Enclosure 1b. Category 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://vscentrum.be/.

Title of the application:

High-throughput screening of ternary tungsten alloys with DFT

Name and first name of the applicant:

Lejaeghere Kurt

Institution:

Ghent University

Research group / department:

EA17 / Center for Molecular Modeling

Title / position:

FWO fellow (PhD student)

e-mail address:

Kurt.Lejaeghere@UGent.be

Total computing time that is needed, in node days:

4000

Total disk storage that is applied for:

182 GiB

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

Ab initio screening of suitable tungsten alloys as first wall material in nuclear fusion reactors

2. Short description of the research project within the framework of which computing time is applied for (max. 1 A4 in Arial 12):

Future fusion reactors critically depend on the selection of appropriate materials for the so-called 'first wall'. These metals come into direct contact with the plasma and are hence exposed to high temperatures and neutron fluxes. The fusion community agrees that tungsten addresses these issues best, but its room-temperature brittleness still needs to be resolved, as this compromises the malleability and long-term stability of fusion reactor components. Alloying may offer a solution.

The number of possible tungsten alloys is nearly unlimited, however, and only a computational approach enables an assessment of a sufficiently large number of candidate compounds within a reasonable time. The FWO fellowship of Kurt Lejaeghere aims at performing such a computational screening in a high-throughput fashion. Whereas this fellowship previously allowed the benchmark of the elemental phases (DOI: 10.1080/10408436.2013.772503 and 10.1103/PhysRevB.89.014304) and the calibration of new methodologies using the tungsten binaries (DOI: 10.1103/PhysRevLett.111.075501), it is now the aim to investigate ternary tungsten alloys, which are more relevant for actual fusion research. To do this, the previously developed high-throughput methods will be used.

3. Financing institution or channel, financing the research project in full or in part (FWO, BOF, IWT, EU, ...): Please attach the confirmation letter as enclosure (see instructions in enclosure 3).

FWO

4. Promoter of the research project:

Prof. dr. ir. Veronique Van Speybroeck

5. Persons mandated by the Applicant to compute on the TIER1 within the framework of the present project: Please provide for every person:

- name and first name
- institution
- research group / department
- title / position
- experience with TIER1/TIER2 infrastructure in Belgium and abroad

Lejaeghere Kurt
Ghent University
EA17 / Center for Molecular Modeling
FWO fellow
experience with local CMM clusters, HPC UGent TIER2 clusters and the TIER1 machine muk (vsc40323)

Van Speybroeck Veronique
Ghent University
EA17 / Center for Molecular Modeling
Full professor
experience with local CMM clusters and HPC UGent TIER2 clusters (vsc40021)

Cottenier Stefaan
Ghent University
EA10 / WE05 / Center for Molecular Modeling
Assistant professor
experience with local CMM clusters and with Leuven and HPC UGent TIER2 clusters (vsc40026)

6. Description of the computing task, justification for the computing time, disk storage and memory that are applied for, and description of the software tools required (max. 3 A4 in Arial 12). Please clearly provide the following in this regard:

- the number of nodes/cores that are applied for per computing task, with a subdivision of the computing time in sub-tasks indicating the sequence of the sub-tasks
- whether these tasks use diversification (OpenMP, MPI, hybrid OpenMP/MPI …)
- the estimated memory use of a computing task (maximum 64GiB/node)
- whether a vSMP system will be used

- the requirements for disk storage (estimated volume in GiB and the total number of files), more specifically for:
  - required input files (data set, parameter files, etc.)
  - SCRATCH volume used during the performing of the computing tasks
  - result files

We aim to simulate 182 ternary tungsten alloys. In a previous TIER1 pilot phase project (4053 node days), we were able to collect enough data to assess ternary tungsten alloys with 4d or 5d transition metal dopants. For fusion purposes, this is only half of the necessary information, as 3d transition metal dopants are plausible too. The purpose is therefore now to extend the search to all tungsten-based transition metal alloys. This will be done by means of a **SCOOP**-based Python script (http://code.google.com/p/scoop/) that uses **ASE** (https://wiki.fysik.dtu.dk/ase/) to distribute the numerous parallel **VASP** calculations (http://www.vasp.at/). All necessary software has already been installed on muk. We added a VASP benchmark in attachment.

