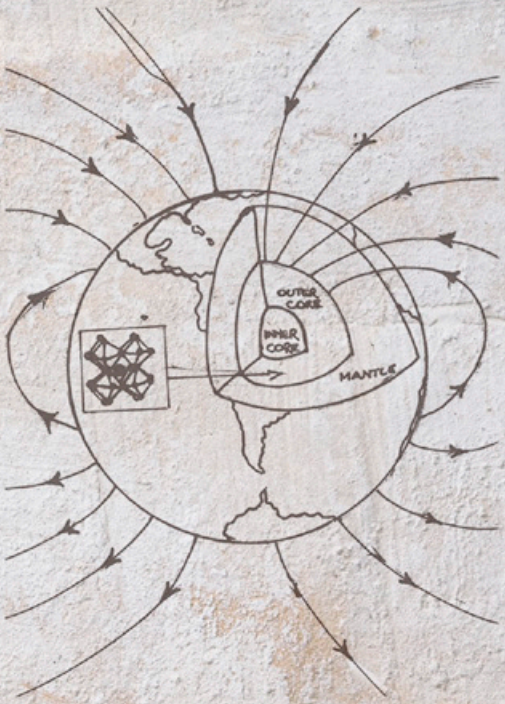

EULER EQUATION

$$\frac{\partial \rho}{\partial t} + \sum_{i=1}^{3} \frac{\partial(\rho u_i)}{\partial x_i} = 0$$

$$\frac{\partial(\rho u_j)}{\partial t} + \sum_{i=1}^{3} \frac{\partial(\rho u_i u_j)}{\partial x_i} + \frac{\partial p}{\partial x_i} = 0$$

$$\frac{\partial E}{\partial t} + \sum_{i=1}^{3} \frac{\partial((E+p)u_i)}{\partial x_i} = 0$$

i, j label the three Cartesian components.
$(x_1, x_2, x_3) = (x, y, z)$ and
$(u_1, u_2, u_3) = (u, v, w)$

PASC18

Platform for Advanced Scientific Computing Conference

Basel Switzerland

2-4 July 2018

```
quicksort(A, i, k):
  if i<k:
    p:=partition(A,i,k)
    quicksort(A,i,p-1)
    quicksort(A,p+1,k)
```

POISSON'S EQUATION

$$\Delta \varphi = f$$

Δ = LAPLACE OPERATOR
f & φ REAL OR COMPLEX-VALUED FUNCTIONS

$$\nabla^2 \varphi = f$$

IN THREE-DIMENSIONAL CARTESIAN COORDINATES:

$$\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}\right)\varphi(x,y,z) = f(x,y,z)$$

When
$f = 0$ we retrieve LAPLACE'S EQUATION

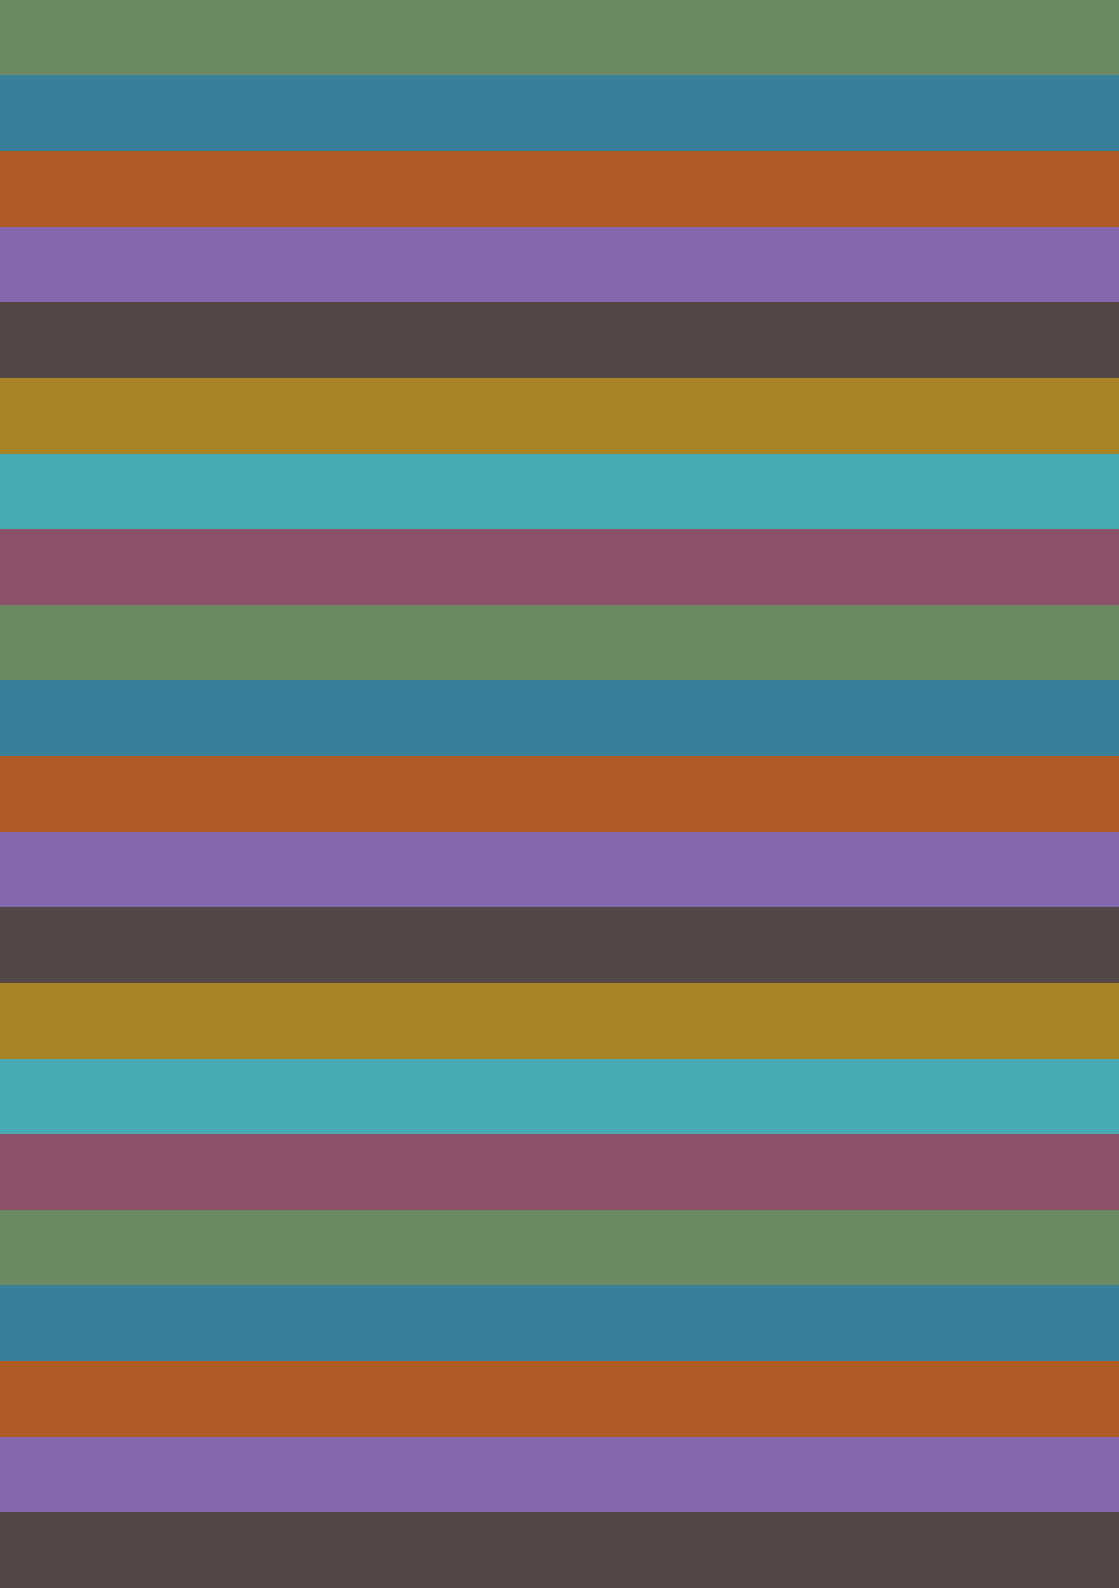


PASC18
Platform for Advanced Scientific Computing
Conference
Basel
Switzerland
2-4 July 2018

Welcome to PASC18

Chair:
Florina Ciorba
(University of Basel, Switzerland)

Chair:
Erik Lindahl
(Stockholm University, Sweden)
Photo: Markus Majetic

We are delighted to welcome you to PASC18 at the Congress Center Basel, Switzerland. The city of Basel is situated on the river Rhine at the intersection of three countries – Switzerland, France and Germany. Home to the oldest university in Switzerland, Basel is considered the cultural capital of the country and Europe's leading research centre for life sciences, medical research, energy engineering and cultural sciences.

PASC18 is the fifth edition of the PASC Conference series, an international platform for the exchange of competences in scientific computing and computational science, with a strong focus on methods, tools, algorithms, application challenges, and novel techniques and usage of high performance computing.

The theme of PASC18 is "Fast and Big Data, Fast and Big Computation", emphasizing the close coupling of data and computation in current and future high-performance computing applications. A panel discussion bringing perspectives from various scientific domains and industry is dedicated to this theme.

The PASC Conference is first and foremost a platform for promoting interdisciplinary communication. At PASC18 we introduce a new session, the interdisciplinary dialogue, where the audience, coming from diverse research fields, can gain insight into a specific field through an interview between prominent computational scientists from different research domains. In this year's dialogue, Petros Koumoutsakos (ETH Zurich) will interview Constantia Alexandrou (University of Cyprus) about her field of research – quantum chromodynamics.

Other program highlights include keynotes from David Bader (Georgia Tech) on massive-scale analytics in real world problems, Marina Becoulet (CEA) on first-principles modelling of magnetohydrodynamics in fusion devices, Alice-Agnes Gabriel (Ludwig Maximilian University of Munich) on extreme-scale earthquake simulations, and Nils P. Wedi (ECMWF) on kilometer-scale weather and climate simulations. Minisymposium, paper and poster presentations complete the technical program, with more than 250 contributions in total from the eight scientific domains represented at the conference.

PASC18 is co-sponsored by the Association for Computing Machinery (ACM) and the PASC Structuring Project, supported by the Council of Federal Institutes of Technology (ETH Board). The PASC Conference is coordinated by the Swiss National Supercomputing Centre (CSCS).

We are grateful to our local hosts – the University of Basel and the City of Basel – and to all participants for contributing to a strong and vibrant program. We thank the following companies and organizations for their support: HPE, IBM, PSI, CRAY, DDN, MARVEL, MICROSOFT, NOVARTIS and NVIDIA.

Conference Chairs

Florina Ciorba (University of Basel, Switzerland)
Erik Lindahl (Stockholm University, Sweden)

Scientific Committee

Minisymposia & Posters Program Chairs

Florina Ciorba (University of Basel, Switzerland)
Erik Lindahl (Stockholm University, Sweden)
Sabine Roller (University of Siegen, Germany)
Jack Wells (Oak Ridge National Laboratory, USA)

Papers Program Chairs

Sabine Roller (University of Siegen, Germany)
Jack Wells (Oak Ridge National Laboratory, USA)

Proceedings Chair

Timothy Robinson (ETH Zurich / CSCS, Switzerland)

Chemistry & Materials
Domain Co-Chair: **Edoardo Di Napoli** (Forschungszentrum Jülich, Germany)
Domain Co-Chair: **Aurora Clark** (Washington State University, USA)
Wibe de Jong (Lawrence Berkeley National Laboratory, USA)
Matteo Giantomassi (Université Catholique de Louvain, Belgium)
Mathieu Luisier (ETH Zurich, Switzerland)
Markus Meuwly (University of Basel, Switzerland)
Kristin Persson (Lawrence Berkeley National Laboratory, USA)
Giovanni Pizzi (EPFL, Switzerland)

Climate & Weather
Domain Co-Chair: **Willem Deconinck** (ECMWF, UK)
Domain Co-Chair: **Rupert Ford** (Science and Technology Facilities Council, UK)
Mario Acosta (Barcelona Supercomputing Center, Spain)
Luca Bonaventura (Politecnico di Milano, Italy)
Katherine Evans (Oak Ridge National Laboratory, USA)
Oliver Fuhrer (MeteoSwiss, Switzerland)

Computer Science & Applied Mathematics
Domain Co-Chair: **Michael Heroux** (Sandia National Laboratories, USA)
Domain Co-Chair: **Martin Schulz** (TU Munich, Germany)
Simone Deparis (EPFL, Switzerland)
Laura Grigori (INRIA, France)
Matthias Müller (RWTH Aachen University, Germany)
Richard Vuduc (Georgia Institute of Technology, USA)
Ulrike Yang (Lawrence Livermore National Laboratory, USA)

Emerging Application Domains
Domain Co-Chair: **Simon Scheidegger** (University of Zurich, Switzerland)
Domain Co-Chair: **Georgia Tourassi** (Oak Ridge National Laboratory, USA)
Rumi Chunara (New York University, USA)
Philipp Eisenhauer (University of Bonn, Germany)
Roger Käppeli (ETH Zurich, Switzerland)
Shannon Quinn (University of Georgia, USA)
Philipp Renner (Lancaster University, UK)
Kerstin Kleese van Dam (Brookhaven National Laboratory, USA)

Engineering
Domain Co-Chair: **Richard Sandberg** (The University of Melbourne, Australia)
Domain Co-Chair: **Jackie Chen** (Sandia National Laboratories, USA)
George Biros (The University of Texas at Arlington, USA)
Steve Plimpton (Sandia National Laboratories, USA)
Philipp Schlatter (KTH Royal Institute of Technology, Sweden)
Maarten van Reeuwijk (Imperial College London, UK)

Life Sciences
Domain Co-Chair: **Dan Jacobson** (Oak Ridge National Laboratory, USA)
Domain Co-Chair: **Abigail Morrison** (Forschungszentrum Jülich, Germany)
Ben Brown (Lawrence Berkeley National Laboratory, USA)
Sharlee Climer (University of Missouri - St. Louis, USA)
Georgios Gkoutos (University of Birmingham, UK)
Susanne Kunkel (KTH Royal Institute of Technology, Sweden)
Sandipan Mohanty (Forschungszentrum Jülich, Germany)

Physics
Domain Co-Chair: **George Lake** (University of Zurich, Switzerland)
Domain Co-Chair: **Sinéad Ryan** (Trinity College, Ireland)
Frank Jenko (University of California, USA)
Tilo Wettig (University of Regensburg, Germany)
Frank Wuerthwein (UC San Diego, USA)

Solid Earth Dynamics
Domain Co-Chair: **Ebru Bozdag** (Colorado School of Mines, USA)
Domain Co-Chair: **Dimitri Komatitsch** (CNRS, France)
Sebastien Chevrot (CNRS, France)
David May (University of Oxford, UK)
Louise Kellogg (UC Davis, USA)
Rene-Edouard Plessix (Shell Technology Center Amsterdam, Netherlands)
James Wookey (University of Bristol, UK)

Organizing Committee

Conference Coordinator
Tatjana Ruefli (ETH Zurich / CSCS, Switzerland)

Assistant Conference Coordinator
Raluca Hodoroaba (ETH Zurich / CSCS, Switzerland)

Technical Program Coordinator
Timothy Robinson (ETH Zurich / CSCS, Switzerland)

Communication Officers
Angela Schneebeli, **Valentina Tamburello** (ETH Zurich / CSCS, Switzerland)

Treasurer
Michele De Lorenzi (ETH Zurich / CSCS, Switzerland)

Sponsorship Program
Vittoria Rezzonico (EPFL, Switzerland)

Submission System Administrators and Webmasters
Monica Frisoni, **Alessandro Prato**, **Stefano Schuppli** (ETH Zurich / CSCS, Switzerland)

IP
Invited Plenary Presentations

IP01 Unraveling Earthquake Dynamics Through Extreme-Scale Multi-Physics Simulations 2
Alice-Agnes Gabriel (Ludwig Maximilian University of Munich, Germany)

IP02 Public Lecture on Massive-Scale Analytics Applied to Real-World Problems 3
David Bader (Georgia Institute of Technology, USA)

IP03 From Weather Dwarfs to Kilometre-Scale Earth System Simulations 4
Nils P. Wedi (ECMWF, UK)

IP04 Challenges in the First Principles Modelling of Magneto Hydro Dynamic Instabilities and their Control in Magnetic Fusion Devices 5
Marina Becoulet (CEA, France)

ID
Interdisciplinary Dialogue

ID01 The Colourful Theory, and Visible and Invisible Matter in the Universe: An Interdisciplinary Dialogue between Constantia Alexandrou and Petros Koumoutsakos 8
Constantia Alexandrou (University of Cyprus, Cyprus), Petros Koumoutsakos (ETH Zurich, Switzerland)

PNL
Panel Discussion

PNL01 Panel Discussion on Big Data vs. Fast Computation – Is HPC Facing a Game Change? 12

AP
ACM PASC18 Papers

AP01 ACM PASC18 Papers Session I 16
AP02 ACM PASC18 Papers Session II 18
AP03 ACM PASC18 Papers Session III 20
AP04 ACM PASC18 Papers Session IV 22

MS
Minisymposia

MS01 Adaptive Parallel Strategies for the Exploration of Challenging Search Spaces with Applications in Particle Simulations and Optimization, Part I 26
MS02 Capability Computing, Performance Portability, and Co-Design in the PASC Projects 29
MS03 Computational Aspects of Heterogeneous Agents Macro 32
MS04 Distributed Training of Deep Neural Net Models for High Energy Physics 35
MS05 Foundations and Applications of Performance Engineering 38
MS06 Large Scale Electronic-Structure Calculations on Modern and Future High-Performance Supercomputers 41
MS07 Machine Learning in Weather and Climate 44
MS08 On the Road to Exascale Computing: Turbulence Simulations of Complex Flows at the PetaFlops Pit Stop, Part I: Applications 47
MS09 Adaptive Parallel Strategies for the Exploration of Challenging Search Spaces with Applications in Particle Simulations and Optimization, Part II 50
MS10 Bridging the Software Productivity Gap for Climate and Weather Models 53
MS11 Computing the Effect of Risk 56
MS12 Engineering Scientific Software in times of Agile Development, Continuous Integration and Cloud Computing 59
MS13 Generative Models and Density Estimator for High Energy Physics 62
MS14 How Fintech and Big Data Change and Challenge the Insurance Sector 65
MS15 Machine Learning and Quantum Chemistry 68
MS16 NP-Hard Computations: Massively Parallelizing Mixed-Integer Linear Programs 71
MS17 On the Road to Exascale Computing: Turbulence Simulations of Complex Flows at the Petaflops Pit Stop, Part II: Methods 74
MS18 Addressing Resilience Challenges for Computing at Extreme Scale 77
MS19 Advances in Computational Geosciences, Part I 80
MS20 Challenges in Porting and Maintaining Atmospheric Codes on Emerging Hardware Architectures 83
MS21 Computational Solutions to Large-Scale Data Management and Analysis Challenges in Personalized Health 86
MS22 Fostering Software Engineering Best Practice within Research Teams 89
MS23 High Performance Graph Algorithms 92
MS24 Plasma I: Exciting Opportunities for Plasma Simulation in the Pre-Exascale Era 95
MS25 Scientific Computing in times of MPI+X: Looking at Multiple "X" with regard to Performance and Portability 98
MS26 Tensor Algebra Computation: Implementations and Applications 101
MS27 Actionable Health Intelligence: From Precision Medicine to Population Health 104
MS28 Advances in Automation and Efficiency for the Exascale Era – Experiences from the Biomolecular Sciences 107
MS29 Advances in Computational Geosciences, Part II 110
MS30 Efficient Parallel Methods in High-Dimensional Approximation and Beyond 113
MS31 How Can We Escape the Data Avalanche in Climate Science? 116
MS32 Increasing Credibility of Simulation and Analytic Software for Science 119

MS33 Machine Learning Schemes with High Extrapolation Accuracy for Materials Discovery 122
MS34 Plasma II: Frontiers in Gyrokinetic Turbulence Simulation on New and Emerging HPC Platforms 125
MS35 Gravitational-Wave Data Analysis with the Current Generation of Advanced Detectors 128
MS36 HPC for HEP: Enabling Big Data from Large Instruments on Leadership Class HPC Infrastructures 131
MS37 HPUQ: Current Challenges in Uncertainty Quantification for Mechanistic Models, Part II: Theory, Methods and Tools 134
MS38 Mass and Energy Transport Phenomena in Solid State 137
MS39 Scalable Solvers for Forward and Inverse Problems in Geophysics 140
MS40 Towards Weather and Climate Simulations at 1-km Resolution 143
MS41 Use of AI to Analyze Complex Biological Systems 146
MS42 Coupling Strategies Towards Exascale for Complex Earth System Modelling 149
MS43 Distributed Asynchronous Parallel Computing: Progress and Challenges for Multi-Physics Applications on Heterogeneous Architectures 152
MS44 Emerging Trends in Statistical Mechanics Applications to Nanostructured Materials 155
MS45 Evolution of Knowledge Management in Astrophysics 158
MS46 HPC beyond HEP: Opening Doors for New Data Intensive Sciences at Leadership Class HPCs Using BigPanDA 161
MS47 HPUQ: Current Challenges in Uncertainty Quantification for Mechanistic Models, Part II: Applications in Life Sciences and Engineering 164
MS48 Unconventional Methods for Partial Differential Equations 167

Poster Sessions

CHM Chemistry & Materials 172
CLW Climate & Weather 172
CSM Computer Science & Applied Mathematics 172
EAD Emerging Application Domains 174
ENG Engineering 174
LIF Life Sciences 174
PHY Physics 175
SED Solid Earth Dynamics 175

