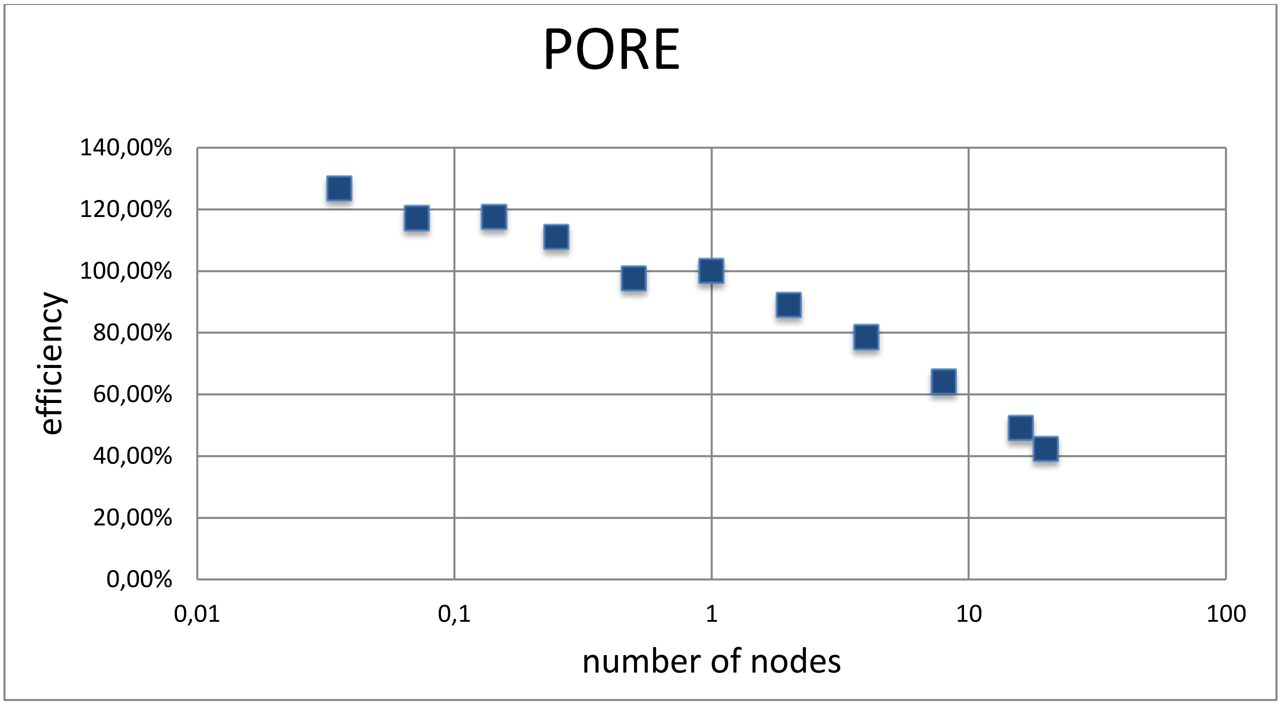

By performing the calculations on a few cores up to multiple nodes, we note that the computational efficiency proceeds in 2 steps. On the one hand, it remains most efficient to perform calculations on 1 or a few cores, and up to the use of an entire node, the efficiency steadily declines. This intranode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