From the pool of possible tungsten alloys, the script selects a limited number of compounds to treat concurrently. Each of these individual ternaries is characterized in a few distinct steps: first the most stable structure (out of 6 possibilities) is identified by means of 5-point equations of state; then that most stable structure is characterized in more detail (7-point equation of state, elastic constants, band structure, density of states). An overview of the workflow is represented schematically below.

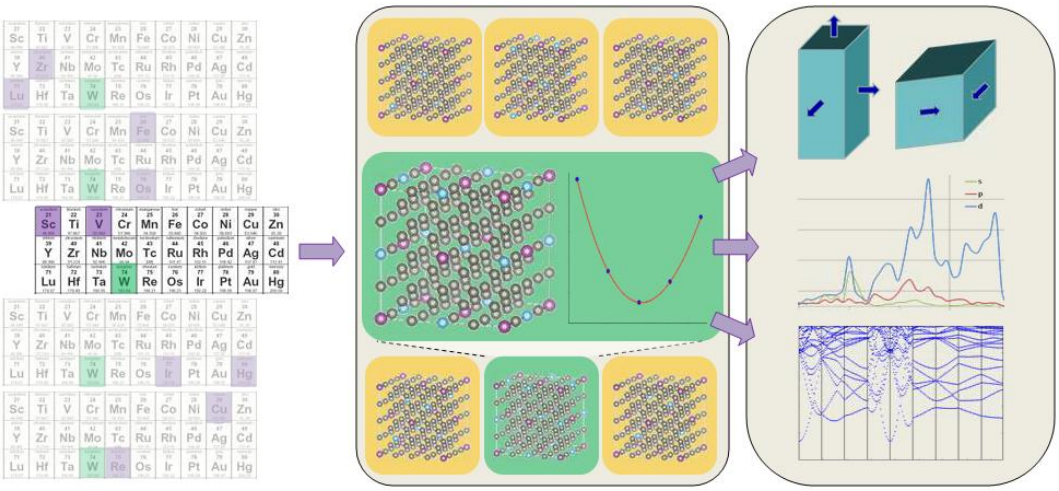

This approach is therefore truly of a high-throughput nature: the requirements for a single calculations are not that intensive, but the challenge lies in the huge number of separate simulations that are needed. The 72-hour wall time is not an issue, as the script contains a restart clause to check which calculations were already performed previously.

If we limit ourselves to transition metal dopants, there are 465 or $\binom{30}{2}$ possible ternary tungsten alloys. 283 of these have already been obtained in the framework of a previous TIER1 proposal and TIER2 calculations. For the remaining set, we need to take into account a computational time of 21 node days per compound:

Total work package = 182 computational taks		
1 task	*Computing time*	*Memory use*
Single relaxation	200 core hours	25 MiB
Selection most stable structure (out of 6 possibilites via a 5-point equation of state)	6 x 5 x relaxation = 6000 core hours	6 x 5 x relaxation = 750 MiB
Determination 7-point equation of state	7 x relaxation = 1400 core hours	7 x relaxation = 175 MiB
Determination bands, elastic constants and density of states	4 x relaxation = 800 core hours	4 x relaxation = 100 MiB
Subtotal	8200 core hours = **21 node days**	**1025 MiB**
Total	182 tasks x 21 node days / task = **3887 node days**	182 tasks x 1025 MiB / task = **182 GiB**

Together, this yields 3887 node days. Note that this proposal assumes a longer computing time and larger SCRATCH memory use for a single relaxation than the previous TIER1 pilot project. This is due to two reasons: (1) the previously allotted time of 4053 node days did not suffice to calculate the 283 ternaries that were already obtained; part of them were calculated on TIER2 infrastructure (2) most of the structures that will be calculated here will be spin-polarized; this roughly doubles the computing time and required memory. To also deal with possible

eventualities, we therefore ask for **4000 node days**. The SCRATCH memory use of **182 GiB** is to be considered as an upper limit. We will write as much information back to the TIER2 as possible after each 72-hour run.

Each distinct VASP calculation will be parallellized over one entire node, using all 16 cores. Virtual memory use will seldom amount to more than **32 GiB RAM per node**. Because each calculation is limited to one node, and because the parallellization between nodes is managed by SCOOP, **no MPI** is required. **No vSMP** scheme will be used either.