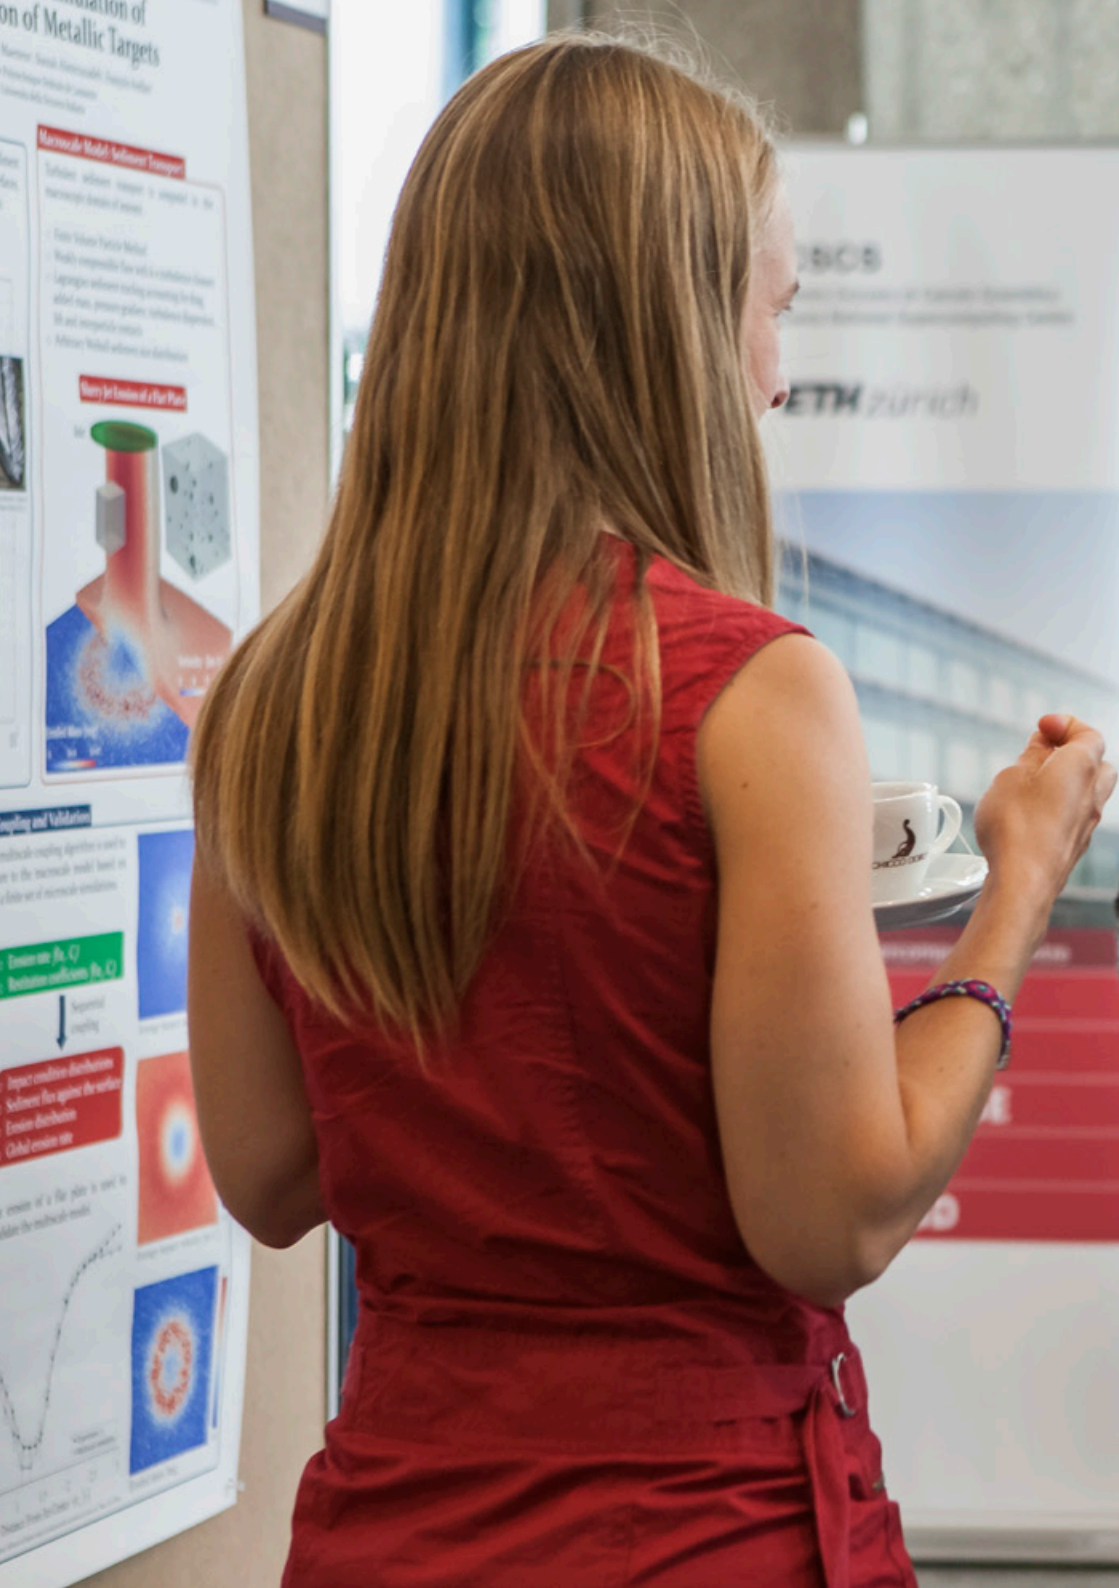
ETH zurich

IP

Invited Plenary Presentations

IP01

Unraveling Earthquake Dynamics Through Extreme-Scale Multi-Physics Simulations

Alice-Agnes Gabriel (Ludwig Maximilian University of Munich, Germany)

Monday,
July 2, 2018

10:20 – 11:10
Montreal Room

Chair:
Dimitri Komatitsch
(CNRS, France)

Earthquakes are highly non-linear multiscale problems, encapsulating geometry and rheology of faults within the Earth's crust torn apart by propagating shear fracture and emanating seismic wave radiation. This talk will focus on using physics-based scenarios, modern numerical methods and hardware specific optimizations to shed light on the dynamics, and severity, of earthquake behavior. It will present the largest-scale dynamic earthquake rupture simulation to date, which models the 2004 Sumatra-Andaman event - an unexpected subduction zone earthquake which generated a rupture of over 1,500 km in length within the ocean floor followed by a series of devastating tsunamis. The core components of the simulation software will be described, highlighting the benefits of strong collaborations between domain and computational scientists. Lastly, future directions in coupling the short-term elastodynamics phenomena to long-term tectonics and tsunami generation will be discussed.

Alice-Agnes Gabriel *is an Assistant Professor of Geophysics at Ludwig Maximilian University of Munich. She received a PhD in seismology from ETH Zurich in 2013. She fuses expertise from Earth science, physics and computational mathematics to study the fundamentals of earthquake physics and develop methodological innovations for seismology. She is specifically interested in simulating waves and rupture processes within arbitrarily complex geological structures to enhance classic probabilistic seismic hazard assessment and a wide range of industry applications. Her career is distinguished by first-rate earthquake scenarios realized on some of the largest supercomputers worldwide.*

IP02

Public Lecture on Massive-Scale Analytics Applied to Real-World Problems

David Bader (Georgia Institute of Technology, USA)

Tuesday,
July 3, 2018

18:30 – 19:30
Montreal Room

Chair:
Bastien Chopard
(University
of Geneva,
Switzerland)

Emerging real-world graph problems include: detecting and preventing disease in human populations; revealing community structure in large social networks; and improving the resilience of the electric power grid. Unlike traditional applications in computational science and engineering, solving these social problems at scale often raises new challenges because of the sparsity and lack of locality in the data, the need for research on scalable algorithms and development of frameworks for solving these real-world problems on high performance computers, and for improved models that capture the noise and bias inherent in the torrential data streams. In this talk, Bader will discuss the opportunities and challenges in massive data-intensive computing for applications in social sciences, physical sciences, and engineering.

David Bader *is Professor and Chair of the School of Computational Science and Engineering at Georgia Institute of Technology, and is regarded as one of the world's leading experts in data sciences. His interests are at the intersection of high performance computing (HPC) and real-world applications, including cybersecurity, massive-scale analytics, and computational genomics. Bader has co-authored over 200 articles in peer-reviewed journals and conferences, and is an associate editor for high-impact publications including IEEE Transactions on Computers, ACM Transactions on Parallel Computing, and ACM Journal of Experimental Algorithmics. He is a Fellow of the IEEE and AAAS, and has served on a number of advisory committees in scientific computing and cyber-infrastructure, including the White House's National Strategic Computing Initiative. Bader has served as a lead scientist in several DARPA programs and is a co-founder of the Graph500 list, a rating of "Big Data" computing platforms. He was recognized as a "Rock Star of HPC" by InsideHPC and as HPCwire's "People to Watch" in 2012 and 2014.*

IP03

From Weather Dwarfs to Kilometre-Scale Earth System Simulations

Nils P. Wedi (ECMWF, UK)

Wednesday,
July 4, 2018

10:00 – 10:50
Montreal Room

Chair:
Willem Deconinck
(ECMWF, UK)

The European Centre for Medium-Range Weather Forecasts (ECMWF) leads a number of Horizon 2020 activities (ESCAPE) with innovation actions for developing a holistic understanding of energy-efficiency for extreme-scale applications using heterogeneous HPC architectures by: (a) defining and encapsulating the fundamental algorithmic building blocks ("Weather and Climate Dwarfs") underlying weather and climate services; (b) combining frontier research on algorithm development with hardware adaptation using DSLs; (c) developing benchmarks and cross-disciplinary Verification, Validation, and Uncertainty Quantification (VVUQ) for weather and climate applications; and (d) synthesizing the complementary skills of global numerical weather prediction with leading European researchers. This talk will illustrate the need for and practicality of producing ensembles of km-scale simulations, summarize progress on accelerating state-of-the-art global weather and climate predictions, and discuss outstanding issues and future directions on producing and analysing big weather data while balancing time-critical customer needs with energy- and time-to-solution.

Nils P. Wedi *has a PhD from Ludwig Maximilian University of Munich and joined ECMWF in 1995. His career at ECMWF encapsulates a diverse range of work both technical and scientific. He leads ECMWF's Earth System Modelling section that addresses all aspects of scientific and computational performance relating to ECMWF's forecast model and the ensemble forecasting system. He develops strategies to secure the scalability of the model on future high-performance computing systems. He is the scientific coordinator of the European H2020 projects ESCAPE and ESCAPE-2 to address the challenges of rising energy cost for computing towards affordable, exascale high performance simulations of weather and climate, and he is a member of the World Meteorological Organization working group on numerical experimentation (WGNE).*

IP04

Challenges in the First Principles Modelling of Magneto Hydro Dynamic Instabilities and their Control in Magnetic Fusion Devices

Marina Becoulet (CEA, France)

Wednesday,
July 4, 2018

16:40 – 17:30
Montreal Room

Chair:
Sinéad Ryan
(Trinity College
Dublin, Ireland)

The main goal of the International Thermonuclear Experimental Reactor (ITER) project is the demonstration of the feasibility of future clean energy sources based on nuclear fusion in magnetically confined plasma. In the era of ITER construction, fusion plasma theory and modelling provide not only a deep understanding of a specific phenomenon, but moreover, modelling-based design is critical for ensuring active plasma control. The most computationally demanding aspect of the project is first principles fusion plasma modelling, which relies on fluid models – such as Magneto Hydro Dynamics (MHD) – or increasingly often on kinetic models. The challenge stems from the complexity of the 3D magnetic topology, the large difference in time scales from Alfvenic (10-7s) to confinement time (hundreds of s), the large difference in space scales from micro-instabilities (mm) to the machine size (few meters), and most importantly, from the strongly non-linear nature of plasma instabilities, which need to be avoided or controlled. The current status of first principles non-linear modelling of MHD instabilities and active methods of their control in existing machines and ITER will be presented, focusing particularly on the strong synergy between experiment, fusion plasma theory, numerical modelling and computer science in guaranteeing the success of the ITER project.

Marina Becoulet *is a Senior Research Physicist in the Institute of Research in Magnetic Fusion at the French Atomic Energy Commission (CEA/IRFM). She is also a Research Director and an International Expert of CEA, specializing in theory and modelling of magnetic fusion plasmas, in particular non-linear MHD phenomena. After graduating from Moscow State University (Physics Department, Plasma Physics Division) in 1981, she obtained a PhD in Physics and Mathematics from the Institute of Applied Mathematics, Russian Academy of Science (1985). She worked at the Russian Academy of Science in Moscow, on the Joint European Torus in the UK, and since 1998 has been employed at CEA/IRFM, France.*

ID

Interdisciplinary Dialogue

ID01

The Colourful Theory, and Visible and Invisible Matter in the Universe: An Interdisciplinary Dialogue between Constantia Alexandrou and Petros Koumoutsakos

Constantia Alexandrou (University of Cyprus, Cyprus)
Petros Koumoutsakos (ETH Zurich, Switzerland)

Monday,
July 2, 2018

18:00 – 18:45
Montreal Room

Chair:
Erik Lindahl
(Stockholm University, Sweden)

Petros Koumoutsakos, a computational science researcher from ETH Zurich, will interview Constantia Alexandrou from the University of Cyprus about her domain of expertise – quantum chromodynamics. Many – if not most – fields in physics employ high performance computing (HPC), yet quantum chromodynamics (QCD) might be the premiere example of an area very difficult to understand outside of the field. In this dialogue, Constantia and Petros will look at what computational QCD achieves through the use of HPC, contextualizing it within a more general discussion of modern-day scientific computing. They will attempt to answer such questions as, "How do we 'compute' theory?" and "Will future computers change the way that theoretical physics 'experiments' are performed?"

Chromodynamics helps us understand our universe

The strong interaction is one of the four forces describing complex phenomena in the evolution of the universe from the quark gluon plasma formed just after the Big Bang at the birth of the cosmos to the formation of neutron stars. The bulk of visible matter in the universe is due to the strong interaction and understanding its properties requires the solution of quantum chromodynamics (QCD), a relativistic quantum gauge theory exhibiting confinement and asymptotic freedom properties that distinguish it from the other known theories. Solving QCD is carried out through large scale simulations using the largest supercomputers such as Piz Daint at the Swiss National Supercomputing Centre and Titan at Oak Ridge National Laboratory.
Recent progress in algorithms and access to larger computers have led to deeper understanding of the strong interactions, such as resolving the thirty-year-old puzzle of the spin of the proton and providing input for probing dark matter in the universe.

Constantia Alexandrou *is Professor of Physics at the University of Cyprus and Institute Professor at the Cyprus Institute. She holds a BA degree in Physics from Oxford University and a PhD in Theoretical Nuclear Physics from the Massachusetts Institute of Technology. She held research positions in Germany and Switzerland before joining the University of Cyprus. Her field of research is the study of the strong interactions using large-scale simulations of quantum chromodynamics. She is the Director of the Computation-based Science and Technology Research Center of the Cyprus Institute and coordinator of two Marie Sklodowska-Curie European Joint Doctoral programs in Computational Science.*

Petros Koumoutsakos *has received an education in Naval Architecture (NTUA Athens, University of Michigan), Aeronautics and Applied Mathematics (Caltech). He has conducted post-doctoral studies at the Center for Parallel Computing at Caltech and at the Center for Turbulent Research at Stanford University and NASA Ames. He was appointed as Chair for Computational Science at ETH Zurich in 2000. Petros is elected Fellow of the American Society of Mechanical Engineers (ASME), the American Physical Society (APS), the Society of Industrial and Applied Mathematics (SIAM) and the Collegium Helveticum. He has held visiting fellow positions at Caltech, the University of Tokyo, MIT and the Radcliffe Institute of Advanced Study at Harvard University. He is recipient of the Advanced Investigator Award by the European Research Council and led the team that won the ACM Gordon Bell prize in Supercomputing (2013). His team researches the how and what of computing as applied to problems ranging from fish swimming to nanotechnology and medicine.*

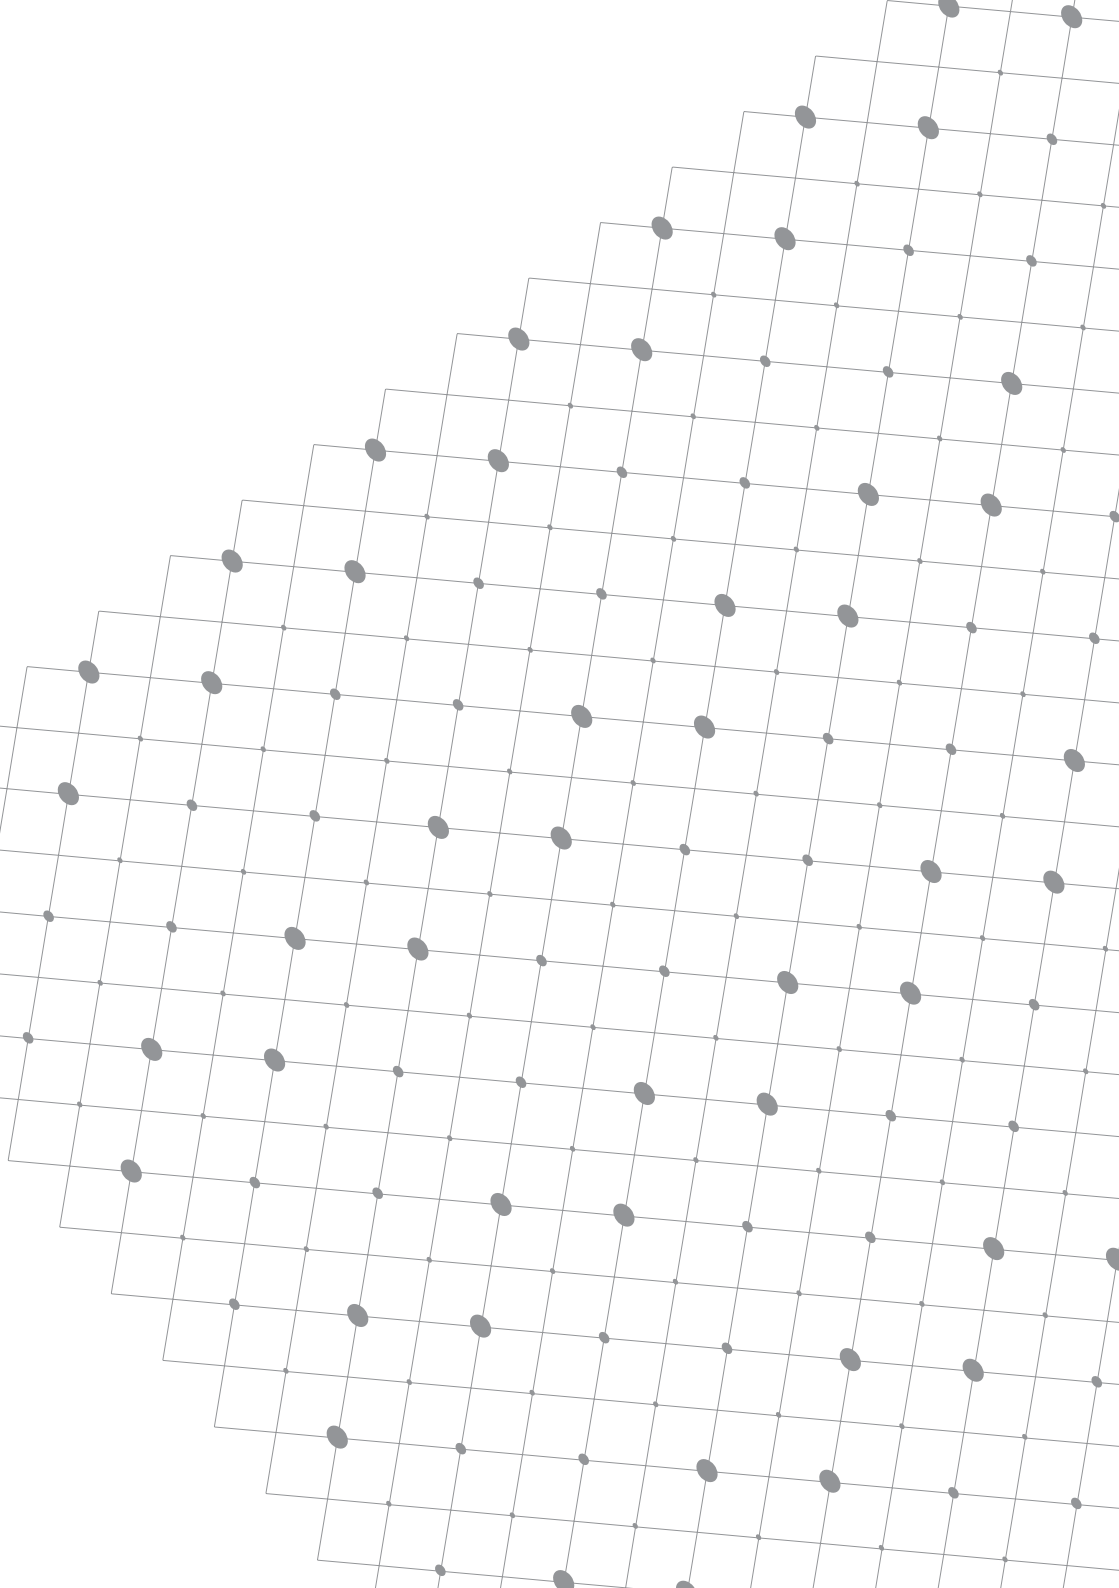

PNL

Panel Discussion

PNL01

Panel Discussion on Big Data vs. Fast Computation – Is HPC Facing a Game Change?

Moderators: ***Florina Ciorba*** *(University of Basel, Switzerland)*
Erik Lindahl *(Stockholm University, Sweden)*

Tuesday, July 3, 2018

09:00 – 10:15 Montreal Room

This panel discussion will address the main theme of PASC18: "Fast and Big Data, Fast and Big Computation". Are these two worlds evolving and converging together? Or is HPC facing a game-changing moment as the appetite for computation in the scientific computing community and industry is for a different type of computation than what we're used to? The panelists will discuss the critical challenges facing key HPC application areas in the next 5-10 years, based on a mix of knowledge and speculation. They will explore whether we need to make radical changes to our practices, methods, tools, and techniques to be able to use modern resources and make faster and bigger progress on our scientific problems. Do the current and projected developments of HPC systems and HPC software match the needs of computational scientists? Can we influence these developments in any meaningful way, or is it just a matter of adapting to the (new) hardware? Do computational scientists need to learn and apply techniques and algorithms from other areas, such as artificial intelligence and machine learning? Or is it that the other areas need to learn how to use and apply HPC to their algorithms?

Panelists: ***Eng Lim Goh*** *(Hewlett Packard Enterprise, USA)*
will bring a perspective from industry.
Nuria Lopez *(ICIQ, Spain)*
will bring a perspective from the chemistry domain.
Matthias Scheffler *(Fritz Haber Institute, Germany)*
will bring a perspective from the physics and materials domains.
Torsten Schwede *(University of Basel, Switzerland)*
will bring a perspective from the life sciences domain.

