Vlaanderen
is supercomputing

VLAAMS **SUPERCOMPUTER CEN**TRUM

USERS DAY 2019

TUESDAY JUNE 4, 2019

vscentrum.be

With support from
fwo

Program

09:50	Welcome
10:00	Lecture dr. Joost Vandevondele (CSCS, ETH Zürich) Developing scientific software for modern high performance computing
11:00	Coffee break
11:20	Information sessions part 1 – VSC staff ▶ Accelerate your research using GPUs ▶ VSC Cloud and data platform ▶ Meet the VSC staff: panel discussion
12:30	Lunch
13:30	Information sessions part 2 – VSC staff ▶ VSC Cloud and data platform ▶ How to get access to PRACE infrastructure ▶ C++ shared memory programming: threading building blocks as an alternative to OpenMP
14:40	One minute poster presentations
15:00	Coffee break + Poster session
16:00	Lecture Prof. Francesco Contino (VUB) Robust Optimisation: challenges to find optimum under uncertainties.
16:30	Poster award and closing
16:40	Reception
18:00	End

Posters

1. The ESA Vitrual Space Weather Modeling Centre. Stefaan Poedts (KU Leuven)
2. Calculating Vibrational spectra in solids: From fingerprinting defects in diamond to the impact of water in molecular crystals. Danny E.P. Vanpoucke, Goedele Roos, Ken Haenen (IMO, IMOMEC, CNRS)
3. Numerical investigation of a turbulent spray jet. Alessandro D'Ausilio, Ivana Stankovic, Bart Merci (Ghent University)
4. First-principles mechanistic study of alkylphenols dealkylation catalyzed by acidic zeolites, Massimo Bocus, Julianna Hajek, Yuhe Liao, Bert F. Sels, Veronique Van Speybroeck (Ghent University, KU Leuven)
5. DF Cooling in Acetonitrile and Dichloromethane. Xiaoyong Zhang, Jeremy N Harvey (KU Leuven)
6. Towards deep active learning using reliable uncertainty estimates. M. Larmuseau, M. Sluydts, T. Dhaene, and S. Cottenier (Ghent University, ePotentia)
7. An Asymptotic-Preserving Multilevel Monte Carlo Method for Particle Based Simulation of Kinetic Equations. Emil Loevbak, Giovanni Samaey, Stefan Vandewalle
8. LES/CMC simulation of the coria rouen spray flame using the Lagrangian point particle approach. Alessandro D'Ausilio, Ivana Stankovic, Bart Merci (Ghent University)
9. Predicting Axial Compressor Stall under Uncertainties by means of Multilevel Monte Carlo. Pieterjan Robbe, Dirk Nuyens, Stefan Vandewalle (KU Leuven)
10. Relaxing Scalability Limits with Speculative Parallelism in Sequential Monte Carlo. Nemeth, Balazs, Haber, Tom, Liesenborgs, Jori, Lamotte, Wim (Hasselt University)
11. Quantifying the Dynamics of Evolutionary Rates Through Time. Membrebe, J.V.; Baele, G.; Suchard, M. ; Lemey, P. (KU Leuven, UCLA)
12. Deep learning on single-cell ATAC-seq data to decipher enhancer logic, Ibrahim Ihsan Taskiran, Liesbeth Minnoye, Zeynep Kalender Atak, Gert Hulselmans, Valerie Christiaens, and Stein Aerts (VIB KU Leuven, KU Leuven)
13. Chemical Compound Space exploration for redox flow batteries. J.L. Teunissen, M. Denayer, F. De Vleeschouwer, F. De Proft (VUB)
14. Single-Layer Structures of a100- and b010-Gallenene: A Tight-Binding Approach. M. Nakhaee, M. Yagmurcukardes, S. A. Ketabi, F. M. Peeters (UAntwerpen)
15. Numerical Characterization of the Flow Field in a Water Model of Liquid Metal-Cooled Nuclear Reactor G. Alessi, S. Buckingham, P. Planquart, L. Koloszar, S. Lopes, A. Villa-Ortiz, C. Spaccapanicca, K. Van Tichelen (SCK-CEN)
16. Toward an Atomistic Understanding of Deep Eutectics Solvents Electrochemical Interfacial Structure. Mesfin Haile Mamme, Samuel L.C. Moors, El Amine Mernissi Cherigui, Monika Lukaczynska, Herman Terryn, Johan Deconinck, Jon Ustarroza, and Frank De Proft (VUB)
17. Urban climate modelling in Sub-Saharan Africa. Brousse Oscar, Van de Walle Jonas, Wouters Hendrik, Demuzere Matthias, Thiery Wim, van Lipzig Nicole PM (KU Leuven)