7. Please indicate why the TIER1 is the appropriate machine to perform the computing task (max. 1/2 A4 in Arial 12):

Although each individual VASP calculation can easily be performed on TIER2 infrastructure, it does not allow us to process that large numbers of jobs. The limited number of available nodes and the long queue times for large jobs make a similar study of 182 ternaries on TIER2 infrastructure impossible within an acceptable time. Even when using 20 16-core nodes non-stop, 200 days of wall-clock time would be required. Only the Tier1 allows us to complete such a computational task within a reasonable time.

8. Summary of the software required to perform the computing task, and possible installation and compilation instructions (max. 2 A4 in Arial 12). Please clearly provide the following per item in this regard:

- a reference to the software's web page
- the software licence system (open source, GPL, etc.)
- if there is no free academic use of the software, state which licence makes the installation and the use valid on the TIER1 by the Applicant (+ add a copy of the signed licence)
- if need be, which licence server will be used (name + IP address)
- whether the software is already available on the TIER1 (see https://vscentrum.be/nl/tier1-rekenen) and, if this is not the case, compilation and installation instructions (possibly with reference to existing TIER2 installation)

VASP

http://www.vasp.at/
non-exclusive academic licence (see attachment)
available on TIER1

ASE

https://wiki.fysik.dtu.dk/ase/
GNU Lesser General Public Licence v2.1+
available on TIER1

SCOOP

http://code.google.com/p/scoop/
GNU Lesser General Public Licence
available on TIER1

9. Period during which the task is to be performed:

As soon as possible. It would be convenient to possess all data by October 1, to guarantee the publication of the results before the FWO deadline of January 31.

10. Describe the results that were obtained within the framework of computing time that was attributed during the past two years on the TIER1 or on other TIER1 or TIER0 supercomputers (max. 2 A4 in Arial 12):

Analysis of the results for 4d and 5d transition metal dopants in tungsten (4053 node days in a pilot phase TIER1 project) showed the favourable effect of an average of 5 d-electrons. Materials that approximate this average, such as $W_{30}PdTa$, combine a better ductility with a low thermal expansion and a good high-temperature resistance at a reasonable cost. 3d dopants are expected to behave differently, due to the smaller atomic volumes and the occurrence of magnetism. For the design of future fusion materials, it is therefore crucial to investigate these materials as well.

Results from this previous TIER1 work have not been published in a journal yet, but have appeared in the PhD thesis of Kurt Lejaeghere (ISBN 978-90-8578-690-0). A publication is planned when the calculations have

been completed for all transition metal ternaries (i.e. combining the results of the previous TIER1 project with those of the present proposal).

Should you have any questions or encounter any difficulties during the electronic submission of an Application, please contact by e-mail:
Associatie KU Leuven: leen.vanrentergem@kuleuven.be ; jan.ooghe@kuleuven.be
Associatie Universiteit Gent: hpc@ugent.be
Associatie Universiteit Hogescholen Antwerpen: stefan.becuwe@uantwerpen.be
Associatie Universiteit Hogescholen Limburg: geertjan.bex@uhasselt.be
Universitaire Associatie Brussel: rosette.vandenbroucke@vub.ac.be ; rosette.vandenbroucke@gmail.com
For the other institutions: dane.skow@herculesstichting.be

Fonds Wetenschappelijk Onderzoek
Research Foundation – Flanders

De heer Kurt Lejaeghere
Beernemsteenweg 42
8750 WINGENE

uw kenmerk	ons kenmerk	contact	datum
	05_05 1.1.375.11	bio@fwo.be cult@fwo.be wt@fwo.be gm@fwo.be med@fwo.be interdisciplinair@fwo.be	23 juni 2010

Betreft: **Toekenning Aspirant 2010-2012**

Geachte heer,

De Raad van Bestuur heeft u vandaag aangesteld als Aspirant met ingang van 1 oktober 2010 tot en met 30 september 2012. Het reglement vindt u als bijlage.

Graag wil ik u feliciteren met deze aanstelling als Aspirant van het FWO.

In het kader van de verdere administratieve opvolging zal het FWO u nog vóór 15 juli 2010 een overeenkomst houdende toekenning van de beurs bezorgen.