Eng Lim Goh *is the VP and CTO, HPC and AI at Hewlett Packard Enterprise. His current research interest is in the progression from data intensive computing to analytic, inductive machine learning, deductive reasoning and artificial specific to general intelligence. In collaboration with NASA he is currently principal investigator of a year long experiment aboard the International Space Station - this project won both the 2017 HPCwire Top Supercomputing Achievement and Hyperion Research Innovation Awards. In 2005, InfoWorld named Dr. Goh one of the 25 Most Influential CTOs in the world. He was included twice in the HPCwire list of People to Watch. In 2007, he was named Champion 2.0 of the industry by BioIT World Magazine and received the HPC Community Recognition Award from HPCwire. Dr. Goh completed his postgraduate work at Cambridge University, UK. He has been granted six U.S. patents with three pending.*

Nuria Lopez *graduated in Chemistry at the University of Barcelona, Spain in 1995, and received her PhD degree in Theoretical Chemistry at the same university in 1999. She held a postdoc position in the group of Prof. Jens K. Nørskov at the Center for Atomic-scale Materials Physics at the Technical University of Denmark, and since 2005 has led a research group focusing on the theoretical modeling of heterogeneous catalysis. She received an ERC grant in 2010 for "Bio2chem-d: Biomass to chemicals: Catalysis design from first principles for a sustainable chemical industry" and now participates in three European projects. She has been a member of the Scientific Steering Committee of PRACE since 2016.*

Matthias Scheffler *is director at the Fritz-Haber-Institut der Max-Planck-Gesellschaft. His research focuses on the developments and use of electronic-structure theory to understand the fundamental aspects of physical and chemical properties of surfaces, interfaces, clusters, nanostructures, and bulk materials. In recent years he has made crucial contributions to advance the field of computational materials science, with a specific focus on big-data analytics. He is coordinating the European Center of Excellence for Novel Materials Discovery (NOMAD) which maintains the largest repository for computational materials science data and a materials encyclopedia. NOMAD also offers big-data tools in order to advance materials science and engineering.*

Torsten Schwede *is Professor for Bioinformatics at the Biozentrum and director of sciCORE, the Center for Scientific Computing, at the University of Basel. At the Swiss Institute of Bioinformatics (SIB), Torsten leads the Data Coordination Centre for the Swiss Personalized Health Network (SPHN). The goal of SPHN is to establish nationwide interoperability of clinical, omics and other health-related data, allowing researchers in Switzerland to share data and collaborate efficiently to promote the development of personalized health in Switzerland. This includes establishing a network of secure HPC infrastructure core facilities to support biomedical research and clinical bioinformatics.*

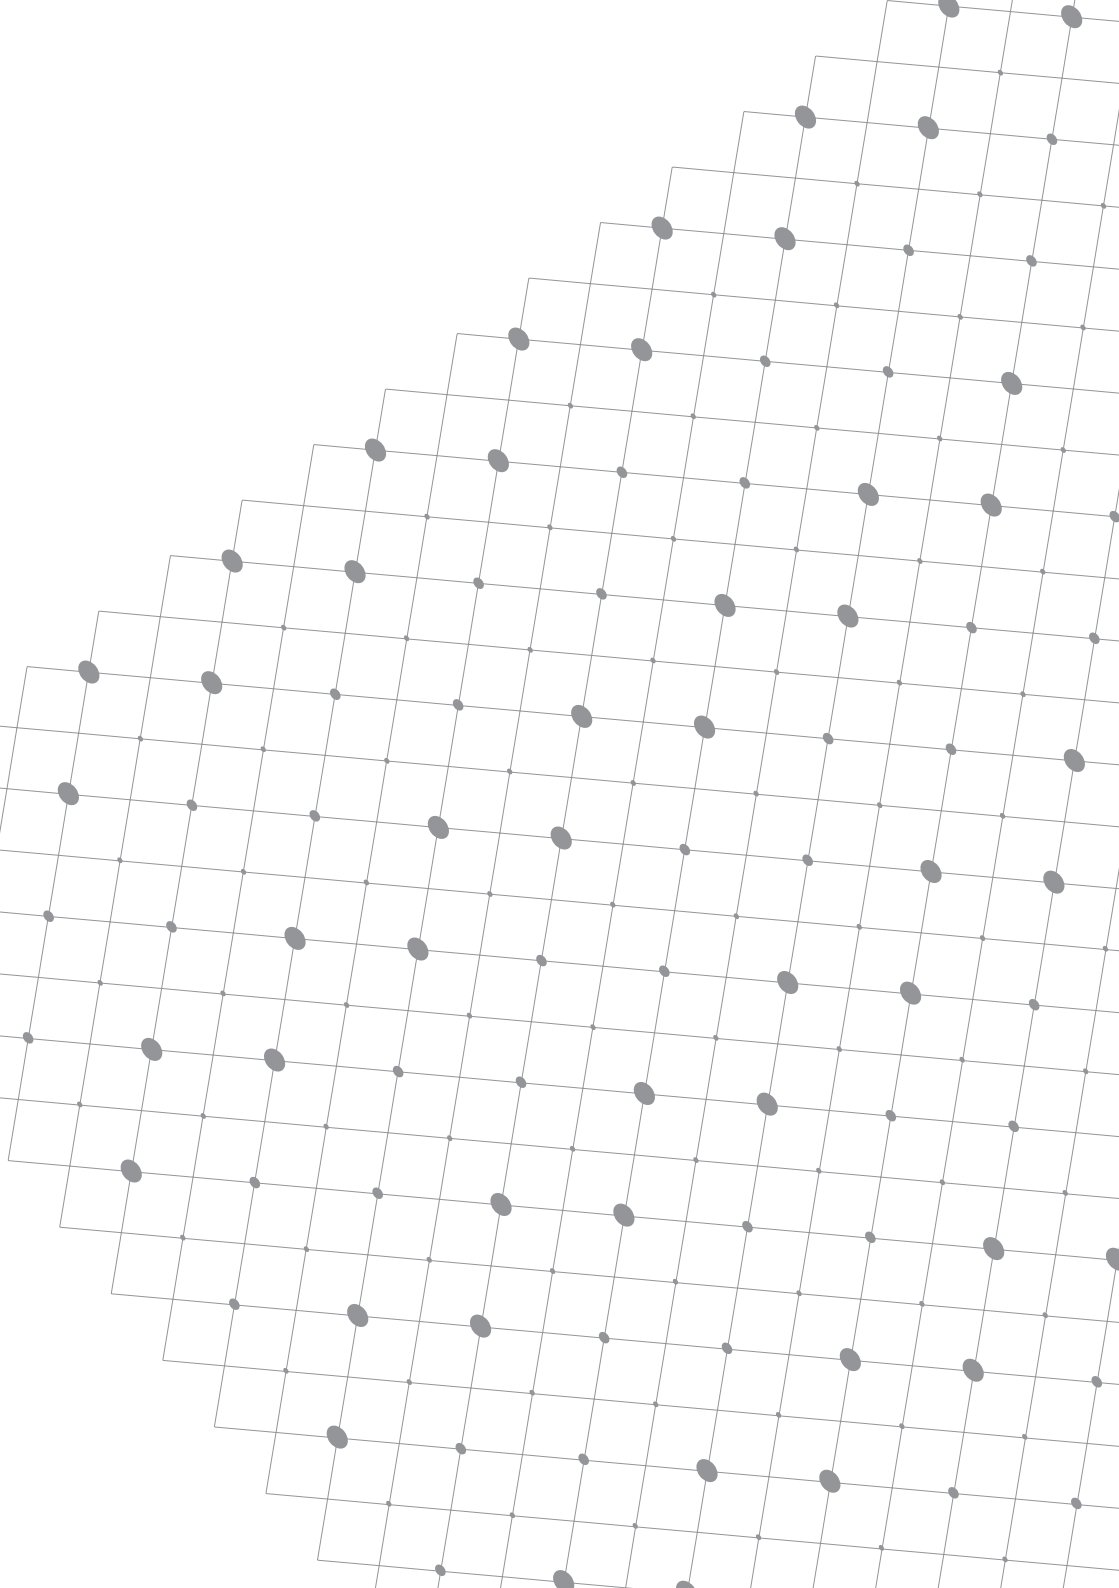

AP

ACM PASC18 Papers

AP01

ACM PASC18 Papers Session I

Chair: ***Sabine Roller*** *(University of Siegen, Germany)*

Extreme Computing for Extreme Adaptive Optics: The Key to Finding Life Outside our Solar System

Hatem Ltaief (King Abdullah University of Science and Technology, Saudi Arabia)

Monday, July 2, 2018

11:10 – 11:40 Montreal Room

The real-time correction of telescopic images in the search for exoplanets is highly sensitive to atmospheric aberrations. The pseudo-inverse algorithm is an efficient mathematical method to filter out these turbulences. We introduce a new partial singular value decomposition (SVD) algorithm based on QR-based Diagonally Weighted Halley (QDWH) iteration for the pseudo-inverse method of adaptive optics. The QDWH partial SVD algorithm selectively calculates the most significant singular values and their corresponding singular vectors. We develop a high performance implementation and demonstrate the numerical robustness of the QDWH-based partial SVD method. We also perform a benchmarking campaign on various generations of GPU hardware accelerators and compare against the state-of-the-art SVD implementation SGESDD from the MAGMA library. Numerical accuracy and performance results are reported using synthetic and real observational datasets from the Subaru telescope. Our implementation outperforms SGESDD by up to fivefold and fourfold performance speedups on ill-conditioned synthetic matrices and real observational datasets, respectively. The pseudo-inverse simulation code will be deployed on-sky for the Subaru telescope during observation nights scheduled early 2018.

Co-Author(s): **Dalal Sukkari**, **David Keyes** (King Abdullah University of Science and Technology, Saudi Arabia), **Olivier Guyon** (National Institutes of Natural Sciences, Japan)

The CLAW DSL: Abstractions for Performance Portable Weather and Climate Models

Valentin Clement (Center for Climate System Modeling, Switzerland)

Monday,
July 2, 2018

11:40 – 12:10
Montreal Room

In order to profit from emerging high-performance computing systems, weather and climate models need to be adapted to run efficiently on different hardware architectures such as accelerators. This is a major challenge for existing community models that represent very large code bases written in Fortran. We introduce the CLAW domain-specific language (CLAW DSL) and the CLAW Compiler that allows the retention of a single code written in Fortran and achieve a high degree of performance portability. Specifically, we present the Single Column Abstraction (SCA) of the CLAW DSL that is targeted at the column-based algorithmic motifs typically encountered in the physical parameterizations of weather and climate models. Starting from a serial and non-optimized source code, the CLAW Compiler applies transformations and optimizations for a specific target hardware architecture and generates parallel optimized Fortran code annotated with OpenMP or OpenACC directives. Results from a state-of-the-art radiative transfer code, indicate that using CLAW, the amount of source code can be significantly reduced while achieving efficient code for x86 multi-core CPUs and GPU accelerators. The CLAW DSL is a significant step towards a performance portable climate and weather model and could be adopted incrementally in existing code with limited effort.

Co-Author(s): **Sylvaine Ferrachat** (ETH Zurich, Switzerland), **Oliver Fuhrer**, **Xavier Lapillonne**, **Carlos E. Osuna** (MeteoSwiss, Switzerland), **Robert Pincus** (University of Colorado Boulder, USA), **Jon Rood** (National Renewable Energy Laboratory, USA), **William Sawyer** (ETH Zurich / CSCS, Switzerland)

AP02

ACM PASC18 Papers Session II

Chair: ***Jack Wells*** *(Oak Ridge National Laboratory, USA)*

A Parallel Solver for Graph Laplacians

Tristan Konolige (University of Colorado Boulder, USA)

Tuesday,
July 3, 2018

11:30 – 12:00
Montreal Room

Problems from graph drawing, spectral clustering, network flow and graph partitioning can all be expressed in terms of graph Laplacian matrices. There are a variety of practical approaches to solving these problems in serial. However, as problem sizes increase and single core speeds stagnate, parallelism is essential to solve such problems quickly. We present an unsmoothed aggregation multi-grid method for solving graph Laplacians in a distributed memory setting. We introduce new parallel aggregation and low degree elimination algorithms targeted specifically at irregular degree graphs. These algorithms are expressed in terms of sparse matrix-vector products using generalized sum and product operations. This formulation is amenable to linear algebra using arbitrary distributions and allows us to operate on a 2D sparse matrix distribution, which is necessary for parallel scalability. Our solver outperforms the natural parallel extension of the current state of the art in an algorithmic comparison. We demonstrate scalability to 576 processes and graphs with up to 1.7 billion edges.

Co-Author(s): **Jed Brown** (University of Colorado Boulder, USA)

Abstractions and Directives for Adapting Wavefront Algorithms to Future Architectures

Robert Searles (University of Delaware, USA)

Tuesday,
July 3, 2018

12:00 – 12:30
Montreal Room

Architectures are rapidly evolving, and exascale machines are expected to offer billion-way concurrency. We need to rethink algorithms, languages and programming models among other components in order to migrate large scale applications and explore parallelism on these machines. Although directive-based programming models allow programmers to worry less about programming and more about science, expressing complex parallel patterns in these models can be a daunting task especially when the goal is to match the performance that the hardware platforms can offer. One such pattern is wavefront. This paper extensively studies a wavefront-based miniapplication for Denovo, a production code for nuclear reactor modeling. We parallelize the Koch-Baker-Alcouffe (KBA) parallel-wavefront sweep algorithm in the main kernel of Minisweep (the miniapplication) using CUDA, OpenMP and OpenACC. Our OpenACC implementation running on NVIDIA's next-generation Volta GPU boasts an 85.06x speedup over serial code, which is larger than CUDA's 83.72x speedup over the same serial implementation. Our experimental platform includes SummitDev, an ORNL representative architecture of the upcoming Summit supercomputer. Our parallelization effort across platforms also motivated us to define an abstract parallelism model that is architecture independent, with a goal of creating software abstractions that can be used by applications employing the wavefront sweep motif.

Co-Author(s): **Sunita Chandrasekaran** (University of Delaware, USA), **Wayne Joubert**, **Oscar Hernandez** (Oak Ridge National Laboratory, USA)

AP03

ACM PASC18 Papers Session III

Chair: ***Michael A. Heroux*** *(Sandia National Laboratories, USA)*

Distributed, Shared-Memory Parallel Triangle Counting

Andrew Lumsdaine (Pacific Northwest National Laboratory, USA)

Tuesday,
July 3, 2018

11:30 – 12:00
Singapore Room

Triangles are the most basic non-trivial subgraphs. Triangle counting is used in a number of different applications, including social network mining, cyber security, and spam detection. In general, triangle counting algorithms are readily parallelizable, but when implemented in distributed, shared-memory, their performance is poor due to high communication, imbalance of work, and the difficulty of exploiting locality available in shared memory. In this paper, we discuss four different (but related) triangle counting algorithms and how their performance can be improved in distributed, shared-memory by reducing in-node load imbalance, improving cache utilization, minimizing network overhead, and minimizing algorithmic work. We generalize the four different triangle counting algorithms into a common framework and show that for all four algorithms the in-node load imbalance can be minimized while utilizing caches by partitioning work into blocks of vertices, the network overhead can be minimized by aggregation of blocks of work, and algorithm work can be reduced by partitioning vertex neighbors by degree. We experimentally evaluate the weak and the strong scaling performance of the proposed algorithms with two types of synthetic graph inputs and three real-world graph inputs. We also compare the performance of our implementations with the distributed, shared-memory triangle counting algorithms available in PowerGraph-GraphLab and show that our proposed algorithms outperform those algorithms, both in terms of space and time.

Co-Author(s): **Thejaka Kanewala** (Indiana University, USA), **Marcin Zalewski** (Pacific Northwest National Laboratory, USA)

MRG8 – Random Number Generation for the Exascale Era

Yusuke Nagasaka (Tokyo Institute of Technology, Japan)

Tuesday,
July 3, 2018

12:00 – 12:30
Singapore Room

Pseudo random number generators (PRNGs) are crucial for numerous applications in HPC ranging from molecular dynamics to quantum chemistry, and hydrodynamics. These applications require high throughput and good statistical quality from the PRNGs – especially for parallel computing where long pseudo-random sequences can be exhausted rapidly. Although a handful of PRNGs have been adapted to parallel computing, they do not fully exploit the features of wide-SIMD many-core processors and GPU accelerators in modern supercomputers. Multiple Recursive Generators (MRGs) are a family of random number generators based on higher order recursion, which provide statistically high-quality random number sequences with extremely long-recurrence lengths, and deterministic jump-ahead for effective parallelism. We reformulate the MRG8 (8th-order recursive implementation) for Intel's KNL and NVIDIA's P100 GPU – named MRG8-AVX512 and MRG8-GPU respectively. Our optimized implementation generates the same random number sequence as the original well-characterized MRG8. We evaluated MRG8-AVX512 and MRG8-GPU together with vender tuned random number generators for Intel KNL and GPU. MRG8-AVX512 achieves a substantial 69% improvement compared to Intel's MKL, and MRG8-GPU shows a maximum 3.36x speedup performance compared to NVIDIA's cuRAND library.

Co-Author(s): **Akira Nukada**, **Satoshi Matsuoka** (Tokyo Institute of Technology, Japan), **Kenichi Miura**, **John Shalf** (Lawrence Berkeley National Laboratory, USA)

AP04

ACM PASC18 Papers Session IV

Chair: ***Olaf Schenk*** *(Università della Svizzera italiana, Switzerland)*

A Massively Parallel Algorithm for the Approximate Calculation of Inverse p-th Roots of Large Sparse Matrices

Michael Lass (Paderborn University, Germany)

Tuesday, July 3, 2018

11:30 – 12:00
Sydney Room

We present the *submatrix method*, a highly parallelizable method for the approximate calculation of inverse *p*-th roots of large sparse symmetric matrices which are required in different scientific applications. We follow the idea of Approximate Computing, allowing imprecision in the final result in order to be able to utilize the sparsity of the input matrix and to allow massively parallel execution. For an $n \times n$ matrix, the proposed algorithm allows to distribute the calculations over n nodes with only little communication overhead. The approximate result matrix exhibits the same sparsity pattern as the input matrix, allowing for efficient reuse of allocated data structures. We evaluate the algorithm with respect to the error that it introduces into calculated results, as well as its performance and scalability. We demonstrate that the error is relatively limited for well-conditioned matrices and that results are still valuable for error-resilient applications like preconditioning even for ill-conditioned matrices. We discuss the execution time and scaling of the algorithm on a theoretical level and present a distributed implementation of the algorithm using MPI and OpenMP. We demonstrate the scalability of this implementation by running it on a high-performance compute cluster comprised of 1024 CPU cores, showing a speedup of 665× compared to single-threaded execution.

Co-Author(s): **Stephan Mohr** (Barcelona Supercomputing Center, Spain), **Hendrik Wiebeler**, **Thomas D. Kühne**, **Christian Plessl** (Paderborn University, Germany)

Balanced Graph Partition Refinement Using the Graph p-Laplacian

Dimosthenis Pasadakis (Università della Svizzera italiana, Switzerland)

Tuesday,
July 3, 2018

12:00 – 12:30
Sydney Room

A continuous formulation of the optimal 2-way graph partitioning based on the *p*-norm minimization of the graph Laplacian Rayleigh quotient is presented, which provides a sharp approximation to the balanced graph partitioning problem, the optimality of which is known to be NP-hard. The minimization is initialized from a cut provided by a state-of-the-art multilevel recursive bisection algorithm, and then a continuation approach reduces the *p*-norm from a 2-norm towards a 1-norm, employing for each value of *p* a feasibility-preserving steepest-descent method that converges on the *p*-Laplacian eigenvector. A filter favors iterates advancing towards minimum edgecut and partition load imbalance. The complexity of the suggested approach is linear in graph edges. The simplicity of the steepest-descent algorithm renders the overall approach highly scalable and efficient in parallel distributed architectures. Parallel implementation of recursive bisection on multi-core CPUs and GPUs are presented for large-scale graphs with up to 1.9 billion tetrahedra. The suggested approach exhibits improvements of up to 52.8% over METIS for graphs originating from triangular Delaunay meshes, 34.7% over METIS and 21.9% over KaHIP for power network graphs, 40.8% over METIS and 20.6% over KaHIP for sparse matrix graphs, and finally 93.2% over METIS for graphs emerging from social networks.