Posters

18. “Eppur si muove! And yet it moves!”:Study of the intermolecular dynamic behaviour of multidentate ferrocenyl phosphines and their corresponding d10 metal complexes in solution. César A. Urbina Blanco, Andrés R. León Garzón, Benjamin Kovács, José Martins, Jean-Cyrille Hierso, Mark Saeys, Titus van Erp and Sander Roet (Ghent University)
19. Computational fluid dynamics-based optimization of steam cracking furnaces using novel technologies. Yu Zhang, Stijn Vangaever, Jens N. Dedeyne, Guy B. Marin, Geraldine J. Heynderickx, Kevin M. Van Geem (Ghent University)
20. PARAMOUR project: Influence of small-scale atmospheric processes on the decadal variability and predictability of the atmosphere-ocean-ice sheet system in the polar regions. Vanden Broucke S. , Sotiriadis S., Van Lipzig N. (KU Leuven)
21. HLA-sorter: a Long Amplicon based Clustering Method to identify HLA Alleles. Koen Herten, Vicky Van Sandt, Aleksander Senev, Greet Peeters, Johan Kerkhofs, Marie-Paule Emonds, Joris R Vermeesch (KU Leuven, Rode Kruis Vlaanderen)
22. Computational fluid dynamics-based studies of multiphase reactive flows. Noel Gomez, Laurien Vandewalle, Sepehr Madanikashani, Shekhar Kulkarni, Geraldine J. Heynderickx, Guy B. Marin, and Kevin M. Van Geem (Ghent University)
23. Coordinated Control of Wind Farms based on Large-Eddy Simulations, Pieter Bauweraerts, Ishaan Sood & Johan Meyers (KU Leuven)
24. New CDFT-based computational approach to elucidate tandem mass spectrometry fragmentation patterns of molecules. Vanhaegenborgh, Y.,Geerling, P.,De Proft, F.,Cauët, E. (VUB)
25. Investigating the photophysical properties of octaphyrins. Tatiana Woller, Paul Geerlings, Frank De Proft, Benoît Champagne, Mercedes Alonso (VUB, Université de Namur)
26. Fast stochastic turbulence generator for aeroacoustic applications. Ali Kadar, Wim Desmet (KU Leuven)
27. Instabilities in magnetized, two-component jets: relativistic 3D simulations. Dimitrios Millas, Charalampos Sinnis, Nektarios Vlahakis (KU Leuven, National and Kapodistrian University of Athens)
28. A systematic DFT study of self-assembled alkane thiols on a gold (111) surface. Anne-Julie Truyens, Jelle Vekeman, and Frederik Tielens (VUB)
29. Convection-permitting climate modeling over the Lake Victoria region. Jonas Van de Walle, Wim Thiery, Oscar Brousse, Niels Souverijns, Matthias Demuzere, Nicole P.M. van Lipzig (KU Leuven, VUB, U Bochum)
30. Computational study of Nickel-catalysed Negishi Arylation of Propargylic bromides. Andrea Darù, Jeremy N. Harvey (KU Leuven)
31. The New Era of the Digital Herbarium. Huybrechts, P., De Smedt, S., Bogaerts, A., Stoffelen, P. Engledow, H., Groom, Q.J. (Meise Botanical Garden)
32. A Nudged Elastic Band Method for Lennard-Jones Clusters. Cansu Utku, Jeremy Harvey (KU Leuven)

Abstracts VSC USERS DAY 2019

Developing scientific software for modern high performance computing (Joost Vandevondele, CSCS)

High-end supercomputers are an essential tool for scientific progress, and our needs will continue to grow. However, with technology reaching the end of Moore's law scaling, architectural diversity and complexity will likely continue to increase. For example, accelerated architectures, using GPUs from various vendors, will dominate the list of high-performance systems installed in the US and world-wide. Furthermore, programming models need to take into account the fact that data movement, not compute, is increasingly the bottleneck, and our choice of algorithms should reflect that reality. Good software engineering will be needed to develop software that scientists can use to rapidly prototype new methods and models, yet is able to exploit the performance of modern hardware. Increasingly, one should separate the concerns of what should be computed and how it should be computed, differentiating between mathematical descriptions, algorithmic descriptions, imperative code and explicit data structures, and low level optimization. Refactoring existing codes to adopt the tools of the big players in HPC (e.g. python, C++) and making code maintainable and reusable are essential to make good use of the resources available for scientific computing.

To illustrate these concepts, examples will be given from a variety of domains. A first example comes from the weather and climate domain, where CSCS and MeteoSwiss pioneered the use of GPUs for operational weather forecast, and developed a set of libraries (GridTools) to build weather models on top of domain specific languages (DSLs). The second example will be in computational chemistry/material science, where CP2K and its sparse matrix multiplication library DBCSR, allow for efficient DFT computations of large systems on GPUs. In the latter context, the potential of machine learning to accelerate both the models as well as the implementation will be demonstrated.