Voor bijkomende inlichtingen kan u steeds terecht bij uw dossierbeheerder op bovenvermeld e-mailadres.

Tenslotte vraag ik u met aandrang om op uw publicaties steeds uw titel Aspirant van het FWO te vermelden.

Ik wens u veel succes toe in uw verdere onderzoeksloopbaan.

Hoogachtend,

dr.ir. Elisabeth Monard
secretaris-generaal

Bijlage: 1

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

Egmontstraat 5
B-1000 Brussel

Tel. +32-(0)2-512 91 10
Fax. +32-(0)2-512 58 90
E-mail post@fwo.be
Internet www.fwo.be

Stichting van Openbaar Nut
Ondernemingsnummer 0880.212.840

Ik aanvaard mijn beurs als Aspirant van het Fonds Wetenschappelijk Onderzoek – Vlaanderen vanaf 01/10/2012 tot en met 30/09/2014 met het oog op het doorvoeren van het wetenschappelijk onderzoek, vermeld in mijn mandaataanvraag, aan de Universiteit Gent.Studiecentrum voor Kernenergie

Ik ben ervan op de hoogte dat ik administratief en juridisch onder de bevoegdheid kom van het Fonds Wetenschappelijk Onderzoek - Vlaanderen, vertegenwoordigd door zijn voorzitter en zijn secretaris-generaal, en dat ik disciplinair onder de bevoegdheid kom van de academische overheid van de universiteit en eventueel de wetenschappelijke instelling waaraan ik mijn werkzaamheden doorvoer.

Ik neem er nota van dat:

a) mijn beurs wordt geschorst gedurende langdurig ziekteverlof, ouderschapsverlof, zwangerschaps- en borstvoedingsverlof, palliatief verlof, medische bijstand.

b) het me niet is toegestaan, behoudens formele machtiging, mijn beurs bij het FWO te cumuleren met welke andere toelage of bezoldiging ook;

c) ik de verbintenis aanga een doctoraat op proefschrift voor te bereiden en dat ik mij voor geen enkele andere studie mag inschrijven, noch de eraan verbonden examens afleggen, zonder vooraf de instemming van het FWO te hebben bekomen.

d)
- het bedrag van mijn beurs beperkt is tot het nettoloon dat ik zou ontvangen als assistent aan een universiteit in de Vlaamse Gemeenschap (salarisschalen van toepassing op het assisterend personeel van de universiteiten in de Vlaamse Gemeenschap volgens het Besluit van de Vlaamse regering van 4 mei 2001 (BS 9 januari 2002);
- het netto beursbedrag per maand minimum 1.823,49 euro bedraagt (index 1,5769 op 1 maart2012);
- de beurs wordt aangepast aan de toegekende geldelijke anciënniteit op basis van artikel 27 van het reglement inzake mandaat Aspirant van het FWO;
- de beurs rekening houdt met de gezinstoestand en onderworpen is aan de sociale zekerheid;
- het vakantiegeld en de jaartoelage overeenstemt met het vakantiegeld en de eindejaarstoelage dat ik zou ontvangen als assistent aan een universiteit in de Vlaamse Gemeenschap.

Ik verbind me er toe strikt het reglement inzake het mandaat Aspirant van het FWO, waarvan ik verklaar een exemplaar ontvangen te hebben, en de voorschriften van mijn academische overheid na te leven.
Op al mijn publicaties en hun overdrukken, zal ik mijn hoedanigheid van Aspirant van het Fonds Wetenschappelijk Onderzoek - Vlaanderen vermelden.

Datum en handtekening voorafgegaan door de woorden: 'gelezen en goedgekeurd'
gelezen en goedgekeurd, 23/07/2012

Kurt Lejaeghere

PS: De keerzijde van dit document eveneens aanvullen.