Co-Author(s): **Toby Simpson**, **Drosos Kourounis**, **Olaf Schenk** (Università della Svizzera italiana, Switzerland), **Kohei Fujita**, **Takuma Yamaguchi**, **Tsuyoshi Ichimura** (University of Tokyo, Japan)

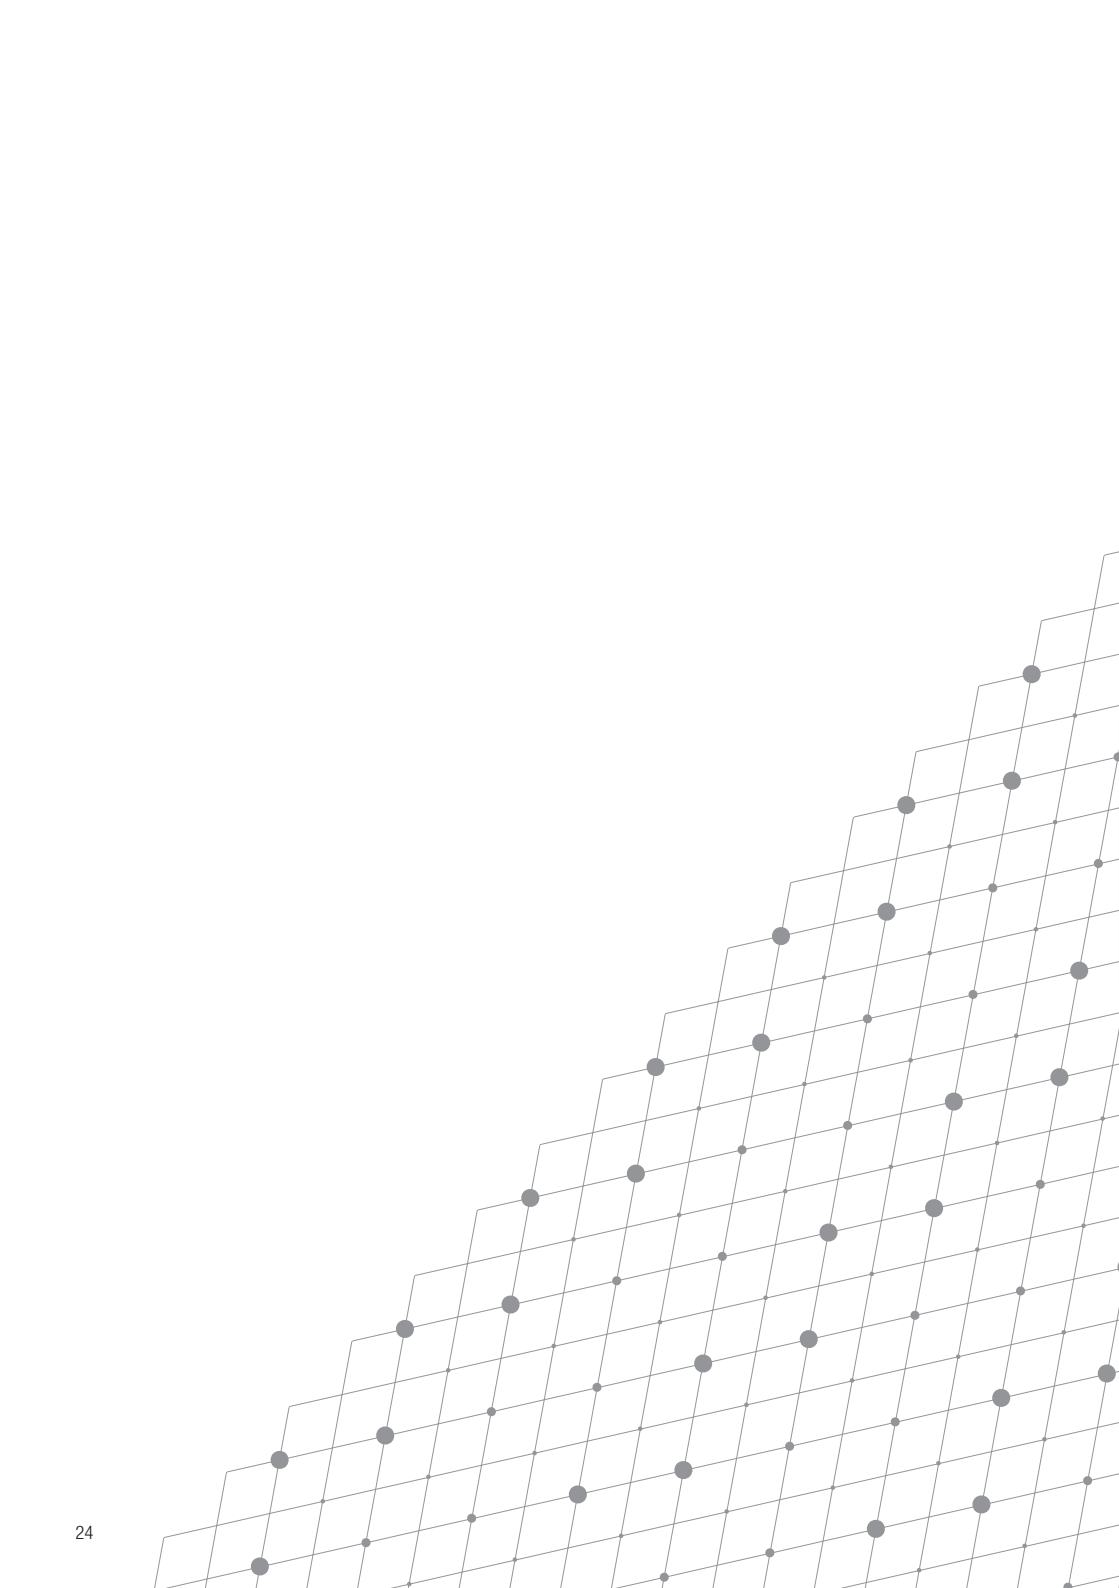

MS

Minisymposia

MS01

Adaptive Parallel Strategies for the Exploration of Challenging Search Spaces with Applications in Particle Simulations and Optimization, Part I

Organizer(s): ***Andreas Vitalis***, ***Marco Bacci***, ***Amedeo Caflisch***
(University of Zurich, Switzerland)

A common problem in numerical optimization and sampling is the detection of relevant states. These could be, for instance, the local minima on a rugged parameter surface or the transition state of a chemical reaction. For most cases, an exhaustive search for the optimal solution is intractable. Here, we focus on parallel sampling and optimization strategies relying on multiple replicas, most prominently, adaptive methods where all simulated replicas use the same propagator and sample the same underlying surface. In these methods, replica intercommunication is used to provide a global assessment as to which replicas are most interesting. This implies, in general, periodic data mining steps across replicas. Furthermore, in order to extract and utilize the gained information in post-processing, data must often be stored, which poses stringent data management and analysis challenges in particular for high-dimensional cases. The minisymposium wishes to discuss the following questions: What are meaningful and easily generalizable tools, strategies, and algorithms to guide the sampling/exploration? How can we maintain scalability and load balance? What types of post-processing algorithms can be applied to the generated data, and are those scalable to provide on-the-fly solutions to direct the exploration?

FAST - Goal-Oriented Adaptive Sampling of Protein Dynamics

Gregory Bowman (Washington University School of Medicine, USA)

Monday,
July 2, 2018

13:00 – 13:30
Samarkand Room

Molecular dynamics simulations are a powerful means of understanding conformational changes. However, it is still difficult to simulate biologically relevant time scales without the use of specialized supercomputers. Here, we introduce a goal-oriented sampling method, called fluctuation amplification of specific traits (FAST), for extending the capabilities of commodity hardware. FAST works by iteratively running a batch of simulations, building a Markov state model (MSM), and then using the last MSM to decide what subset of the states that have been discovered so far it would be most valuable to run the next set of simulations from. Importantly, the ranking function we use to choose starting points for each batch of simulations includes an exploitation term that favors states with desirable geometric properties and an exploration term that favors poorly sampled states. FAST outperforms conventional simulations and other MSM-based adaptive sampling algorithms by at least an order of magnitude. Furthermore, FAST yields both the proper thermodynamics and kinetics because, in contrast to many other enhanced sampling algorithms, the Hamiltonian used during individual simulations is unperturbed. Therefore, we expect FAST to be of great utility for a wide range of applications.

Applications and Advancements of the Progress-Index Guided Sampling Method in Molecular Dynamics Simulations

Marco Bacci (University of Zurich, Switzerland)

Monday,
July 2, 2018

13:30 – 14:00
Samarkand Room

Computer simulations of molecules offer unparalleled spatial and temporal resolution to the characterization of atomistic processes. However, the complexity and ruggedness of the free energy landscape often hamper the usefulness of brute-force molecular dynamics as most of the simulation time is spent in a few metastable states. To help in overcoming these limitations, we have recently developed the Progress Index Guided Sampling (PIGS) method. PIGS is a multi-replica unsupervised adaptive sampling protocol that aims to maximize phase space coverage by reseeding redundant replicas with interesting ones. Interesting replicas are detected on-the-fly by using a heuristic, which is informed by scalable data-mining algorithms that take as input a user-defined representation of the simulated system. Therefore, PIGS allows focusing the sampling enhancement on selected regions of interest without the need for reaction coordinates or external potentials. This also means that it is a straightforward task to retrieve in post-processing the thermodynamics and kinetics of the system within a Markovian approximation of the true dynamics. Here we show results from real-life simulations of biomolecules in explicit solvent performed with sampling enhancement on segments of different length. Additionally, we present algorithmic advancements, especially a fully scalable implementation of PIGS in the simulation engine GROMACS.

Co-Author(s): **Cassiano Langini**, **Andreas Vitalis**, **Amedeo Caflisch** (University of Zurich, Switzerland)

iMapD: Intrinsic Map Dynamics Exploration for Uncharted Effective Free Energy Surfaces

Roberto Covino (Max Planck Institute of Biophysics, Germany)

Monday,
July 2, 2018

14:00 – 14:30
Samarkand Room

Molecular dynamics (MD) simulations explore the configurational space of physical systems at their natural pace. Simulations extensively revisit typical configurations until rare and interesting transition events occur. Biasing the simulator away from the region already explored can, therefore, drastically accelerate the discovery of new regions, and is often the only way to gain access to all relevant states. We propose iMapD, an enhanced exploration simulation framework, where MD and machine learning adaptively bootstrap each other. Machine learning guides the search for important configurations by processing information from previous explorations. This search proceeds iteratively in an algorithmically orchestrated fashion without advance knowledge of suitable collective variables. The enhanced exploration occurs through strategically initialized short unbiased simulations, and does not rely on any unphysical force steering the dynamics of the system. Applied to a molecular sensor of lipid saturation in membranes, a dimer dissociation pathway not seen in millisecond long equilibrium simulations is discovered at the second iteration. In combination with path sampling techniques, iMapD enables us to characterize even the slowest dynamics of the system.

Co-Author(s): **Hendrik Jung**, **Gerhard Hummer** (Max Planck Institute of Biophysics, Germany), **Eliodoro Chiavazzo** (Politecnico di Torino, Italy), **Ioannis Kevrekidis** (Johns Hopkins University, USA)

Exploiting Task-Based Parallelism in Bayesian Uncertainty Quantification and Stochastic Optimization

Panagiotis Hadjidoukas (ETH Zurich, Switzerland)

Monday,
July 2, 2018

14:30 – 15:00
Samarkand Room

The mapping of Uncertainty Quantification (UQ) to computing architectures is a very challenging process and at the same time an essential aspect for all fields of simulation science. In UQ the aggregate scientific knowledge is obtained by ensembles of simulation runs, created dynamically by the employed UQ algorithm and scheduled on the available compute nodes. ∏4U is a computational framework that exploits the capabilities of massively parallel and hybrid computer architectures for large scale Bayesian uncertainty quantification, reliability analysis and stochastic optimization. At the core of the framework, a platform-agnostic task-parallel library supports nested parallelism and provides automatic load balancing on computing architectures that range from multicore systems to hybrid GPU clusters. The software is open-source and includes HPC implementations of algorithms such as Transitional Markov Chain Monte Carlo and Approximate Bayesian Computation. Experimental results using representative applications demonstrate the flexibility and excellent scalability of the proposed framework.

MS02

Capability Computing, Performance Portability, and Co-Design in the PASC Projects

Organizer(s): ***Joost VandeVondele*** *(ETH Zurich / CSCS, Switzerland)*

Selected PASC projects will present the scientific challenge they aim to solve by using high-end supercomputers, and in particular the computational approach adopted. The topics include astrophysics (smooth particle hydrodynamics), numerical weather and climate (stencils and grids), linear algebra for electronic structure (sparse matrix and tensor operations), and biomedical applications (fluid-structure interaction, machine learning).

The focus will be on aspects concerning (i) capability computing: how to scale to several hundreds/thousands of compute nodes, including the use of communication optimal algorithms and asynchronous communication; (ii) performance portability: how to address the growing diversity in hardware on a compute node, including generic software design, and auto-tuning; and (iii) co-design in these projects: how to engage with vendors to optimally exploit current hardware, and to provide feedback that has or will influence next-generation hardware.

Topics include: side by side comparisons of multi-core, many core, and GPU compute nodes; optimization techniques for flops, memory bandwidth, or network performance; JIT compilation of machine specific kernels; programming approaches such as the use of domain specific languages (DSLs), remote memory access (RMA), task based programming.

SPH-EXA: Optimizing Smooth Particle Hydrodynamics for Exascale Computing

Florina Ciorba, **Ruben Cabezon** (University of Basel, Switzerland)

Monday,
July 2, 2018

13:00 – 13:30
Sydney Room

Understanding fluid and plasma behavior under complex physical conditions forms the basis of highly important research questions. Numerical simulations of fluids in astrophysics and computational fluid dynamics are among the most computationally demanding calculations in terms of sustained floating point operations per second, which are expected to benefit from the upcoming Exascale high-performance computers. A well-known hydrodynamics solver is Smooth Particle Hydrodynamics (SPH). The parallelization of codes implementing the SPH method is not trivial due to the nature of the physics and the algorithms involved. The SPH-EXA project targets the design of a scalable and fault tolerant SPH-EXA mini-app. The scientific insights from the optimized executions of the SPH-EXA mini-app will be incorporated into current SPH-based production codes in the fields of astrophysics (SPHYNX, ChaNGa), and CFD (SPH-flow), resulting in, what we call, the SPH-EXA version of those codes. The SPH-EXA mini-app will employ advanced parallelization methods, scalable dynamic load balancing within single compute nodes and across massive numbers of nodes, and fault-tolerance mechanisms to sustain its scalable execution. An essential outcome of this project is a repository of experiments to enable verification, reproducibility, and portability of the execution and simulation results to other SPH-EXA codes.

Co-Author(s): **Lucio Mayer** (University of Zurich, Switzerland), **David Imbert** (NEXTFLOW Software, France)

Portability and Scalability of the COSMO Weather and Climate Model on Heterogeneous Architectures

Carlos E. Osuna (MeteoSwiss, Switzerland)

Monday,
July 2, 2018

13:30 – 14:00
Sydney Room

The clear evidence of the importance of high horizontal resolutions in the quality and accuracy of weather and climate simulations is demanding unprecedented computational capacity. Previous developments from PASC projects resulted in a GPU capable version of the COSMO model that provides significant speedup in the time to solution on NVIDIA GPUs, which allowed the first operational GPU enabled weather forecast system at MeteoSwiss as well as European-scale decadal climate simulations at unprecedented resolutions of 2 km. In order to improve the performance portability of COSMO in heterogeneous systems, recent efforts are supporting and optimizing COSMO for Xeon Phi KNL architecture and further improving the performance on accelerators exploiting advanced optimizations like task parallelism on the model that improve the performance on strong scalability regimes on massively parallel accelerators. Additionally, recent developments of a toolchain allow to combine all these advanced optimizations with a performance model specific to the domain and configuration of the model. We present results and performance comparisons for the COSMO 1km resolution configuration on Xeon Phi KNL and NVIDIA P100 systems.

Co-Author(s): **Stefan Moosbrugger**, **Oliver Fuhrer** (MeteoSwiss, Switzerland), **Felix Thaler** (ETH Zurich / CSCS, Switzerland), **Torsten Hoefler** (ETH Zurich, Switzerland)

Implementing a Sparse Tensor Linear Algebra Library for Electronic Structure Calculations

Juerg Hutter (University of Zurich, Switzerland)

Monday, July 2, 2018

14:00 – 14:30
Sydney Room

Sparse matrix-matrix multiplication is an essential building block for a wide range of algorithms in various scientific fields. For this task, the sparse matrix library DBCSR (Distributed Block Compressed Sparse Row) has been developed. Its multi-layered structure automatically takes care of and optimizes several computational aspects like parallelism (MPI, OpenMP, CUDA), data (cache) locality and on-the-fly filtering. As part of the PASC project, we are extending the library to include tensor algebra, based on the realization that most tensor operations can be mapped on matrix multiplications. First, we introduce the library, describing the repository on GitHub, how to compile it, the test methods, the tutorial, and the API in Fortran and C/C++. Then we give details on the implemented solutions to tackle scalability on large node-counts, based on a communication optimal algorithm with dynamically distributed load-balancing, implemented with remote memory access MPI communications. At the node level, we present a novel approach for the generation of optimal kernels based on autotuning and JIT compilation. Finally, we report the performance results, in terms of time-to-solution and energy-to-solution, of DBCSR on systems with Intel Xeon CPUs, Intel Xeon Phi Knights Landing (KNL) processors, and systems with NVIDIA GPUs.

Co-Author(s): **Alfio Lazzaro**, **Ilia Sivkov**, **Patrick Seewald** (University of Zurich, Switzerland)

AV-FLOW: A High-Performance Library for Fluid-Structure Interaction with Complex Materials and Transitional Flow

Dominik Obrist (University of Bern, Switzerland)

Monday, July 2, 2018

14:30 – 15:00
Sydney Room

The flow systems of the heart and the great blood vessels comprise complex materials (soft tissue) and flows at moderately high Reynolds numbers which may undergo transition from laminar to turbulent flow. Computational modelling of such fluid-structure interaction (FSI) problems requires efficient high-fidelity solvers for structure and flow as well as a robust scheme for coupling the two phases. We present a new FSI framework based on the immersed boundary method which has been developed for modelling cardiovascular flow systems. This high-performance library is optimized for parallel execution on the Cray XC40/50 system at CSCS. The structural and flow solvers use geometric domain decomposition for parallelization on multi-core multi-node platforms. The coupling between the structure and flow uses a parallel transfer library to minimize communication between the different computing cores. Compute intensive kernels were written in CUDA to make use of the GPGPUs on the nodes of the Cray XC40/50. We show performance benchmarks and different FSI test cases including a benchmark for solid inertia and a problem with transitional flow past an obstacle made of a complex material with fibers.

MS03

Computational Aspects of Heterogeneous Agents Macro

Organizer(s): ***Felix Kubler*** *(University of Zurich, Switzerland)*

Discrete-time, infinite-horizon, general equilibrium models are routinely used in macroeconomics and in public finance for exploring the quantitative features of model economies and for counterfactual policy analysis. One important question concerns the importance of household heterogeneity for the amplification and propagation of macroeconomic shocks.

In this session we bring together leading young researchers in the field to present alternative approaches to the computation of equilibria in dynamic stochastic models with heterogeneous agents and/or with financial frictions. Three of the papers directly propose new methods for the solution of models with a continuum of ex post heterogeneous agents.

Exploiting MIT Shocks in Heterogeneous-Agent Economies: The Impulse Response as a Numerical Derivative

Kurt Mitman (Stockholm University, Sweden)

Monday,
July 2, 2018

13:00 – 13:30
Nairobi Room

We propose a new method for computing equilibria in heterogeneous-agent models with aggregate uncertainty. The idea relies on an assumption that linearization offers a good approximation; we share this assumption with existing linearization methods. However, unlike those methods, the approach here does not rely on direct derivation of first-order Taylor terms. It also does not use recursive methods, whereby aggregates and prices would be expressed as linear functions of the state, usually a very high-dimensional object (such as the wealth distribution). Rather, we rely merely on solving nonlinearly for a deterministic transition path: we study the equilibrium response to a single, small "MIT shock'' carefully. We then regard this impulse response path as a numerical derivative in sequence space and hence provide our linearized solution directly using this path. The method can easily be extended to the case of many shocks and computation time rises linearly in the number of shocks. We also propose a set of checks on whether linearization is a good approximation. We assert that our method is the simplest, most transparent linearization technique among currently known methods. The key numerical tool required to implement it is value-function iteration, using a very limited set of state variables.

Co-Author(s): **Timo Boppart**, **Per Krusell** (Stockholm University, Sweden)

Solving Heterogeneous Agent Models with Nonconvex Optimization Problems: Linearization and Beyond

Michael Reiter (Institute for Advanced Studies, Austria)

Monday,
July 2, 2018

13:30 – 14:00
Nairobi Room

In this talk I present a methodology for the solution of dynamic stochastic general equilibrium models with heterogeneous agents, with an emphasis on models with nonconvex decision problems. First I present an implementation of a linearization method that makes the solution of large models feasible, using dimension reduction methods both on the states and on the equilibrium variables of the model. The linearized solution serves as a starting point to compute global approximation solutions, by providing a guess of the value functions and a suitable collocation grid. The method is applied to a two-asset model where households hold a financial asset and face a discrete housing choice.

Comparative Valuation Dynamics in Models with Financing Restrictions

Fabrice Tourre (Northwestern University, USA)

Monday,
July 2, 2018

14:00 – 14:30
Nairobi Room

This contribution develops a theoretical framework to nest many recent dynamic stochastic general equilibrium economies with financial frictions into one common generic model. Our goal is to study the macroeconomic and asset pricing properties of this class of models, identify common features and highlight areas where these models depart from each other. In order to characterize the asset pricing implications of this family of models, we study their term structure of risk prices and risk exposures, the natural extension of impulse response functions in economic environments exhibiting non-linear dynamics. Given our continuous time setup with a Brownian information structure, our study requires us to solve systems of non-linear partial differential equations of up to 4 state variables; the occasionally binding nature of our financial frictions give rise to a free boundary problem in the 4-dimensional state space. We use finite difference schemes coded in C++ and an iterative procedure to compute the equilibrium dynamics, the stationary distribution, the shock exposure and cost elasticities, and rho-mixing coefficients of our model.

Co-Author(s): **Paymon Khorrami** (University of Chicago, USA)

Self-Justified Equilibria: Existence and Computation

Felix Kubler (University of Zurich, Switzerland)

Monday,
July 2, 2018

14:30 – 15:00
Nairobi Room

In this talk we introduce "self-justified" equilibrium as a solution concept in stochastic general equilibrium models with a large number of heterogeneous agents. In each period agents trade in assets to maximize the sum of current utility and forecasted future utility. Current prices ensure that markets clear and agents forecast the probability distribution of future prices and consumption on the basis of current endogenous variables and the current exogenous shock. The forecasts are self-justified in the sense that agents use forecasting functions that are optimal within a given class of functions and that can be viewed as optimally trading off the accuracy of the forecast and its complexity. We show that self-justified equilibria always exist and we develop a computational method to approximate them numerically. By restricting the complexity of agents' forecasts we can solve models with a very large number of heterogeneous agents. Errors can be directly interpreted.

Co-Author(s): **Simon Scheidegger** (University of Zurich, Switzerland)

MS04

Distributed Training of Deep Neural Net Models for High Energy Physics

Organizer(s): ***Jean-Roch Vlimant*** *(California Institute of Technology, USA)*
Sofia Vallecorsa *(CERN, Switzerland)*
Wahid Bhimji *(Lawrence Berkeley National Laboratory, USA)*

The power of artificial neural nets at executing challenging tasks learned from data is very attractive to fields of physical science and especially to High Energy Physics. In recent years, there have been a significant number of articles reporting promising results with applying deep learning to HEP challenges. By virtue of the very large number of parameters of artificial neural nets, with deep and wide architecture, trained with stochastic gradient descent, it is mandatory to process a lot of representative data in order to obtain accurate models. Training of such models requires a tremendous amount of computing, and commonly takes days, if not weeks, to converge. GP-GPU technology has enabled a lot of this computation, thanks to the high level of parallelisation of the formalism of artificial neural net, but more however can be gained in parallelised calculation of the stochastic gradient descent. Supercomputing facilities are particularly suited for distributed training of deep neural nets, thanks to their large computation power and excellent connectivity. This minisymposium will address the current state-of-the-art and present performances in training models for high energy physics, with a particular view to software availability and to foster further utilization of supercomputers

Large Scale Training for Model Optimization

Felice Pantaleo (CERN, Switzerland)

Monday,
July 2, 2018

13:00 – 13:30
Osaka Room

In the recent years, several studies have demonstrated the benefit of using deep learning to solve typical tasks related to high energy physics data taking and analysis. The computational need for inference of a model once trained is rather modest and does not usually need specific treatment. The training of neural net models requires a lot of data, especially for deep models with numerous parameters. Training of such models has been made tractable with the improvement of optimization methods and the advent of GPUs well adapted to tackle the task of training neural nets. It is important to scale up the available network-training resources and to provide tools for optimal large-scale trainings. One of the avenues to further accelerate the training is via data parallelism, in which the computation of the gradients is computed on multiple subsets of the data in parallel and used collectively to update the model toward the optimum parameters. Several frameworks exist for performing distributed training, all with their strengths and limitations. In this context, our development of a new training workflow, which scales on multi-node/multi-GPU architectures with an eye to deployment on high performance computing machines is described.