Robust Optimisation: challenges to find optimum under uncertainties (Francesco Contino, VUB)

In a design phase, engineers often rely on simulations and optimise a set of parameters towards one or several objectives. This procedure becomes expensive as they tackle the many requirements (flexibility, efficiency, resilience, ...) of novel concepts and try to reach the required level of fidelity. Yet, more often than not, they rely on deterministic simulations, i.e. not considering the uncertainties of the conditions of these simulations. In some instances, the quality of an optimum can be dramatically affected by a slight change in the working conditions (wear of the geometry, change of boundary, material properties, ...). Taking these uncertainties into account leads to robust design optimisation but adds a significant computational overhead, also known as the curse of dimensionality.

Based on the research performed in the BURN group, and within the SBO project EUFORIA, we have been able to couple the best optimisation strategy with efficient uncertainty quantification methods. This leads to a significant decrease of the computational work and simulations that would not be affordable now become realistic. This presentation includes the challenges that still remain and where the VSC is a key instrument to achieving the next step in the industry.

PRACE Project Access and Peer Review Process: How to prepare a good PRACE proposal (Luca Marsella en Maria Grazia Giuffreda van CSCS (vanuit Zwitserland)

The overarching goal of PRACE is to provide a federated European supercomputing infrastructure that is science driven and globally competitive. It builds on the strengths of European science providing high-end computing and data analysis resources. To provision a federated world-class Tier-0 supercomputing infrastructure PRACE has put in place a strong transparent peer review process exclusively based on scientific excellence, which will be presented together with the Project Access.

The PRACE Peer Review of proposals submitted for a PRACE Project Access on a Tier-0 system includes two types of assessment: scientific and technical. The two assessments are carried out in a coordinated effort by two groups of experts, as the technical review is used by the scientific one, which ultimately determines whether a proposal is successful or not. The presentation will discuss the key elements of a successful proposal, considering both the scientific and the technical aspects. The focus will be set on the latter, illustrating the most recent template for submitting Tier-0 applications and highlighting how to present the technical information that is necessary in a successful proposal.

VSC Cloud and Data platform

End 2017 the Flemish Government decided to invest 30 million euro to create a new Tier-1 Supercomputing platform based on the Flemish Supercomputer Centre (VSC) 2018-2022 program “Supercomputing as a Service”. This projects aims to keep investing in compute power but also to meet the growing demands towards data and flexibility in science, offering a platform based on three integrated components: Compute, Data and Cloud.

In this session we present the progress of the project since the implementation started a year ago. This comprises the extension of BrENIAC for the Compute component, a detailed overview of the current Cloud infrastructure and its usage, as well as a demonstration of Globus Online to transfer (large) data files to and from the Tier-1 and Tier-2 HPC systems.

Accelerate your research with GPUs

The use of GPUs in high performance computing is becoming a mature practice due to significant performance advantages for some applications/algorithms. Such performance gains stem from available hardware support in modern data centers, in addition to the software stack which can offload some part of work on GPUs. This can result in roughly 5x to 100x times faster code, without compromising precision, depending on how the GPUs are exploited during the application runtime. Currently, VSC offers few tens of GPU nodes across different machines, and encourages users to make efficient use of such devices to leverage performant computing in science and industry.
During this presentation, you will learn how you can level up to using (multiple) GPUs in your application. The examples and scaling tests will address various fields of current research (AI, chemistry, etc), and can be adapted to your workflow after some relevant adjustments.

C++ shared memory programming: threading building blocks as an alternative to OpenMP

For the previous two decades, OpenMP has been the de facto starndard for shared memory programming in the context of scientific computing. All major C, C++ and Fortrancompilers offer support for OpenMP out of the box. It is a directive-based extension to each of those languages, implying that parallelization can be applied gradually without requiring a major rewrite of the serial code base. However, as the complexity of the underlying hardware increases, that promise becomes harder and harder to keep, and redesign is often required to achieve significant efficiencylevels. Since the OpenMP standard is mostly feature identical for C, C++ and Fortran, modern C++ features and OpenMP are not alwaysa good match. Lastly, it is quite hard to handle OpenMP parallelization at vairous levels of the application, i.e., combining OpenMPconstructs in your code with a library parallelised with OpenMP is non-trivial.

Threading Building Blocks (TBB), originally developed by Intel and now open source aims to mitigate the latter two problems. It is a template library that works well with many STL components and modern C++ features, and it is task-based, with an efficient runtime scheduler allowing for efficient parallelization at various levels.

In this information session, we will cover a number of TBB features so that you have an overview of what TBB can do for your C++ code base.

Meet the VSC staff

This session is an opportunity to discussand exchange ideas with VSC staff members. A few questions to guide the discussion:

- You get access to infrastructure that is an order of magnitude more capable, how would that change your research?
- How can the VSC facilitate better interaction between the researchers and other users?

USERS DAY 2019

This event is supported by