VASP on Muk: performance

The new VSC Tier1 machine, Muk, consists of:

- 528 computing nodes, with two 8-core Intel E5-2670 (Sandy Bridge) processors each (2.60 GHz clock frequency)
- 64 GB RAM and a 500 GB local 7.2k RPM disk per node
- 400 TB local scratch (DDN SFA10k with 580 of 1 TB 7.2k RPM (data) and 20 of 15k RPM disks (metadata)), accessible through a GPFS-shared file system
- a Mellanox Connect-X FDR Infiniband netword
- 4 login nodes, each identical to a computing node

Test problem: Ge (100) nanowire with 320 electronic bands and 8 irreducible k-points in a 3-step geometry optimization (LDA, 345 eV cut-off)

1 Optimal settings for VASP

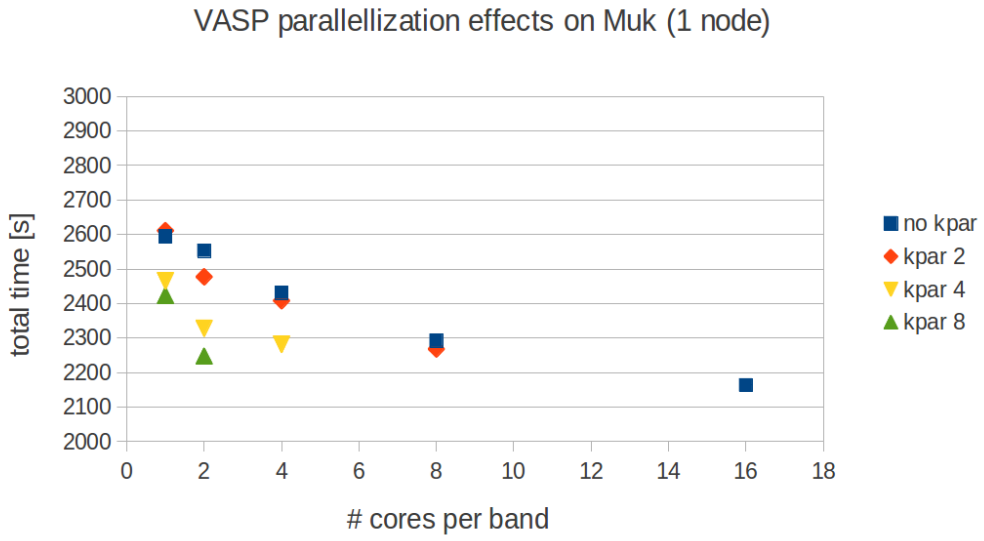

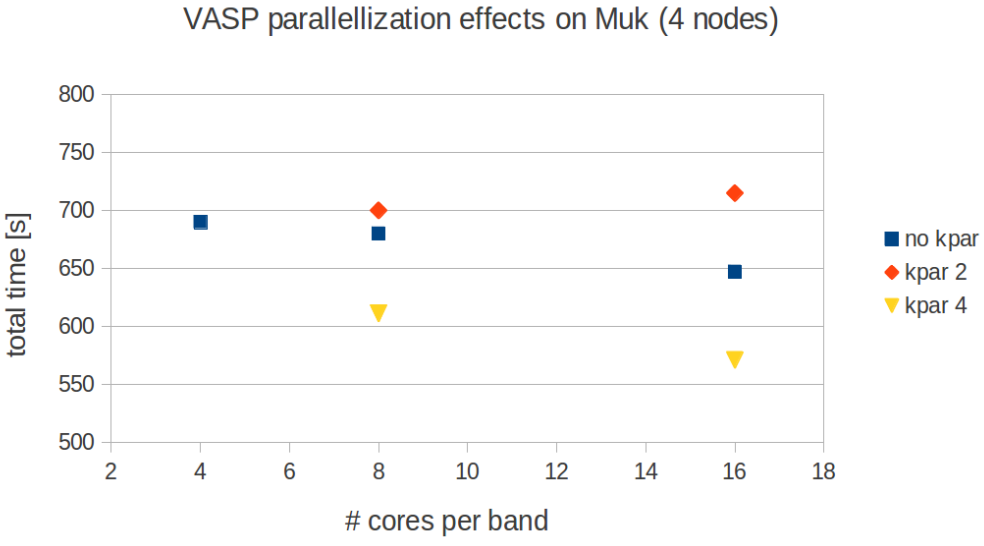

The total duration of a computational run decreases as the number of cores per band (within a node) increases. The optimal setting for NPAR is hence equal to the number of nodes. K-point parallellization is very efficient, but for multinode calculations, it is best to use it only when more than 2 k-points are computed simultaneously. The optimal setting for KPAR is hence equal to the number of k-points, or as large as possible (if this number is larger than 2).