Co-Author(s): **Maurizio Pierini** (CERN, Switzerland), **Jean-Roch Vlimant**, **Thong Nguyen** (California Institute of Technology, USA)

Training Generative Adversarial Models over Distributed Computing System

Gul Rukh Khattak (CERN, Switzerland)

Monday,
July 2, 2018

13:30 – 14:00
Osaka Room

In the High Energy Physics field, simulation of the interaction of particles in detectors material is a computing intensive task, even more so with complex and fined grained detectors. The complete and most accurate simulation of particle/matter interaction is primordial while calibrating and understanding the detector, but is seldomly required at physics analysis level, once several detector effects can hide slight imperfection in simulation. Some level of approximation is therefore acceptable and less computationally intensive approaches can be implemented. We present a fast simulation based on conditional generative adversarial networks. We use a dataset composed of the energy deposition from electron, photons, charged and neutral hadrons in a fine grained digital calorimeter. The training of these models is quite computing intensive, even with the help of GPGPU, and we propose a method to train them over multiple nodes and GPGPU using a standard message passing interface. We report on the scalings of time-to-solution. Further tuning of hyper-parameter of the models are rendered tractable and we present the physics performance of the best model obtained via a Bayesian optimization using gaussian processes. We demonstrate how a high performance computing center can be utilized to globally optimize these kinds of models.

Co-Author(s): **Sofia Vallecorsa**, **Federico Carminati** (CERN, Switzerland), **Jean-Roch Vlimant** (California Institute of Technology, USA)

Extreme Scale Deep Learning at NERSC

Thorsten Kurth (Lawrence Berkeley National Laboratory, USA)

Monday,
July 2, 2018

14:00 – 14:30
Osaka Room

We present various studies on very large scale distributed deep learning on HPC systems including the ~10k node Intel Xeon-Phi-based Cori system at NERSC. We explore CNN classification architectures and generative adversarial networks for HEP problems using large images corresponding to full LHC detectors and high-resolution cosmology convergence maps. We have explored distributed scaling in different deep-learning frameworks, including Caffe, TensorFlow and PyTorch with different communication layers, i.e. Google RPC or MPI-based approaches such as Intel MLSL, Uber Horovod and Cray's CPE ML Plugin. We describe various approaches for scaling out the training of single models up to the full Cori system. We further discuss recent work contrasting performance with different frameworks, systems and system architectures.

Co-Author(s): **Wahid Bhimji** (Lawrence Berkeley National Laboratory, USA)

Practical Scaling Techniques

Peter Messmer (NVIDIA Inc., Switzerland)

Monday,
July 2, 2018

14:30 – 15:00
Osaka Room

The need for large scale training of neural networks is stemming from the advent of ever growing labeled datasets in data science combined with the successes of deep learning at achieving super-human performance at pattern recognition tasks and others. Fast and powerful GP-GPU have enabled such trainings thanks to an impressive level of parallelisation of computation. There remain however large problems which may take days to weeks to converge. To this end, additional level of parallelisation across computing units are used for additional speed up. We present an overview of the practical techniques which can be used for scaling throughput of model training.

MS05

Foundations and Applications of Performance Engineering

Organizer(s): ***Gerhard Wellein***, ***Georg Hager***
(Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany)
Helmar Burkhart *(University of Basel, Switzerland)*

Achieving hardware and energy efficiency is important for current large-scale numerical simulations and will be a key component in the exascale era. In a world of heterogeneous, highly parallel computer architectures with deep memory hierarchies, complex application scenarios, and a broad spectrum of algorithms, a thorough analysis and understanding of the complex interaction of software, data structures, algorithms, and hardware features, a.k.a. performance engineering, is required for implementing codes that allow for portable performance on the computer generations to come. The minisymposium addresses a broad range of topics in performance engineering for modern HPC architectures, ranging from recent advances in performance models and tools supporting a "white-box" performance engineering approach to application performance tuning cases studies and "black-box" solutions. The presentations will point out the potentials and limitations of performance engineering activities and demonstrate the wide spectrum of performance models used in the performance engineering, including simple performance expectations, automatic model parameter selections, and analytic models.

Performance Engineering - Why and How?

Georg Hager (Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany)

Monday,
July 2, 2018

13:00 – 13:30
Singapore Room

We give an overview of Performance Engineering (PE) techniques used in scientific computing. Starting from a motivation based on resource efficiency, we demonstrate how PE can support computational science along several directions of thrust: Classification, insight, prediction, and optimization. There is a wide range of PE techniques, all of which have a modeling component of some kind. Such models come in all shapes and sizes, but we classify them on a scale from black to gray to white: Black-box models ignore all or most of the actual "inner workings" of hardware-software interactions and try to classify or predict interesting metrics automatically, based mostly on measurements. White-box models try to derive useful predictions from first principles, i.e., known properties of the hardware and the software. The wide range of "gray-box" models in between bridge the gap and use the best of both worlds. Examples from physics and high performance computing are given.

Towards a Discipline of Performance Engineering: Lessons Learned from Stencil Kernel Benchmarks

Danilo Guerrera (University of Basel, Switzerland)

Monday,
July 2, 2018

13:30 – 14:00
Singapore Room

An accurate measure of performance is often challenging, and the measurement process is seldom well documented and accurate, raising a credibility problem concerning the collected data. In this talk, we show how performance models and tools can work together. The selected test cases are kernels belonging to the stencil pattern, that is present in several scientific applications, ranging from geophysics to astronomy, fluid dynamics, image processing, and weather forecasts. First, we comment on how to pass from a description of a stencil to pseudo code and move to a modeling phase based on the "Kerncraft" tool for automatic Roofline and Execution-Cache-Memory performance modeling. After the automatic generation of compilable source code, we will focus on how to ensure the reproducibility of the performance results of its execution, using "PROVA!" a distributed workflow and system management tool for reproducible research. Knowing that a specific code performs accordingly with the model(s) can drive to the identification of relevant bottlenecks and therefore to potential optimizations. The ultimate goal is to generalize our approach to modeling, predicting and benchmarking, to a general application context.

Holistic Performance Engineering for Sparse Iterative Solvers

Jonas Thies (German Aerospace Center, Germany)

Monday,
July 2, 2018

14:00 – 14:30
Singapore Room

In many applications, sparse (linear and/or eigenvalue) solvers take up a large fraction of the overall runtime. We believe that the increasingly complex hardware of today's and future HPC systems has led to a gap in the understanding of the performance achieved by actual applications, many of which are still using a monolithic 'MPI only' approach despite the heterogeneous nature of the hardware. We have developed a new sparse solver library PHIST (https://bitbucket.org/essex/phist/) that defines a simple "kernel interface" layer inspired by MPI. Algorithms implemented in PHIST are portable in terms of software and performance as they only call building blocks of linear algebra via this interface. We have introduced simple performance models for these basic building blocks at the interface level, so that regardless of the backend providing the implementation, an overview of the optimization potential on the kernel level can be obtained, and performance pitfalls in the application (e.g. strided memory accesses) may be revealed. Available backends for PHIST include established libraries such as Trilinos/Epetra or PETSc, as well as more recent "MPI+X" approaches as implemented in Trilinos/Tpetra or our own kernel library GHOST (https://bitbucket.org/essex/ghost).

Machine Learning Framework for Performance Coverage Analysis

Tanzima Z. Islam (Western Washington University, USA)

Monday,
July 2, 2018

14:30 – 15:00
Singapore Room

Proxy applications are written to represent subsets of performance behaviors of larger, and more complex applications that often have distribution restrictions. They enable easy evaluation of these behaviors across systems, e.g., for procurement or co-design purposes. However, the intended correlation between the performance behaviors of proxy applications and their parent codes is often based solely on the developer's intuition. In this paper, we present novel machine learning techniques to methodically quantify the coverage of performance behaviors of parent codes by their proxy applications. We have developed a framework, VERITAS, to answer these questions in the context of on-node performance: (a) which hardware resources are covered by a proxy application and how well, and (b) which resources are important, but not covered. Since 2016, a more general machine learning framework has been developed around VERITAS which leverages deep learning techniques to automatically learn feature space and present information in a more intuitive fashion.

Co-Author(s): **Jayaraman J. Thiagarajan**, **Abhinav Bhatele**, **Todd Gamblin** (Lawrence Livermore National Laboratory, USA), **Martin Schulz** (TU Munich, Germany)

MS06

Large Scale Electronic-Structure Calculations on Modern and Future High-Performance Supercomputers

Organizer(s): ***Stefan Goedecker*** *(University of Basel, Switzerland)*
Andre Schleife *(University of Illinois at Urbana-Champaign, USA)*
Matthieu Verstraete *(Université de Liege, Belgium)*

Modern electronic-structure methods provide parameter-free simulations of properties across the whole spectrum of Physics, Chemistry, and Materials Science. They have become a bedrock of advanced analytical methods and are used systematically to interpret advanced experiments and complex interactions, with ever growing perspectives for more "realistic" systems, including defects, thermal and external fields, and transient phenomena. This minisymposium explores the next generation of electronic-structure software, which will lead users to exascale supercomputers, through highly efficient and highly parallel algorithms. We will showcase recent advances in ground- and excited-state calculations, spectroscopic quantities and transport, and adaptive methods which exploit different algorithms for different systems such as periodic/localized or many/few electrons per atom.

First-Principles Electron Transport with Phonon Coupling: Large Scale at Low Cost

Tue Gunst (Technical University of Denmark, Denmark)

Monday, July 2, 2018

13:00 – 13:30
Boston 3 Room

In the race towards high-performance nanometer-scaled devices the electronics industry now faces a major challenge from phonon-assisted tunneling. Despite the rapid size-reduction in experiments, system sizes still fall outside what is feasible for existing device models including electron-phonon coupling from first-principles. Therefore, the role of phonon-assisted tunneling in sub-10-nanometer gate-length devices has not been accurately quantified so far. We present a method that include phonon-assisted tunneling in large-scale first-principles calculations using a single "special thermal displacement" of the atomic coordinates at almost the same cost as elastic transport calculations [1] . We apply the method to ultrascaled silicon devices and demonstrate the importance of phonon-assisted band-to-band and source-to-drain tunneling. In a diode the phonons lead to a rectification ratio suppression in good agreement with experiments, while in an ultrathin body transistor the phonons increase off currents by four orders of magnitude, in agreement with our state-of-the-art perturbation theory calculations. In addition, electron-phonon coupling of nanostructured devices in operation conditions can change significantly from its bulk value [2] . This makes the method an appealing design tool for next-generation devices and nanomaterials. [1]T. Gunst *et al.*, Phys. Rev. B 96, 161404(R) (2017). [2]T. Gunst *et al.*, Phys. Rev. Lett. 118, 046601 (2017).

Large-Scale First-Principles Electronic Structure Calculations in Petascale and Exascale Supercomputers: A Real-Space Density Functional Theory Code

Jun-Ichi Iwata (The University of Tokyo, Japan)

Monday, July 2, 2018

13:30 – 14:00
Boston 3 Room

First-principles electronic structure calculation based on the Density Functional Theory (DFT) has been an indispensable tool for many fields of material science and engineering. With the development of supercomputers, the size of the targets of the first-principles DFT calculations becomes larger and larger, and nowadays, the target systems with a few hundreds to a thousand of atoms have been computable with standard plane-wave based DFT program codes. However, the computable size is not yet satisfactory to clarify the properties of materials in the situations close to realistic applications. We'd like to introduce our program code RSDFT, which has been developed to perform large-scale first-principles calculations on massively-parallel supercomputers including the Japanese flagship machine K computer. RSDFT is based on the real-space finite-difference pseudopotential method. Contrary to the standard plane-wave methods, the real-space method does not need to use Fast Fourier Transformations, which requires heavy communication burden in parallel computations, and therefore RSDFT shows rather good scalability even in the computations with tens of thousands of compute nodes. It has also been started to develop RSDFT for the next Japanese flagship computer called post-K computer. We aim to make the first-principles calculations of tens-of-thousand-atom systems easy as a daily work.

Co-Author(s): **Atsushi Oshiyama** (The University of Tokyo, Japan)

Potentialities of Wavelet Formalism towards a Reduction of the Complexity of Large Scale Electronic Structure Calculations

Luigi Genovese (CEA, France)

Monday, July 2, 2018

14:00 – 14:30
Boston 3 Room

For the last few years, the BigDFT software package has implemented a linear scaling Kohn-Sham density functional theory optimization algorithm based on Daubechies wavelets, where a minimal set of localized support functions are optimized in situ and therefore adapted to the physico-chemical properties of the system under investigation. We illustrate, from a general perspective, a quantitative method to identify and assess the partitioning of a large quantum-mechanical system into fragments. Our approach reduces arbitrariness in the fragmentation procedure and enables the possibility of assessing quantitatively whether the corresponding fragment multipoles can be interpreted as observable quantities associated with a system moiety. Such an approach is based on general grounds and its implementation is unrelated to the wavelet formalism. However, we show that the use of a minimal set of in situ-optimized basis functions allows at the same time a proper fragment definition and an accurate description of the electronic structure.

Co-Author(s): **Stephan Mohr** (Barcelona Supercomputing Center, Spain), **Laura Ratcliff** (Imperial College London, UK)

ABINIT on Pre-Exascale Supercomputers: Hybrid Parallelism and Numerical Stability

Marc Torrent (CEA, France)

Monday, July 2, 2018

14:30 – 15:00
Boston 3 Room

ABINIT is one of the most widely used electronic structure codes, implementing plane-wave based Density-Functional Theory. With multiple levels of parallelism, changing the code with every hardware evolution is a tedious task. To avoid obsolescence and allow adaptivity, an abstract layer for intensive low-level computing tasks has been introduced. Low-level sections have been rewritten specifically for a few hardware types. A global change of the hybrid parallelism is necessary to adapt the code to the new and future many-core architectures, as well as to memory bandwidth. A positive side effect of memory sharing is a better convergence of the diagonalization algorithm. Performances will be shown on Intel Xeon Skylake and Intel Xeon Phi KNL. Vectorization with large vectors and multithreading with more and more tasks induce a non-predictability of the floating-point operations that increase numerical noise and instabilities. We tackle this issue with the use of stochastic arithmetic to estimate the number of significant digits of each code section. Doing this, we can identify numerically sensitive code sections.

Co-Author(s): **Jordan Bieder** (CEA, France), **Yohan Chatelain**, **Pablo Oliveira** (Université de Versailles Saint-Quentin-en-Yvelines, France)

MS07

Machine Learning in Weather and Climate

Organizer(s): ***Peter Dominik Dueben***, ***Willem Deconinck*** *(ECMWF, UK)*
Rupert Ford *(Science and Technology Facilities Council, UK)*

The increasingly large amounts of data being produced by weather and climate simulations and earth system observations is sometimes characterised as a deluge. This deluge of data is both a challenge and an opportunity. The main opportunities are to make use of this wealth of data to 1) improve knowledge by extracting additional knowledge from the data and 2) to improve the quality of the models themselves by analysing the accuracy, or lack thereof, of the resultant simulation data. An example of the former case is improved prediction of large scale phenomena such as El Nino. An example of the latter is the improvement of a Physics parameterisation scheme through detailed analysis of the errors in a large number of datasets.

One way to realise these opportunities is to use machine learning approaches. As machine learning in weather and climate is a relatively new topic this minisymposium introduces the audience to how machine learning could be used in weather and climate and outlines its implications in terms of computing costs. To ground the ideas in concrete examples it also illustrates the use of machine learning in the weather and climate domain with practical examples.

Deep Learning in Weather and Climate, Part 1: The Domain Perspective

Peter Dominik Dueben (ECMWF, UK)

Monday, July 2, 2018

13:00 – 13:30
Rio Room

From the perspective of Earth System modelling, the use of machine learning, an in particular deep learning, is still in its infancy. There are many possible ways how deep learning could improve model quality or generate significant speed-ups for simulations. However, it has yet to be shown that deep learning can hold what it is promising for this application and its specific needs. This talk will provide an overview how deep learning may impact Earth System modelling in the future. We will provide examples how these methods have been used until today and discuss both limitations and prospects for their application. We will present results when using deep learning to improve model simulations for a toy model of atmospheric dynamics (the Lorenz'95 model). We will also show preliminary results that use deep neural networks that are trained from global atmospheric data to represent atmospheric dynamics and networks that are designed to speed-up parts of a weather forecast model at full complexity.

Deep Learning in Weather and Climate, Part 2: The Computing Perspective

Christoph Angerer (NVIDIA Inc., Germany)

Monday, July 2, 2018

13:30 – 14:00
Rio Room

In this presentation we will discuss ways in which deep neural networks can be integrated with traditional climate and weather simulations. In particular, we will be focusing on the design, implementation, and training of a deep convolutional neural network and its integration with the IFS Forecast Model inside RAPS as a new stand-alone radiation scheme. This work is a case study for how AI and large scale simulation may be applied on a cooperative basis and let the strengths of each converge to form a new tool for science.

Integrating Machine Learning Algorithms and HPDA Frameworks to Run Predictive Analytics on Large-Scale Climate and Weather Datasets

Alessandro D'Anca (CMCC, Italy)

Monday,
July 2, 2018

14:00 – 14:30
Rio Room

This work relates to the integration of a recurrent neural network algorithm (Long Short-Term Memory - LSTM) into Ophidia, a datacube-oriented High-Performance Data Analytics framework. More specifically, Ophidia provides a (big) datacube abstraction to the end users, while it physically builds its core set of functionalities (namely '*operators*') on top of an array-database system. Operators in Ophidia run in parallel over a cluster to tackle big data challenges on massive scientific datasets. At the array-database level, Ophidia allows end-user developing her own analytics functions (namely '*primitives*'), which by definition represent a sequential array-based data transformation. By implementing the LSTM algorithm as a *primitive* running over a long time series, machine learning capabilities can be integrated into Ophidia taking advantage of a HPDA approach applied over large-scale datasets. A couple of case studies have been considered: the former relates to the output of a WRF model running over the Brazilian region of Curitiba, whereas the latter includes both simulated data, through an unstructured grid forecasting model running at CMCC by the Ocean Predictions and Applications Division, and observations over the Apulia region in the South-East of Italy. Preliminary insights about the proposed approach seems promising and will be presented.

Co-Author(s): **Sandro Fiore** (CMCC Foundation, Italy)

Using Self-Organising Maps to Understand Relationships between Clouds and Cloud Controlling Factors

Samantha V. Adams (Met Office, UK)

Monday,
July 2, 2018

14:30 – 15:00
Rio Room

The long term global warming predicted for a doubling of carbon dioxide is known as the 'climate sensitivity'. Many future regional impacts of climate change become more serious for larger climate sensitivities. Current estimates of climate sensitivity from different climate models around the world vary by more than a factor of two, from approximately 2 to 5K. The reason for such a large range of estimates is due to the uncertainty around how clouds will change as the climate warms. Changes in clouds are hard to predict because they depend non-linearly on many interacting environmental factors. In this work we are interested in establishing whether machine learning can provide new insights into a) the factors controlling cloud changes and b) how various climate models represent these relationships. We use the Self-Organising Map (SOM), an unsupervised learning technique well suited to analysing high-dimensional data, to explore relationships between cloud controlling factors and compare the results to standard linear correlation. We find that potentially interesting new relationships not shown by linear correlation are revealed by the SOM technique.

Co-Author(s): **Mark Webb** (Met Office, UK)

MS08

On the Road to Exascale Computing: Turbulence Simulations of Complex Flows at the PetaFlops Pit Stop, Part I: Applications

Organizer(s): ***Ramesh Balakrishnan*** *(Argonne National Laboratory, USA)*
Philipp Schlatter *(KTH Royal Institute of Technology, Sweden)*

Wall resolved LES is prohibitively expensive, and a first principles based wall modeled LES, that is free of tunable parameters, is far from becoming an alternative to RANS as a predictive tool for design. On the spatial discretization front, most higher-order and spectral methods have had many of their successes in the realm of low to moderate Reynolds number flows, and on relatively less complex geometries. The bulk of the complex flow high Re simulations still employ nominally second-order accurate schemes in unstructured mesh finite volume flow solvers, and finite-element based solvers with modest polynomial order. Even in these fairly mature solvers, evidence suggests that merely running larger cases with an increased number of grid points and on larger computational domains does not always guarantee a better solution (in the LES sense). There is, therefore, a need for a two-level effort whereby benchmark higher-order simulations can serve to inform sub-grid models to improve the predictive capability of existing mature flow solvers. We hope to motivate a discussion along these lines with examples of results of higher-order simulations, as well as their role in assessing and improving the predictive capabilities of sub-grid models with more conventional flow solvers.

Study of the Cyclic Flow Variability in an Internal Combustion Engine Using Spectral Elements

George Giannakopoulos (ETH Zurich, Switzerland)

Monday,
July 2, 2018

13:00 – 13:30
Darwin Room

The accurate simulation of the in-cylinder processes is a critical step towards the realization of efficient internal combustion engines (ICEs). Advanced, high-fidelity simulations provide a deep insight of the generation and characteristics of in-cylinder turbulence, which is essential for the reduction of cycle-to-cycle variability (CCV) that affects adversely the efficiency of spark-ignited engines. In this work, the capabilities of the high-order, massively-parallel spectral element solver Nek5000 have been extended to perform a series of wall-resolved, implicit large eddy simulations (LES) of the gas exchange process during motored operation of a laboratory-scale ICE. Thirty consecutive engine cycles were simulated, providing a sufficient number of high-fidelity flow and temperature fields for statistical analysis. The presentation will discuss how turbulence is generated during the intake stroke by the high-speed valve jet and demonstrate the stochasticity of this process. In addition, the increased pressure during the compression stroke dramatically changes the thermodynamic properties that were found to strongly affect turbulence. This creates complex, unsteady, cycle-specific flow features that influence the momentum and thermal boundary layer structure and therefore heat transfer in the entire cylinder.