2 Muk scaling

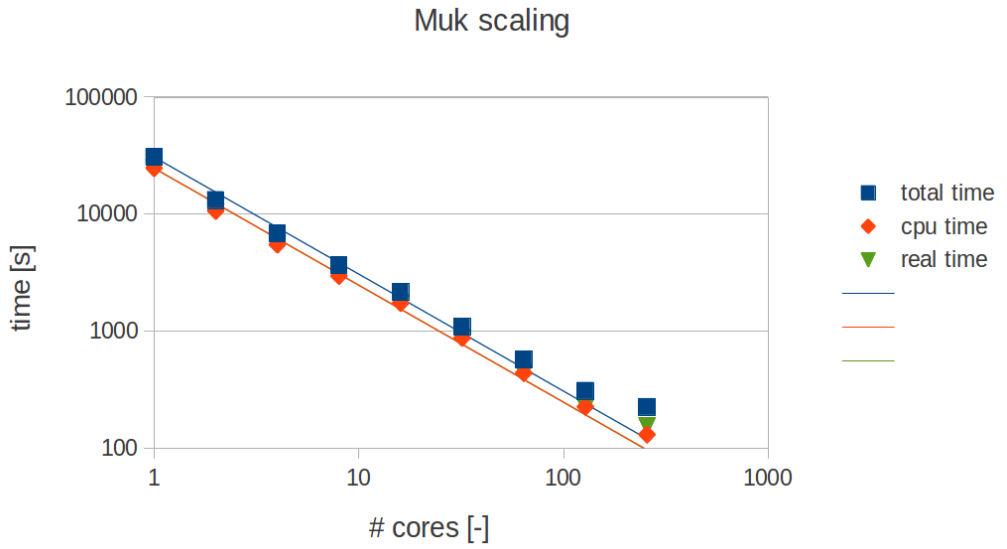

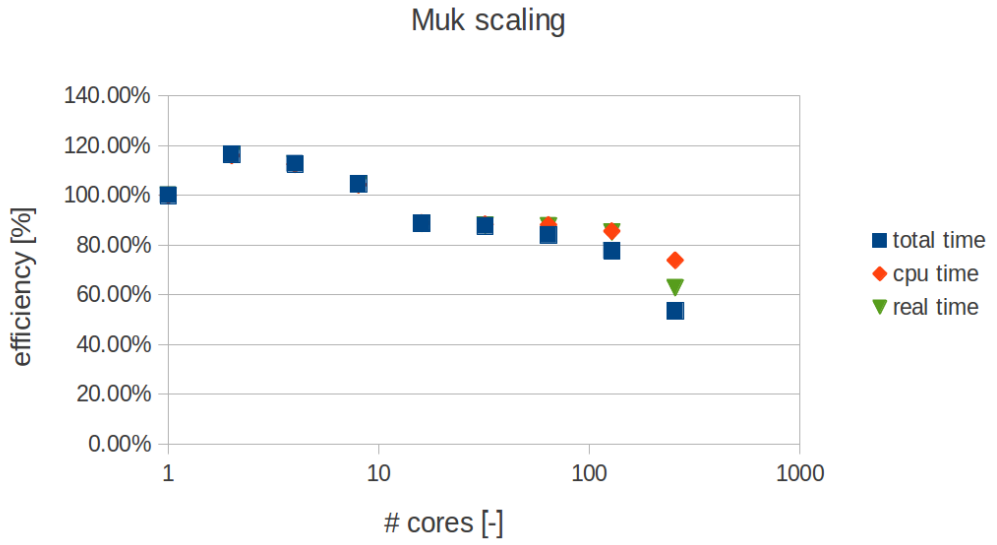

Muk displays an almost ideal scaling behaviour up to 100 cores. For even more cores, the cpu time keeps scaling reasonably well, but other issues, such as data transfer, consume increasingly more time.

3 Limitations

Preliminary tests with VASP have been performed using cells with up to 1728 atoms (CuZr supercells). 1000-atom structures can still be computed on 512 cores (32 full nodes), but for the 1728-atom cells,

it is already necessary to revert to hybrid-mode calculations (partial use of the node), in particular 512 cores or 128 partially filled nodes.

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)