Co-Author(s): **Saumil Patel** (Argonne National Laboratory, USA), **Christos Frouzakis**, **Konstantinos Boulouchos** (ETH Zurich, Switzerland), **Paul Fischer** (University of Illinois Urbana-Champaign, USA), **Ananias Tomboulides** (Aristotle University of Thessaloniki, Greece)

Direct Numerical Simulation and Large Eddy Simulation of Canonical Flows for Wind Engineering Applications

Ramesh Balakrishnan (Argonne National Laboratory, USA)

Monday,
July 2, 2018

13:30 – 14:00
Darwin Room

With the availability of massively parallel computing platforms, large eddy simulation (LES) of flows in the atmospheric boundary layer (ABL) over complex terrain is being adopted as the tool of choice for predicting wind loads on wind turbines, and the transport of pollutants in urban layouts. Owing to the complex nature of the flow, however, it is unlikely that we will ever be able to resolve the details of the terrain and capture the relevant flow physics at the surface. As a result, it becomes necessary to model the sub-grid physics in the near surface region. Further, even for the idealized case of complex terrain without surface roughness, the high Reynolds number LES would require a sub-grid model that can account for the effects of anisotropic turbulence and flow separation. Prior LES of the neutral ABL over hill like geometries indicate that while the predictions match mean field measurements, detailed features such as flow separation, and the formation of ABL shear rolls are difficult to capture when the surface effects are modeled. This talk presents simulations of canonical wall bounded flows at high Reynolds numbers that were carried out to improve subgrid models for flows over complex terrain.

Co-Author(s): **Aleksandr Obabko** (Argonne National Laboratory, USA), **Paul Fischer** (University of Illinois Urbana-Champaign, USA)

Using a High Order Flow Solver for Generating DNS and LES Reference Databases for the Development of Turbulence Models

Ariane Frere (Cenaero, Belgium)

Monday, July 2, 2018

14:00 – 14:30
Darwin Room

Despite the fast increase of available computational resources, turbulence modeling is still required for the simulation of practical flows, due to the impossibility to capture the wide range of turbulent scales directly. Consequently, the validation of these turbulence model is critical. The work presented here concerns the generation of high-fidelity databases dedicated to the improvement of wall-shear stress modeled LES and RANS for applications in aeronautics and turbomachinery. As a first step, instrumentation guidelines for both DNS and LES, including a standardized set of statistical fields, as well as procedures to ensure statistical convergence and estimate confidence intervals, are presented. Subsequently two applications are presented, which have recently been computed by Cenaero using the in-house high order DGM solver Argo, namely the NACA4412 airfoil at Re=1.6M near stall, and the MTU-T161 turbine cascade at Re=90K and 200K. DGM maintains high accuracy on unstructured meshes and features strong computational efficiency at large scale. It is therefore a promising approach for generating high resolution DNS and LES on the complex geometry typical for turbomachinery. The respective computational resources for this project were allocated by PRACE on JuQueen at the Jülich Supercomputing Centre and MareNostrum IV at the Barcelona Supercomputing Center.

Co-Author(s): **Michel Rasquin**, **Koen Hillewaert** (Cenaero, Belgium)

Wall Resolved and Wall Modeled Simulations of Separated Flow over Airfoils

Ramesh Balakrishnan (Argonne National Laboratory, USA)

Monday, July 2, 2018

14:30 – 15:00
Darwin Room

Under realistic flight conditions, the turbulent flow physics over an airfoil is characterized by a wide range of time and length scales. The specific flow phenomenon, that is the subject of this presentation, is that of flow separation at high Reynolds number and high angles of attack (AoA). Due to its importance for airfoil/wing design, the accurate prediction of flow separation remains a challenge. Industrial prediction tools for separated flows rely, for the most part, on the RANS approach in which the predictive quality is heavily dependent on the effectiveness of the closure models to capture the underlying sub-grid physics. The availability of petascale computing resources, and simulation techniques, such as large eddy simulation (LES) of flows over complex geometry, open the possibility of using the results of wall resolved simulations to inform and calibrate lower-order sub-grid models used in RANS computations. The current talk will present a benchmark study of two state-of-the-art parallel flow solvers, Argo and PHASTA, which are employed to perform high-resolution wall resolved LES, and detached eddy simulations, respectively, of separated flow over the NACA-4412 airfoil for a Reynolds number of 1.6 million and AoAs of 12 and 13.87 degrees.

Co-Author(s): **Ariane Frère**, **Michel Rasquin**, **Koen Hillewaert** (Cenaero, Belgium), **Kenneth Jansen** (University of Colorado Boulder, USA), **Jun Fang,** (Argonne National Laboratory, USA)

MS09

Adaptive Parallel Strategies for the Exploration of Challenging Search Spaces with Applications in Particle Simulations and Optimization, Part II

Organizer(s): ***Andreas Vitalis***, ***Marco Bacci***, ***Amedeo Caflisch***
(University of Zurich, Switzerland)

A common problem in numerical optimization and sampling is the detection of relevant states. These could be, for instance, the local minima on a rugged parameter surface or the transition state of a chemical reaction. For most cases, an exhaustive search for the optimal solution is intractable. Here, we focus on parallel sampling and optimization strategies relying on multiple replicas, most prominently, adaptive methods where all simulated replicas use the same propagator and sample the same underlying surface. In these methods, replica intercommunication is used to provide a global assessment as to which replicas are most interesting. This implies, in general, periodic data mining steps across replicas. Furthermore, in order to extract and utilize the gained information in post-processing, data must often be stored, which poses stringent data management and analysis challenges in particular for high-dimensional cases. The minisymposium wishes to discuss the following questions: What are meaningful and easily generalizable tools, strategies, and algorithms to guide the sampling/exploration? How can we maintain scalability and load balance? What types of post-processing algorithms can be applied to the generated data, and are those scalable to provide on-the-fly solutions to direct the exploration?

Task-Based Parallelization of Replica Exchange Transition Interface Sampling in OpenPathSampling

David W. H. Swenson (University of Amsterdam, Netherlands)

Monday,
July 2, 2018

15:30 – 16:00
Samarkand Room

Path sampling methods, such as transition path sampling and transition interface sampling, are powerful tools for studying rare events. They perform Monte Carlo simulations in the space of trajectories, focusing the simulation effort on the transition itself to avoid spending long waiting times in the stable states. Since they are Monte Carlo approaches, they can use multiple walkers, but some approaches also use replicas from different path ensembles. In particular, replica exchange transition interface sampling (RETIS) involves simultaneously sampling trajectories from several path ensembles. However, even within a single ensemble, the lengths of the sampled trajectories can vary and are unpredictable. This makes load balancing an extremely challenging problem. This presentation describes the parallelization of RETIS in the software package OpenPathSampling using dask.distributed, a Python package for task-based programming. The task-based approach enables parallelization that provides optimal use of computational resources, not only by load balancing, but also by allowing the allocated resources to be scaled up or down according to the needs of the simulation. While the approach is described here in the context of path sampling, the same technique could be applied to many trajectory-based simulation methods.

Replica-Exchange Enveloping Distribution Sampling (RE-EDS) to Calculate Multiple Free-Energy Differences in a Single Simulation

Sereina Z. Riniker (ETH Zurich, Switzerland)

Monday,
July 2, 2018

16:00 – 16:30
Samarkand Room

Enveloping distribution sampling (EDS) allows the calculation of free-energy differences between multiple end states from a single simulation. A reference-state Hamiltonian is simulated which envelopes the Hamiltonians of the end states. The challenge when using EDS is the determination of optimal parameters for the reference-state Hamiltonian. Previously, the choice of parameters for an EDS simulation with multiple end states was a non-trivial problem that limited the application of the methodology. To overcome these limitations, we have generalized the replica-exchange EDS (RE-EDS) methodology to arbitrary systems. By exchanging configurations between replicas with different parameters for the reference-state Hamiltonian, major parts of the problem to choose optimal parameters are circumvented. Algorithms to estimate the energy offsets and optimize the replica distribution have been developed. Our approach was tested successfully using a system consisting of nine inhibitors of phenylethanolamine N-methyltransferase (PNMT), which were studied previously with thermodynamic integration and EDS.

On the Interpretation of Non-Equilibrium MD Trajectories

Tanja Schilling (University of Freiburg, Germany)

Monday,
July 2, 2018

16:30 – 17:00
Samarkand Room

As a researcher in statistical physics, one may often be interested in reducing the complexity of a many-particle system to the study of a set of relevant observables. If the system is in equilibrium, a systematic way to derive an equation of motion for the "relevant" observables from the microscopic dynamics has been known for some time as the "Mori-Zwanzig" formalism, which leads to the Langevin equation. In contrast, if the dynamics is not stationary, it is not a priori clear which form the equation of motion for an averaged observable will have. We adapt Mori-Zwanzig formalism to derive the equation of motion for a non-equilibrium trajectory-averaged observable as well as for its non-stationary auto-correlation function. We also derive a fluctuation-dissipation-like relation which relates the memory kernel and the autocorrelation function of the fluctuating force. In addition, we show how to relate the Taylor expansion of the memory kernel to experimental data, thus allowing to construct the equation of motion from direct measurements.

Co-Author(s): **Thomas Voigtmann** (German Aerospace Center, Germany), **Hugues Meyer** (University of Luxembourg, Luxembourg)

Dynamic Histogram Analysis to Determine Free Energies and Rates from Biased Simulations

Lukas S. Stelzl (Max Planck Institute of Biophysics, Germany)

Monday,
July 2, 2018

17:00 – 17:30
Samarkand Room

Transitions between metastable states govern many fundamental processes in biology, such as biomolecular folding. The underlying free energy surfaces can be obtained from simulations using enhanced sampling methods. We present an algorithm to calculate free energies and rates from enhanced sampling simulations on biased potential energy surfaces. Inputs are the accumulated times spent in each state or bin of a histogram, and the transition counts between them. For each of the states/bins optimal unbiased free energies are obtained by maximizing the likelihood of a master equation (i.e., first-order kinetic rate model). Unbiased rate coefficients for transitions between states can then be estimated. The resulting "dynamic histogram analysis method extended to detailed balance" (DHAMed) improves on the DHAM method. DHAMed yields accurate free energies in cases where the common weighted-histogram analysis method (WHAM) for umbrella sampling fails because dynamics within the windows is slow. We illustrate DHAMed with applications to proteins and RNAs and accurately estimate free energies from sets of short trajectories, providing a way forward for computational drug design. Our rate formalism can be used to construct Markov state models from biased simulations and we demonstrate its practical applicability by determining RNA folding kinetics from replica exchange molecular dynamics.

Co-Author(s): **Adam Kells**, **Edina Rosta** (King's College London, UK), **Gerhard Hummer** (Max Planck Institute of Biophysics, Germany)

MS10

Bridging the Software Productivity Gap for Climate and Weather Models

Organizer(s): ***Xavier Lapillonne*** *(MeteoSwiss, Switzerland)*
Valentin Clement *(Center for Climate Systems Modeling, Switzerland)*

Numerical weather prediction and climate models are large and complex software applications that need to run efficiently on today's and future massively parallel computer systems. The rapid change in these computing architectures and the increase in diversity are seriously affecting the ability to retain a single source code that runs efficiently in different architectures. Several weather models have successfully adapted their codebases to many-core and heterogeneous architectures like GPUs and Xeon Phi using a combination of multiple traditional programming models for parallel architectures like OpenMP, OpenACC and MPI. However porting existing large community codes to multiple architectures is a daunting task and leads to codes that are more complex and difficult to maintain. As a result in the past years numerous new technologies and approaches are emerging in order to provide new programming models, like domain-specific languages (DSLs) or source-to-source translation tools that can increase the productivity of development in weather codes while providing a high degree of performance portability. In this minisymposium we propose a discussion with some of the most prominent novel approaches where the new advances in programming models used for heterogeneous architectures in weather and climate models will be presented.

Experience on Porting Atmosphere Kernels on Many-Core Processors and Accelerators

Lin Gan (Tsinghua University, China)

Monday,
July 2, 2018

15:30 – 16:00
Rio Room

This talk includes a summary of our previous work that ports different atmosphere kernels onto various state-of-the-art platforms, including the Sunway TaihuLight system. Performance portability for atmosphere codes is no doubt a big challenge, so great efforts have to be made and patience is required as well. In addition to some experiences and lessons, we also take this opportunity to discuss on the novel Sunway processors. For Sunway system, different software is being developed to make it easy for applications to be ported.

Performance Portability for Next Generation HPC Architectures in E3SM via the Kokkos Programming Model

Luca Bertagna (Sandia National Laboratories, USA)

Monday,
July 2, 2018

16:00 – 16:30
Rio Room

This work converts the atmospheric dynamical core (HOMME) of the Energy Exascale Earth System Model (E3SM) from the current CPU-centric implementation, in Fortran 90, to a new performance-portable implementation, in C++ with the Kokkos performance-portability framework. HOMME simulates the dynamics and physical processes of the atmosphere. It is the most computationally demanding part of E3SM. Kokkos provides performance-portable multidimensional arrays and intraprocess parallel execution constructs. These form an abstraction layer over the hardware architecture of a compute node within a supercomputer. We will present results for the performance of our implementation on conventional CPU, Intel Xeon Phi, and Nvidia GPU; compare performance with the original Fortran on CPU and Xeon Phi; and discuss details of the implementation.

Co-Author(s): **Andrew Salinger**, **Irina Tezaur**, **Andrew Bradley**, **Michael Deakin**, **Daniel Sunderland**, **Oksana Guba** (Sandia National Laboratories, USA)

Experience Applying the PSyclone Configurable Domain Specific Compiler to the Met Office LFRic Model

Rupert Ford (Science and Technology Facilities Council, UK)

Monday,
July 2, 2018

16:30 – 17:00
Rio Room

Earth-system models tend to be large, complex codes developed by large teams of scientists over periods of years. However, the scale of the problems to be simulated calls for the highest levels of computational performance. Achieving good performance when both computer architectures and the underlying code base are constantly evolving is a complex challenge. In recent years, the use of Domain-Specific Languages (DSLs) as a potential solution to this problem has begun to be investigated. The UK Met Office's LFRic project is developing a new, Finite Element dynamical core and has adopted a DSL approach. In this talk we will describe this work and the functionality of the domain-specific compiler, PSyclone, which has been developed to process the (serial) code written by the natural scientists and generate the code required to run on massively parallel machines.

Co-Author(s): **Andrew R. Porter**, **Sergi Siso** (Science and Technology Facilities Council, UK)

Novel Programming Models for Large Geophysical Fluid Dynamics Models

Carlos E. Osuna (MeteoSwiss, Switzerland)

Monday,
July 2, 2018

17:00 – 17:30
Rio Room

Running operationally high resolution (~1km) global weather and climate models will be a milestone for the scientific community since there is clear evidence of the importance of high horizontal resolutions in the quality and accuracy of the simulations. Yet achieving this will pose serious computational challenges for large scientific codes that are developed using traditional programming models such as OpenMP and MPI. In order to adapt models to run efficiently on modern computing architectures and accelerators, numerous domain specific languages (DSL) and libraries that abstract architecture dependent optimizations have been proposed, like the GridTools libraries used operationally for running COSMO on GPUs. Yet these tools are specific to a domain or model, and have little reuse among them of architecture specific optimizers which leads to high maintenance costs. We present a novel programming model based on the GridTools ecosystem of libraries, a toolchain that allows to develop and interoperate various DSL frontends by providing domain and architecture specific optimizers. It aims at standardizing tools for performance portability by proposing a standard intermediate representation for weather and climate codes. We demonstrate the toolchain for the COSMO regional model and evaluate performance results compared to the operational model running on NVIDIA GPUs.

Co-Author(s): **Tobias Wicky**, **Stefan Moosbrugger**, **Torsten Hoefler** (ETH Zurich, Switzerland), **Oliver Fuhrer** (MeteoSwiss, Switzerland)

MS11

Computing the Effect of Risk

Organizer(s): ***Michel Juillard*** *(Banque de France, France)*

Many important economic phenomena relate to the notion of risk. Economic actors not only make decisions as a function of their current situation but also depending on their expectation of future developments. Because economic systems are not deterministic, future economic events are usually treated as stochastic phenomena. The general form of the problem at hand is to determine how the probabilistic distribution of future economic events influences current decisions. The wider the distribution, the more risk in today's decisions. The papers in this session present different computation challenges involved in attempting to describe the effect of risk on economic decisions.

Approximating Equilibria with Ex-Post Heterogeneity and Aggregate Risk

Elisabeth Proehl (University of Geneva, Switzerland)

Monday, July 2, 2018

15:30 – 16:00
Montreal Room

Dynamic stochastic general equilibrium models with ex-post heterogeneity due to idiosyncratic risk have to be solved numerically. This is a nontrivial task as the cross-sectional distribution of endogenous variables becomes an element of the state space due to aggregate risk. Existing global solution methods have assumed bounded rationality in terms of a parametric law of motion of aggregate variables in order to reduce dimensionality. In this paper, I remove that assumption and compute a fully rational equilibrium dependent on the whole cross-sectional distribution. Dimensionality is tackled by polynomial chaos expansions, a projection technique for square-integrable random variables, resulting in a nonparametric law of motion. I establish conditions under which the method converges and approximation error bounds. To illustrate the method, I compute the Aiyagari-Bewley growth model and the Huggett model with aggregate risk. In the former, I find that the bounded rationality assumption leads to significantly more inequality than in a fully rational equilibrium. Furthermore, more risk sharing in form of redistribution can lead to higher systemic risk. In the latter model, I find that prices increase with more stringent selling constraints, but are also more negatively skewed.

The Extended Perturbation Method

Martin M. Andreasen (Aarhus University, Denmark)

Monday, July 2, 2018

16:00 – 16:30
Montreal Room

This presentation introduces the extended perturbation method, which improves upon standard perturbation by removing approximation errors under certainty equivalence. Using the neoclassical growth model and a New Keynesian model, we show that extended perturbation achieves higher accuracy than standard perturbation when using third order approximations. We also show that extended perturbation generates stable approximations even when standard perturbation explodes. This paper also adds to the literature on downward nominal wage rigidities in the New Keynesian model, by showing that this friction only plays a significant role when using standard perturbation but not when using the more accurate extended perturbation approximation.

Co-Author(s): **Anders Kronborg** (Danish National Bank, Denmark)

Back in Time. Fast. Improved Time Iterations

Pablo Winant (Bank of England, UK)

Monday,
July 2, 2018

16:30 – 17:00
Montreal Room

We consider a new solution algorithm to solve nonlinear economic models using projections. For Bellman problems, our method is a variant of Howard's improvement steps. Contrary to the original improvements, it generalizes to models specified by equilibrium conditions in which case it is equivalent to the Newton-Raphson algorithm applied to one big nonlinear system of equations, without requiring the explicit inversion of the (memory-hungry) Jacobian. In particular, convergence is quadratic, i.e. much faster than regular time-iterations. Convergence of each gradient improvement step requires the (local) contractivity of the time-iterations operator. We show how this property relates to eigenvalues coming from local perturbation analysis, and how to estimate the local spectral radius of this operator close to a solution candidate. Gradient improvements can be implemented easily, essentially by composing the same elements as time-iterations. Our timing comparisons still suggest it performs much faster, especially when the number of dimensions or the number of grid points increase.

Taking Risk into Account with Higher-Order Approximations

Michel Juillard (Banque de France, France)

Monday,
July 2, 2018

17:00 – 17:30
Montreal Room

In a nonlinear model, expectation of future shocks entails expected benefits or expected losses. Rational agents can make decisions today so as to maximize expected benefits or minimize expected loss. These behaviors are related to economic concepts such as precautionary saving, asset prices, risk premium, term premium. By contrast, linear models are characterized by certainty equivalence, and in such environments, agents are indifferent to future uncertainty. One of the major benefits of using higher order approximation of a certain class of economic models, is the ability to analyse attitude towards risk. Computing higher-order approximation of DSGE models involves several computational challenges. Derivatives of the original model must be evaluated. These high dimensional objects must be stored in a convenient manner. Above second order, computations involve tensor algebra. A key component is a fast implementation of the Faa Di Bruno formula for the derivatives of the composition of two functions. Until now, all these steps have only been programmed in C++ in dynare++. They represent challenging tasks for a rather new programming language such as Julia and an interesting test case.

MS12

Engineering Scientific Software in times of Agile Development, Continuous Integration and Cloud Computing

Organizer(s): ***Guido Juckeland***
(Helmholtz-Zentrum Dresden-Rossendorf, Germany)

The dawn of Web 2.0 applications, smartphones/tablets and the omnipresent yet invisible cloud computing have dramatically changed our perception of the IT landscape in the last decade. At the same time these technologies delivered an abundance of new tools that are more suited to the needs of domain scientists: GitHub/GitLab with their inherent support for agile programming, user space package managers for easy software installations even when facing complex dependencies, and web portals to HPC systems or private clouds so that no prior knowledge is needed to use state-of-the-art compute resources. This minisymposium will showcase all these tools and how they are used in real scientific workflows. The shown best practices are easy to reproduce since they are all based on freely available software packages, so that the audience can use the presentations both as an inspiration but also as a kickstart for their own better science.

HPC-as-a-Service to Domain Scientists

Sunita Chandrasekaran (University of Delaware, USA)

Monday,
July 2, 2018

15:30 – 16:00
Sydney Room

Applications-as-a-service operating modes have changed the computing landscape in a multi-disciplinary research laboratory both from a user's and an HPC operator's perspective. Lately, applications are being offered as web-based user interfaces regardless of the actual location of the computation. To this end, the cloud computing revolution has had a wonderful side effect that everybody can now easily accept that certain tasks are transparently performed elsewhere - this talk will give an example from the bioinformatics application domain showing how cloud resources can be used for DNA sequencing. As such more and more HPC centers offer web-portals to access their systems along with applications developers also offering a web-based front-end so that the "obscure green font on black screen magic" of a typical SSH session is hidden from the end user. This enables both new groups to use HPC systems but also provides users a more error-proof and efficient way of using installed applications. This talk will highlight the criticality of an applications-as-a-service mode and will also discuss how Docker containers and Jupyter notebooks can be used for this easy-to-use application-as-a-service notion. The parallel benchmark suite from SPEC HPG organization will be used for demo purposes.

The Reality of Scientific Software Development is Agile - Best Practices and Lessons Learned

Guido Juckeland (Helmholtz-Zentrum Dresden-Rossendorf, Germany)

Monday,
July 2, 2018

16:00 – 16:30
Sydney Room

The reality of the development of scientific software is often far from the clearly structured, cascading work packages that grant applications require, but rather in the spirit of agile programming: Start from a working minimal prototype, always have running code, work in sprints (typically towards the end of reporting periods). Those characteristics actually match rather well onto the concept of agile programming. This talk will explain the principles of agile software development, the existing software tool support using the free-of-charge GitLab Libre Edition as an example, and the implementation of the processes, including team communication, continuous integration and publication of documentation.

Co-Author(s): **Tobias Frust** (Helmholtz-Zentrum Dresden-Rossendorf, Germany)

Using Jetstream and High Performance Remote Research Desktops to Lower the Barrier of Entry for HPC Resources

Robert Henschel (Indiana University, USA)

Monday,
July 2, 2018

16:30 – 17:00
Sydney Room

Indiana University operates two environments designed to lower the barrier of entry for HPC resources. One is the NSF Jetstream project, the first NSF funded cloud designed for those who have not previously used high performance computing resources. Jetstream provides users with long running virtual machines with a customizable software stack to meet the needs of non-traditional HPC applications. The other environment is a research desktop solution that is making high performance Linux desktops available remotely. The desktops contain all the normal HPC command line tools and allow for direct job submission to the HPC machines, but also provide access to interactive applications like Matlab, Comsol Multiphysics, R-Studio and Jupyter. The goal of both projects is to lower the barrier of entry and broaden adoption of traditional HPC and high-throughput computing environments. The talk will provide an architectural overview, use cases and experiences for operating such environments.

Co-Author(s): **Dave Hancock** (Indiana University, USA)

Spack: A Package Manager for Scientific Software

Todd Gamblin (Lawrence Livermore National Laboratory, USA), **Massimiliano Culpo** (EPFL, Switzerland)

Monday,
July 2, 2018

17:00 – 17:30
Sydney Room

On mainstream Linux distributions, package managers simplify the software installation process by providing pre-built, generic binaries. Users can leverage a wide variety of libraries and applications without knowing how to build from source. On HPC systems, users typically build from source, and software is notoriously complex. Building even a moderately sized parallel simulation code can be a major effort. Scientists who use applications codes must typically also know how to build them from scratch, along with tens or hundreds of dependency libraries. Spack is an open source package manager that handles the complexity of HPC environments and allows scientists to automatically and reproducibly install complex software stacks. It allows users to experiment with different compilers, optimizations, build options, and dependency versions, without in-depth build knowledge. Spack is built to handle the complexities of the HPC environment that seldom arise on commodity systems, such as swapping compilers and ABI-incompatible dependencies, cross-compilation, compiler runtime libraries, and optimized binaries. Spack has a rapidly growing community, with over 240 contributors at organizations worldwide. In this talk, we will introduce Spack, show how it can make scientists more productive, and give an overview of ongoing Spack projects and its development road map.

MS13

Generative Models and Density Estimator for High Energy Physics

Organizer(s): ***Sofia Vallecorsa*** *(CERN, Switzerland)*
Jean-Roch Vlimant *(California Institute of Technology, USA)*
Michela Paganini *(Yale University, USA)*

The Large Hadron Collider at CERN is smashing high density bunches of protons near the speed of light at a frequency of 40 MHz. Most of the thousands of particles emitted at each bunch crossing are measured and collected with building-sized detectors consisting of multiple sub-detectors each serving its own purpose. The simulation of the signal created by a particle interacting with such a detector is typically done with very detailed simulations and needs to be stepped infinitesimally over meters of material. This simulation is as much computing intensive as the geometry is complex. In current and future detector design, the fine-grained simulation of such a detector is taking a great part of the full computing budget of experiments and poses a computing challenge. While a great deal of effort is being made to parallelise such software, one possible avenue to reduce the computational requirements is with generative models from the field of deep learning. Such generative models have seen success in conditionally generating images and video of various types. We present how such models are built and trained, and how well they can capture the physics of particle interaction and help generate realistic samples for high energy physics analysis.

The Success of Deep Generative Models

Jakub Tomczak (University of Amsterdam, Netherlands)

Monday,
July 2, 2018

15:30 – 16:00
Osaka Room

Deep generative models allow us to learn hidden representations of data and generate new examples. There are two major families of models that are exploited in current applications: Generative Adversarial Networks (GANs), and Variational Auto-Encoders (VAE). The principle of GANs is to train a generator that can generate examples from random noise, in adversary of a discriminative model that is forced to confuse true samples from generated ones. Generated images by GANs are very sharp and detailed. The biggest disadvantage of GANs is that they are trained through solving a minimax optimization problem that causes significant learning instability issues. VAEs are based on a fully probabilistic perspective of the variational inference. The learning problem aims at maximizing the variational lower bound for a given family of variational posteriors. The model can be trained by backpropagation but it was noticed that the resulting generated images are rather blurry. However, VAEs are probabilistic models, thus, they could be incorporated in almost any probabilistic framework. We will discuss basics of both approaches and present recent extensions. We will point out advantages and disadvantages of GANs and VAE. Some of most promising applications of deep generative models will be shown.

Generative Models for Application-Specific Fast Simulation of LHC Collision Events

Maurizio Pierini (CERN, Switzerland)

Monday,
July 2, 2018

16:00 – 16:30
Osaka Room

We investigate the possibility of using generative models (e.g., GANs and variational autoencoders) as analysis-specific data augmentation tools to increase the size of the simulation data used by the LHC experiments. With the LHC entering its high-luminosity phase in 2025, the projected computing resources will not be able to sustain the demand for simulated events. Generative models are already investigated as the mean to speed up the centralized simulation process. Here we propose to investigate a different strategy: training deep networks to generate small-dimension ntuples of numbers (physics quantities such as reconstructed particle energy and direction), learning the distribution of these quantities from a sample of simulated data. In one step, one would then be able to generate the outcome of the full processing workflow (generation + simulation + reconstruction + selection).

Co-Author(s): **Dominick Olivito**, **Bobak Hashemi**, **Nick Amin** (UC San Diego)

Using Generative Models for Fast Clusters Simulations in the TPC Detector for the ALICE Experiment

Kamil Deja (Warsaw University of Technology, Poland)

Monday,
July 2, 2018

16:30 – 17:00
Osaka Room

Simulation of the events happening in the particle detector is a key component of many High Energy Physics experiments. Currently used Monte Carlo techniques allow to do it accurately, but their precision often comes at the expense of relatively high computational cost. In this work, we present a proof-of-concept solution for simulating clusters that occur after particle collision in the TPC detector in the ALICE Experiment at CERN. The new method we propose, dubbed ParticleGAN for simplicity, leverages recently developed Generative Adversarial Networks to learn the trajectories of particle tracks after collision. Although the quality of generated events is not even with the currently used solutions yet, ParticleGAN offer up to 10^3 speedups over the existing approaches. This applies also to other evaluated generative models namely Variational Autoencoders and variants of GANs. In this work we outline current bottlenecks of the proposed approach and discuss further steps that can allow to deploy the proposed generative models for simulation in production.

Co-Author(s): **Tomasz Trzcinski**, **Łukasz Graczykowski** (Warsaw University of Technology, Poland)

Generative Models for Simulating Highly Granular Calorimeters

Tobias Golling (University of Geneva, Switzerland)

Monday,
July 2, 2018

17:00 – 17:30
Osaka Room

Machine Learning techniques have been used in different applications by the HEP community: in this talk, we discuss the case of detector simulation. The need for simulated events, expected in future High Luminosity LHC experiments is increasing dramatically and requires new fast simulation solutions. We will describe R&D activities, aimed at reproducing the detector response and replace standard Monte Carlo simulation with generative models, typically used in computer vision applications. Two common aspects characterize many of these applications: the representation of input data as regular arrays of numerical values and the use of raw data as the input information to feed the network. Next generation HEP experiments are expected to be more and more characterized by detector components that could comply to this paradigm. Calorimeters of the ILC and CLIC detector concepts are effectively 3D arrays of sensors. We will introduce the first application of three-dimensional convolutional Generative Adversarial Networks and of Variational Auto Econders to the simulation of highly granular calorimeters. Finally we will present detailed validation studies comparing results to Monte Carlo simulation, showing the very good agreement we obtain for high level physics quantities and calorimeter response.

Co-Author(s): **Sofia Vallecorsa**, **Federico Carminati**, **Gul Rukh Khattak** (CERN, Switzerland), **Dalila Salamani** (University of Geneva, Switzerland)

MS14

How Fintech and Big Data Change and Challenge the Insurance Sector

Organizer(s): ***Jean-Michel Benkert***, ***Michelle Allgöwer***
(Baloise Group, Switzerland)

In the past couple years a large number of fintech and more recently insurtech startups have been founded and are challenging established players in the financial services sector. At Baloise – a Swiss company providing insurance services in Switzerland, Belgium, Germany and Luxembourg as well as banking services in Switzerland – we view startups as potential partners on our digital transformation journey rather than competition. This minisymposium aims to demonstrate what problems companies such as Baloise face in terms of digitizing their business and making use of their large amounts of data. To do so, the four sessions will cover the innovation framework Baloise employs in order to rapidly test prototypes, a presentation by Brainalyzed, a startup which aims to optimize and automatize investment decisions at Baloise using AI, a presentation about the challenges arising in the context of data warehouses and legacy systems, and a panel discussion with all the speakers.

Open Innovation at Baloise

Jean-Michel Benkert (Baloise Group, Switzerland)

Monday,
July 2, 2018

15:30 – 16:00
Nairobi Room

In recent years a large number of fintech and more recently insurtech startups have been founded and are challenging established players in the financial services sector. At Baloise – a Swiss company providing insurance services in Switzerland, Belgium, Germany and Luxembourg as well as banking services in Switzerland – we view startups as potential partners on our digital transformation journey rather than competition. Baloise has developed an open innovation framework with the goal of enabling easy and fast cooperation with startups and other external partners as well as intrapreneurs. In this session we will present this open innovation framework and its evolution over time as we have tailored it to the requirements of startups.

Artificial Intelligence for Automated Investment Management

Gunter Fischer (Brainalyzed, Germany)

Monday,
July 2, 2018

16:00 – 16:30
Nairobi Room

Due to increasing digitization in all sectors, the amount of available data is almost unlimited. The challenge is not only to manage this data, but to make it usable. Therefore, data analysis becomes a key success factor for organizations. Especially in the financial sector, data-driven applications are necessary to keep up with the fast-moving financial market and growing competition. The answer to low interest rates and high volatility in the market are automated data-driven investment processes. Data analysis using artificial intelligence (AI) is therefore becoming increasingly important. In this session we will give insights and some practical examples how we worked together with Baloise Asset Management to use some of their data to enhance the investment management process. We will show how the scalability of the learning solution helps to analyze even very complex problems in a short time and what our vision of AI in the financial world looks like.

The Challenges of Big Data for a Traditional Insurance Company

Christoph Geering (Baloise Group, Switzerland)

Monday,
July 2, 2018

16:30 – 17:00
Nairobi Room

With a company history of over 150 years, our IT landscape has grown in a highly fragmented way and consists of numerous legacy systems which have evolved over the last couple of decades covering a wide range of computer languages. Therefore a greenfield approach in terms of big data is out of question and the integration of data originating from these systems represents a costly and time-consuming challenge for Baloise. Securing the availability of internal data on one side and meeting the fast growing business requirements in connection with external (big) data integration on the other side is the balancing act of our digital transformation in the domain of business intelligence. How Baloise tackles these challenges and how the company benefits from cooperation with startups using artificial intelligence to boost this transformation will be explained in this session of the minisymposium. In the second part of this session an insight into projected use cases will be delivered to illuminate Baloise's strategic approaches related to machine learning and big data.

Panel Discussion on How Fintech and Big Data Change and Challenge the Insurance Sector

Jean-Michel Benkert (Baloise Group, Switzerland)

Monday,
July 2, 2018

17:00 – 17:30
Nairobi Room

Join us for a panel discussion on how fintech and big data change and challenge the insurance sector. The panelists are Dr. Gunter Fischer from Brainalyzed, an AI startup in the fintech space, Christoph Geering, responsible for business intelligence at Baloise Switzerland, and Dr. Jean-Michel Benkert, Innovation Manager at Baloise Group.

Co-Author(s): **Christoph Geering** (Baloise Group, Switzerland), **Gunter Fischer** (Brainalyzed, Germany)

MS15

Machine Learning and Quantum Chemistry

Organizer(s): ***Roland Lindh*** *(Uppsala University, Sweden)*

Machine Learning is right now a booming field of computer science which finds applications in the development of computer-human interfaces, in the analysis of medical data of huge populations, in the maintenance of cars, planes and elevators, and self-driving cars, to mention a few. During the last twenty years the field has gone through a development and refinement which has been spectacular. For some reason the use of the technology in pure science has been lagging behind; however, we are now starting to see the use of machine learning in the field of quantum chemistry. Here, the approach will enhance the performance of standard quantum chemical calculations – improving convergence, could serve as a tool for post-analysis of huge sets of *ab initio* results, or could simply replace computationally expensive procedures. Machine learning offers practical alternatives where standard quantum chemical simulations would be prohibitive. During the last few years a small number of quantum chemistry groups have explored the potential of machine learning – the results have been extraordinary and spectacular. Here in this minisymposium we would like to inspire by presenting four different applications in which machine learning is fundamental to success.

Quantum Machine Learning in Chemical Compound Space

Anders S. Christensen (University of Basel, Switzerland)

Monday,
July 2, 2018

15:30 – 16:00
Boston 3 Room

Many of the most relevant chemical properties of matter depend explicitly on atomistic and electronic details, rendering a first principles approach to chemistry mandatory. Alas, even when using high-performance computers, brute force high-throughput screening of compounds is beyond any capacity for all but the simplest systems and properties due to the combinatorial nature of chemical space, i.e. all compositional, constitutional, and conformational isomers. Consequently, efficient exploration algorithms need to exploit all implicit redundancies present in chemical space. I will discuss recently developed statistical learning approaches for interpolating quantum mechanical observables in compositional and constitutional space. Results for our models indicate remarkable performance in terms of accuracy, speed, universality, and size scalability.

Neural Networks Learning Quantum Chemistry

Olexandr Isayev (University of North Carolina, USA)

Monday,
July 2, 2018

16:00 – 16:30
Boston 3 Room

The development of accurate and transferable machine learning (ML) potentials for predicting molecular energetics is a challenging task. The process of data generation to train such ML potentials is a task neither well understood nor researched in detail. In this talk, we will present a fully transferable deep learning potential that is applicable to complex and diverse molecular systems well beyond the training dataset. Recently we introduced ANAKIN-ME (Accurate NeurAl networK engINe for Molecular Energies) or ANI in short.[doi: 10.1039/C6SC05720A] ANI is a new *method and sampling procedure* for training NNPs that utilizes a special kind of symmetry functions to build single-atom atomic environment vectors (AEV) as a molecular representation. To train ANI potential we use fully automated approach for the generation of datasets. [arXiv:1801.09319] It is based on the concept of active learning (AL). We show the use of our proposed AL technique develops a universal ANI potential, which provides very accurate energy and force predictions on the entire COMP6 benchmark. This universal potential achieves a level of accuracy on par with the best ML potentials for single molecule or materials while remaining applicable to the general class of organic molecules comprised of the elements CHNOSFCl.

Neural Network Representations of Non-Equilibrium Potential Energy Surfaces Sampled in Virtual Reality

David Glowacki (University of Bristol, UK)

Monday,
July 2, 2018

16:30 – 17:00
Boston 3 Room

I will outline recent developments in our group aimed at developing efficient potential energy surface (PES) representations of molecular geometries which are far from equilibrium. Recent progress developing a framework for interactive molecular dynamics in a multi-user virtual reality environment (combining rigorous cloud-mounted physical atomistic simulation with commodity virtual reality hardware) enables us to visualize and sample, with atomic-level precision, the structures and dynamics of complex molecular structures 'on the fly'.(arXiv:1801.02884) From within this framework, we can run real-time molecular dynamics (using density functional and semi-empirical theory), accelerating the sampling of high-energy reaction pathways. Combined, these reactive pathways represent a test set of molecular geometries whose energies and forces we then calculate at higher levels (e.g., explicitly correlated local coupled cluster theory), and fit using neural networks. The resultant PES provides coupled cluster quality energies at the cost of classical force fields, enabling us to run thousands of trajectories and thereby make comparisons with experimental dynamical observables in non-equilibrium regimes. I will illustrate this coupled virtual-reality-machine-learning workflow by focusing on recent applications where we have been studying heterogeneous reaction dynamics wherein cyano radicals (CN) undergo reactive scattering at the surfaces of liquids which are composed of long chain hydrocarbons.

Co-Author(s): **Silvia Amabilino**, **Lars Bratholm**, **Simon Bennie** (University of Bristol, UK)

Predicting the Stability of Solids with Density Functional Theory and Machine Learning

Miguel A. L. Marques (Martin Luther University Halle-Wittenberg, Germany)

Monday,
July 2, 2018

17:00 – 17:30
Boston 3 Room

We use a combination of machine learning techniques and high-throughput density-functional theory calculations to explore ternary compounds with the AB2C2 composition. We chose the two most common intermetallic prototypes for this composition, namely the tI10-CeAl2Ga 2 and the tP10-FeMo2B2 structures. We find that there may be ~10 times more stable compounds in these phases than previously known. These are mostly metallic and non-magnetic. While the use of machine learning reduces the overall calculation cost by around 75%, some limitations still exist, in particular for compounds involving the second-row of the periodic table or magnetic elements.

MS16

NP-Hard Computations: Massively Parallelizing Mixed-Integer Linear Programs

Organizer(s): ***Sharlee Climer*** *(University of Missouri - St. Louis, USA)*
Daniel Jacobson *(Oak Ridge National Laboratory, USA)*

Difficult combinatorial problems permeate virtually every area of the sciences, business, and government and many of these problems can be cast as mixed-integer programs (MIPs). A MIP is a mathematical definition of a problem that is comprised of a set of constraints and an objective function. In general, MIPs are NP-hard and require exponential amounts of computation time in the worst case. However, search strategies, such as branch-and-bound, branch-and-cut, and cut-and-solve have evolved to provide optimal solutions for many instances.

Although there has been great progress in this field and computational power has dramatically increased over the years, many important MIPs remain intractable and the use of massive parallelization appears to be a promising means to address this great need. However, many challenges lie ahead. This minisymposium will elucidate some of these challenges, while highlighting progress in this field. It includes a round table discussion with Michael Chan, Sharlee Climer, Daniel Jacobson, Sarah Powers, and Daniel Rehfeldt, and is open to conference participants. The goal of the discussions will be to explore and integrate high-performance expertise with domain-specific insights with an aim to identify strategies that may resolve these pressing challenges.

ug[SCIP-Jack, MPI]: A Massively Parallel Steiner Tree Solver

Daniel Rehfeldt (Zuse Institute Berlin, Germany)

Monday,
July 2, 2018

15:30 – 16:00
Singapore Room

The Steiner tree problem in graphs is a classical combinatorial optimization problem that commonly arises in practical applications as one of many variants. The general-purpose solver SCIP-Jack can solve the classical Steiner tree problem as well as 11 related problems to optimality. Furthermore, the solver comes with shared and distributed parallelization extensions by means of the UG framework that allow the parallelization of its branch-and-bound search. In this talk we briefly introduce the UG framework and go on to show how it is be combined with SCIP-Jack. The resulting solver ug[SCIP-Jack,MPI] has been able to solve several well-known Steiner tree instances for the first time to optimality.

Co-Author(s): **Yuji Shinano**, **Thorsten Koch** (Zuse Institute Berlin, Germany)

Parallel Cut-and-Solve: A Method for Solving Mixed-Integer Programs Utilizing Distributed Computational Power

Michael Chan (University of Missouri - St. Louis, USA)

Monday,
July 2, 2018

16:00 – 16:30
Singapore Room

A great number of problems can be cast as mixed-integer programs (MIPs), but their difficulty prevents many instances from being solved to optimality. The large amounts of distributed computational power becoming more readily available may provide a solution for tackling previously insolvable instances. But current MIP solvers' implementations using branch-and-cut can only be effectively parallelized to a certain degree. Here we present a potential method for parallelizing MIPs on a large scale using cut-and-solve and demonstrate our approach for a combinatorial genetics problem.

Looking Back to Look Forward in Solving Mixed-Integer Linear Programs

Sarah Powers (Oak Ridge National Laboratory, USA)

Monday,
July 2, 2018

16:30 – 17:00
Singapore Room

Numerous problems spanning a variety of fields and disciplines can be formulated as a Mixed-Integer Linear Program. Initially, only very small cases could be solved exactly. With the advent of greater computational power, the boundaries and limitations have continued to increase. However, continued progress is needed to solve many current day problems of interest (e.g., bioinformatics). The use of parallelization and other methods present a potential for further advancing the field and will be discussed during this talk.

Round Table Discussion: Embracing the Complexity Presented by Combinatorial Problems

Sharlee Climer (University of Missouri - St. Louis, USA)

Monday,
July 2, 2018

17:00 – 17:30
Singapore Room

Many real-world combinatorial problems can be cast as Mixed-Integer Linear Programs (MIPs). A MIP is a mathematical definition of a problem that is comprised of a set of decision variables, some or all of which are required to have integral values; a linear objective function to be minimized or maximized; and a set of constraints, all of which are linear equalities or inequalities. MIPs are generally NP-hard problems, yet progress in the field has led to limited successes in solving moderate to large size instances. The application of cutting planes is integral for state-of-the-art solvers that use Branch-and-Cut, but this application is inherently sequential. This round-table discussion will focus on the challenges faced when massively parallelizing computations for solving MIPs and explore strategies for circumventing these challenges.

Co-Author(s): **Michael Chan** (University of Missouri - St. Louis, USA), **Daniel Jacobson** (Oak Ridge National Laboratory, USA), **Daniel Rehfeldt** (Zuse Institute Berlin, Germany)

MS17

On the Road to Exascale Computing: Turbulence Simulations of Complex Flows at the Petaflops Pit Stop, Part II: Methods

Organizer(s): ***Philipp Schlatter*** *(KTH Royal Institute of Technology, Sweden)*
Ramesh Balakrishnan *(Argonne National Laboratory, USA)*

Computational Fluid Dynamics (CFD) is a natural driver for exascale computing both for academic and industrial cases, and has the potential for substantial societal impact, like reduced energy consumption, alternative sources of energy, improved health care, and improved climate models. This minisymposium focuses on algorithms and methods applicable on the way to exascale for CFD simulations. Application cases were discussed in Part I, whereas in Part II we focus on some of the relevant methodological aspects. The main driver is the EU funded Horizon 2020 project ExaFLOW and will feature presentations showcasing their work on addressing key algorithmic challenges in CFD in order to facilitate simulations at exascale, e.g. accurate and scalable solvers, data reduction methods (compression) and strategies to ensure fault tolerance and resilience. In particular, the talks in this minisymposium will highlight the following topics: Adaptive mesh refinement and adjoint-based error estimators, resilience in transient flow solvers, efficient communication operators using PGAS and mixed CG-HDG formulations for higher-order simulations.

Adaptive Mesh Refinement Based on Adjoint Error Estimators for Nek5000

Philipp Schlatter (KTH Royal Institute of Technology, Sweden)

Monday,
July 2, 2018

15:30 – 16:00
Darwin Room

The complex nature of turbulent fluid flows implies that the computational resources needed to accurately model relevant flow problems are virtually unbounded. Computational Fluid Dynamics (CFD) is therefore a natural driver for exascale. Adaptive mesh refinement (AMR) is identified as one key aspect in CFD: The solution of such problems is a-priori unknown, such that the mesh structure necessarily needs to be solution-dependent. We focus on AMR implemented in Nek5000, a direct numerical simulation code based on the spectral element method (SEM). The two main ingredients for AMR are tools for automatic mesh refinement and error estimators. New capabilities in Nek5000 enable the use of h-refinement method. Two methods are considered for estimating the error. The first is local and based on the spectral properties of the solution on each element. These spectral error indicators come with low overhead and are easily implemented but they provide only local error measure. The second method is goal-oriented adjoint error estimators which are based on similar work done for finite elements. AMR capabilities are demonstrated in Nek5000 and applied to steady and unsteady test cases, such as the lid-driven cavity, the flow past a cylinder and wing profiles, in two and three dimensions.

A Minimally Intrusive Low-Memory Approach to Resilience and Multi-Level Check-Pointing for Existing Transient Solvers

Chris D. Cantwell (Imperial College London, UK)

Monday,
July 2, 2018

16:00 – 16:30
Darwin Room

Exascale systems have the potential to allow much larger, more accurate and scale-resolving simulations of transient processes than can be performed on current petascale systems. However, with a much larger number of components, exascale computers are expected to suffer a node failure every few minutes. Many existing parallel simulation codes are not tolerant of these failures and existing resilience methodologies would necessitate major modifications or redesign of the application. We describe a novel, minimally intrusive approach to adding fault tolerance and scalable check-pointing to existing complex scientific simulation codes, demonstrated on the scalable, production-ready spectral/hp element framework Nektar++. Our approach combines the proposed user-level failure mitigation extensions to the Message-Passing Interface (MPI), with the concepts of message-logging and remote in-memory multi-level checkpointing. Logging MPI communication reduces the storage requirement of static data and allows a spare MPI process to rebuild these data structures independently of other processes. Remote in-memory checkpointing avoids disk I/O contention on large parallel filesystems. Forward-path and recovery-path performance of the resilience algorithm is analysed through experiments using the solver for the incompressible Navier-Stokes equations, and strong scaling of the approach is observed.

Co-Author(s): **Allan S. Nielsen** (EPFL, Switzerland)

Efficient Gather-Scatter Operations in Nek5000 Using PGAS

Niclas Jansson (KTH Royal Institute of Technology, Sweden)

Monday,
July 2, 2018

16:30 – 17:00
Darwin Room

Gather-scatter operations are one of the key communication kernels used in the computation fluid dynamics (CFD) application Nek5000 for fetching data dependencies (gather) and spreading results to other nodes (scatter). The current implementation used in Nek5000 is the Gather-Scatter library, GS, which utilises different communication strategies: nearest neighbour exchange, message aggregation, and collectives, to efficiently perform communication on a given platform. GS is implemented using non-blocking, two-sided message passing via MPI and the library has proven to scale well to hundreds of thousands of cores. However, the necessity to match sending and receiving messages in the two-sided communication abstraction can quickly increase latency and synchronisation costs for very fine-grained parallelism, in particular for the unstructured communication patterns created by unstructured CFD problems. ExaGS is a re-implementation of the Gather-Scatter library, with the intent to use the best available programming model for a given architecture. We present our current implementation of ExaGS, based on the one-sided programming model provided by the Partitioned Global Address Space (PGAS) abstraction, using Unified Parallel C (UPC). A detailed description of the library and implemented algorithms are given, together with a performance study of ExaGS when used together with Nek5000, and its co-design application Nekbone.

Co-Author(s): **Nick Johnson**, **Michael Bareford** (University of Edinburgh, UK)

Developing Methods for Exascale CFD Simulations at High Orders

David Moxey (University of Exeter, UK)

Monday,
July 2, 2018

17:00 – 17:30
Darwin Room

The established use of many-core computing hardware in current hardware, as well as its outlook to exascale computing platforms, poses a significant challenge for existing codes and numerical methods in achieving high levels of performance and scalability. In this presentation, we outline the approaches being taken in the Nektar++ spectral/hp element framework: a high-order framework that is used in a variety of fields, including the simulation of high-fidelity flow simulations for industry applications. We will present developments that aim to help improve on-node performance through modification of the underlying methods to improve data locality, as well as techniques that can be used to improve parallel scalability at high core counts.

Co-Author(s): **Martin Vymazal**, **Chris Cantwell**, **Spencer Sherwin** (Imperial College London, UK)

MS18

Addressing Resilience Challenges for Computing at Extreme Scale

Organizer(s): ***Aurelien Cavelan***, ***Florina Ciorba***
(University of Basel, Switzerland)

This minisymposium will discuss faults, errors, and failures that occur in extreme-scale computing systems. We want to increase awareness that resilience is a critical topic and that there are efforts and results that offer solutions to scientists and users of extreme-scale computing systems. Hardware level faults fall into two categories: hard faults are a consequence of permanent component failures (requiring repairs), while soft, or transient faults result from single upset events (e.g. a bit-flip in a memory cell) and have impermanent effects on the system. Hard faults typically result in system-wide failures, and in the absence of a fault management mechanism, an executed application is interrupted and its data is lost. The standard approach to cope with failures is to checkpoint, rollback and recover applications. However, it is expected that this approach may no longer be a viable solution on the upcoming Exascale systems. In contrast to hard faults, transient faults cannot always be detected and they can lead to Silent Data Corruptions (SDCs). Significant research efforts have been pursued to develop efficient SDC detectors, but there is not a perfect solution yet. The speakers selected for this minisymposium will give an overview of current solutions and future challenges.

Characterizing Faults, Errors and Failures in Extreme-Scale Computing Systems

Christian Engelmann (Oak Ridge National Laboratory, USA)

Tuesday, July 3, 2018

13:30 – 14:00 Montreal Room

Building a reliable supercomputer that achieves the expected performance within a given cost budget and providing efficiency and correctness during operation in the presence of faults, errors, and failures requires a full understanding of the resilience problem. The Catalog project develops a fault taxonomy, catalog and models that capture the observed and inferred conditions in current supercomputers and extrapolates this knowledge to future-generation systems. To date, the Catalog project has analyzed billions of node hours of system logs from supercomputers at Oak Ridge National Laboratory and Argonne National Laboratory. This talk provides an overview of our findings and lessons learned.

Easy and Efficient Multilevel Checkpointing for Extreme Scale Systems

Leonardo Bautista (Barcelona Supercomputing Center, Spain)

Tuesday, July 3, 2018

14:00 – 14:30 Montreal Room

Extreme scale supercomputers offer thousands of computing nodes to their users to satisfy their computing needs. As the need for massively parallel computing increases in industry, computing centers are being forced to increase in size and to transition to new computing technologies. While the advantage for the users is clear, such evolution imposes significant challenges, such as energy consumption and reliability. In this talk, we will discuss how to guarantee high reliability to high performance applications running in extreme scale supercomputers. In particular, we cover the tools necessary to implement scalable multilevel checkpointing for tightly coupled applications. This includes an overview of failure types and frequency in current HPC systems. The talk will also cover the theoretical analysis necessary to achieve optimal utilization of the computing resources. Moreover, we will discuss the internals of the FTI library tool, to study how multilevel checkpointing is implemented today.

Recent Results and Open Problems for Resilience at Scale

Yves Robert (École normale supérieure de Lyon, France)

Tuesday,
July 3, 2018

14:30 – 15:00
Montreal Room

The talk will address the following three questions: (i) fail-stop errors: checkpointing or replication or both? (ii) silent errors: application-specific detectors or plain old trustworthy replication? (iii) workflows: how to avoid checkpointing every task?

Panel Discussion on Upcoming Challenges at Exascale

Aurelien Cavelan (University of Basel, Switzerland)

Tuesday,
July 3, 2018

15:00 – 15:30
Montreal Room

This panel discussion will summarize the three preceding talks and offer guidelines and recommendations for the resilience challenges and opportunities available to scientists and users of extreme-scale computing systems.

Co-Author(s): **Leonardo Bautista** (Barcelona Supercomputing Center, Spain), **Yves Robert** (École normale supérieure de Lyon, France), **Christian Engelmann** (Oak Ridge National Laboratory, USA)

MS19

Advances in Computational Geosciences, Part I

Organizer(s): ***Ebru Bozdag*** *(Colorado School of Mines, USA)*
Dimitri Komatitsch *(CNRS, France)*

Recent advances in theory and numerical methods in parallel to the availability of high-quality massive data sets and high-performance computing provide unprecedented opportunities to improve our understanding of Earth's interior and its mechanism. The goal of this session is to bring computational and Earth scientists together to form a platform to discuss the current status, challenges and future directions in computational geosciences highlighting numerical simulations, the state-of-the-art HPC applications and their scientific outcomes. Contributions include, but are not limited to, the areas of earthquake engineering, passive and active-source seismic imaging, geodynamical modelling, magneto-fluid dynamics, etc. in conjunction with computational approaches such as numerical solvers, large-scale workflow, big data, optimisation strategies, etc. on HPC systems.

High-Resolution 3D Viscoelastic Full Waveform Imaging of a Real Seismic Dataset: The Volve Oil Field Studied up to 12 Hz

Dimitri Komatitsch (CNRS, France)

Tuesday,
July 3, 2018

13:30 – 14:00
Darwin Room

We will present recent advances on high-resolution 3D viscoelastic full waveform imaging, focusing in particular on a real seismic dataset for the Volve oil field, which we study up to 12 Hz on a large GPU cluster: the Piz Daint machine at CSCS in Switzerland. We will present both the workflow used and the final high-resolution pictures obtained.

Co-Author(s): **Vadim Monteiller** (CNRS, France)

Elastic Full Waveform Inversion with Active Seismic Data

Rene-Edouard Plessix (Shell Technology Center Amsterdam, Netherlands)

Tuesday,
July 3, 2018

14:00 – 14:30
Darwin Room

In exploration geophysics, seismic full waveform inversion is nowadays regularly applied. Most of the time, the acoustic approximation is made when active seismic data are inverted. This assumption reduced considerably the computation requirement. Indeed, depending on the type of acquisition (streamer versus Ocean bottom node), the number of simulations per iteration can range from several thousands to hundreds of thousands and the grid size can contain from millions to hundreds of millions of points. However, the acoustic approximation limits the range of applications. In complex geology with large earth parameter contrasts, ignoring the elastic effects can lead to significant artefacts in the full waveform results. In this presentation, I will discuss some of the challenges we face when considering the elastic propagation. The first one is obviously the large increase in computational cost, the second one is the multi-parameter inversion aspect. I will also discuss some examples to illustrate the importance of considering complex physical phenomena in our simulation during full waveform inversion.

Accelerating Low-Order Unstructured Finite Element Earthquake Simulation by Time-Parallel Computation on Recent HPC Architectures

Kohei Fujita (University of Tokyo, Japan)

Tuesday,
July 3, 2018

14:30 – 15:00
Darwin Room

Implicit low-order unstructured finite-element method is suitable for accurate modeling of time-history earthquake problems in complex geometry domains. However, it is costly due to large and random data access. To circumvent this bottleneck, we developed a time-parallel method that reduces the number of solver iterations by using sparse matrix vector products with multiple right-hand sides. This leads to reduction in total data transfer and random data access, and thus faster time-to-solution on recent architectures. We demonstrate the performance of the developed method and show application runs on earthquake problems.

Co-Author(s): **Tsuyoshi Ichimura**, **Takuma Yamaguchi**, **Muneo Hori**, **Lalith Maddegedara** (University of Tokyo, Japan)

Computational Models of Magnetic Field Generation in the Earth

Andy Jackson (ETH Zurich, Switzerland)

Tuesday,
July 3, 2018

15:00 – 15:30
Darwin Room

Earth's magnetic field is generated by fluid motion in the outer core by a process termed self-exciting dynamo action. In this process, electrically conducting fluid flows through a magnetic field, inducing electrical currents that reinforce the original magnetic field. The driving force for this is thought to be thermal convection. This process can be simulated on the computer in a self-consistent way, albeit in a parameter regime that is somewhat distant from planetary settings. In particular, the values of viscosity used are too large, and the prospects for reducing these viscosities to more appropriate values are remote. Despite this, the approach has met with great success and has demonstrated that magnetic fields can be generated in this way. Many features are quite Earth-like, most likely because the magnetic Reynolds number (the ratio of magnetic induction to magnetic diffusion) is in the correct regime. We will contrast conventional models with a different approach in which both inertia and viscosity are omitted from the equations at the outset. This approach, whilst in its infancy, holds the promise of providing complementary models of planetary magnetic field generation.

MS20

Challenges in Porting and Maintaining Atmospheric Codes on Emerging Hardware Architectures

Organizer(s): ***Richard Loft*** *(National Center for Atmospheric Research, USA)*
Oliver Fuhrer *(MeteoSwiss, Switzerland)*

Weather and climate models provide society with increasingly reliable weather forecasts and climate projections: critical information that can save both lives and money. With the advent of accelerators in high performance computing, several efforts around the world have begun porting weather and climate models to these emerging hardware architectures using different approaches and programming models. However, little attention has been given to the impact different porting approaches have on the maintainability of the codes. This minisymposium will provide an overview of the porting and maintainability experiences from four different community models in the US and Europe and will include a discussion of the pros/cons of different approaches in view of maintaining performance-portable production atmospheric community models.

Porting and Maintaining a GPU-Enabled and Performance-Portable Version of the Model for Prediction Across Scales (MPAS)

Richard Loft (National Center for Atmospheric Research, USA)

Tuesday, July 3, 2018

13:30 – 14:00
Rio Room

This talk will discuss our efforts to build a portable and maintainable CPU and GPU-enabled version of the Model for Prediction Across Scales (MPAS), a global atmospheric model currently used for meteorological studies and, in the future, climate research. Our approach uses a combination of OMP and OpenACC directives to achieve performance portability. We have focused on three target architectures, namely: traditional multi-core processors (e.g. Intel Xeon and IBM Power), many core processors like the Intel Xeon Phi, and of course NVIDIA GPUs. Leveraging tools that accelerate the optimization and verification process, our team has managed to keep the port synchronized with developer updates originating from the core MPAS science team and maintain readability and excellent performance across the three architectures. The results are encouraging, suggesting a path forward for our community models based on exposing parallelism to standard directives systems.

Co-Author(s): **Raghu-Raj Kumar** (National Center for Atmospheric Research, USA)

Experiences of Porting and Maintaining the ICON Model on Accelerators

William Sawyer (ETH Zurich / CSCS, Switzerland)

Tuesday, July 3, 2018

14:00 – 14:30
Rio Room

The dynamical core of the ICON model was first ported to accelerators within a PRACE 2IP project using the evolving OpenACC standard for accelerator directives. Building on this GPU porting effort the PASC-funded ENIAC project aims to port the full model using OpenACC compiler directives. While this technique can achieve good performance gains, some further improvement may be achieved by using hardware specific language or optimizations. In addition there are concerns about the long term maintainability: for example, would the investment be lost if OpenMP-4.5 usurped OpenACC as the de facto standard? The CLAW source-to-source translator ensures that OpenACC, OpenMP or other paradigms can be employed at the backend. CLAW also allows for a single-column abstraction of the physical parameterizations: this is more pleasing for the scientific developer, while allowing the introduction of domain- and hardware-specific optimizations. Finally, we envision that the inherently more static dynamical core will be reimplemented in a performance-portable platform-agnostic manner, e.g., using the GridTools framework used for the COSMO model. GridTools is implementing support for the underlying icosahedral grid. We present ICON-component examples for each of these paradigms, along with the resulting performance. We derive therefrom a long-term strategy for a maintainable ICON implementation for GPUs.

NOAA Model Development Activities Targeting Exascale

Mark Govett (NOAA, USA)

Tuesday,
July 3, 2018

14:30 – 15:00
Rio Room

In the last two years, significant efforts have been made to adapt the National Weather Service's Finite Volume-cubed (FV3) model to run efficiently on GPU processors. The main criteria for adapting and parallelizing the FV3 code have been (1) minimize changes so the code remains acceptable to the scientists, (2) maintain portability and performance on the CPU, and (3) demonstrate bitwise exact results between the original code, and GPU enabled code. Code adaption and parallelization of the FV3 were based on successful experiences with the Non-hydrostatic Icosahedral Model (NIM), which demonstrated good performance portability between CPU, GPU and MIC processors with a single source code. Work on adapting the FV3 for GPU architectures has proven much more difficult than prior work with the NIM code. This talk will explain some of the complexities and challenges with the FV3 parallelization. We will also describe recent efforts to design and develop prototype models for diverse, highly parallel exascale computing systems expected in 5-10 years.

Experience and Challenges with Maintaining a GPU-Capable Version of COSMO in a Production Environment at MeteoSwiss and ETH

Xavier Lapillonne (MeteoSwiss, Switzerland)

Tuesday,
July 3, 2018

15:00 – 15:30
Rio Room

Advances in computer technologies can greatly enhance our climate and weather modeling capabilities by allowing to increase grid resolution or the complexity of physical processes being accounted for. In order to benefit from hardware improvements, models need to be adapted which is a challenging task. This talk will first present the port of the climate and weather model COSMO to heterogeneous GPU architectures which was achieved using different technologies. A rewrite using a domain specific language (DSL) allowing a high-level hardware agnostic formulation of the model equations was considered for some components while OpenACC compiler directives have been used for the rest of the model. Performance results and implication to the production environment at MeteoSwiss for weather forecast as well as for climate simulation on the leadership-class heterogeneous HPC system Piz Daint in Switzerland will be shown. Finally, challenges related to the different technologies employed, a domain-specific language and compiler directives, in terms of maintenance, performance portability, compiler support, and user acceptance will be discussed.

MS21

Computational Solutions to Large-Scale Data Management and Analysis Challenges in Personalized Health

Organizer(s): ***Leila Tamara Alexander***, ***Torsten Schwede***
(Swiss Institute of Bioinformatics, Switzerland)

Personalized health aims to provide the right treatment at the right time for each individual. A major premise is that empowerment with more knowledge leads to better decision-making. Tailored, predictive interventions have the potential to change from a reactive to a preventative approach, thereby significantly extending the duration of health. To achieve this, highly performant computational environments to store, transfer, analyse and integrate data produced at astonishing rates are prerequisites. Today personalized health becomes a practical reality, and the exploding 'big data' in healthcare provides exciting IT opportunities.

This minisymposium highlights computational approaches to enable the next steps in biomedical discovery. The introductory keynote will outline the current clinical and scientific needs for advanced computational approaches. The second session is about computational methods that already revolutionise medical practice through virtual reality systems for personalized medical surgeries. The third talk discusses the utilisation of HPC for finding new anti-cancer therapies, from discovery to clinical trials. Our final speaker will present the latest developments in workflow environment used by the Swiss Personalized Health Network (SPHN) to support Swiss researchers in exploring the next generation of data-driven health and care innovations.

Semantic Interoperability Challenges for Sharing and Reusing Large Amounts of Heterogeneous Data

Marie-Christine Jaulent (INSERM, France)

Tuesday,
July 3, 2018

13:30 – 14:00
Samarkand Room

Although the Big Data approach seems promising in various analytic uses, sharing or integrating data within the same analysis space remains a complex task, as existing data is highly heterogeneous and difficult to compare. In this presentation, we address the Variety and Veracity dimensions of Big Data when integrating, sharing and reusing large amounts of heterogeneous data for data analysis and decision making applications in the healthcare domain. Many issues are raised by the necessity to conform Big Data to standards in order to make data more interoperable both by humans or computations such as data mining. We discuss how ontologies (computerized meaning) can contribute to the improvement of information sharing and address the problem of data sharing together with semantic interoperability data frameworks.

Challenges of Volume Rendering in a Virtual Reality Environment

Philippe Cattin (University of Basel, Switzerland)

Tuesday,
July 3, 2018

14:00 – 14:30
Samarkand Room

Volume Rendering of medical three-dimensional data is challenging for Virtual Reality due to the computational complexity and the required high frame rate of 90 frames per eye. In this presentation I will show how we achieved this goal using standard GPU hardware and what possible applications of the technology are in the Medical Field.

HPC-Supported Therapy Development in Oncology

Olivier Michielin (University of Lausanne, Switzerland)

Tuesday,
July 3, 2018

14:30 – 15:00
Samarkand Room

Oncology is being revolutionized by technological breakthroughs that permit unprecedented in-depth analysis of the tumour tissue; until now, decision making was based on the analysis of a few well known molecular alterations. Recent technologies are now providing complete interrogation of germline and somatic mutations (genomics), of the gene expression level (transcriptomics), of the resulting protein levels (proteomics), of the metabolic status (metabolomics), of the antigens presented at the surface of tumour cells (immuno-peptidomics), as well as many additional omics to complete this very rich data set. The clinical decision process is therefore expected to heavily rely on computational approaches in the years to come, with machine learning technologies playing a key role. In addition, once key targets are identified through such -omics approaches, computational methods are also key to accelerate the drug discovery process. Two computer-based drug discovery projects will be discussed, both aiming at providing improved immunotherapies for melanoma patients. These projects will illustrate the contribution of high performance computing to the drug design and to the protein design approaches.

Achieving Workflow Interoperability for Personalized Health Research in Switzerland

Thierry Sengstag (Swiss Institute of Bioinformatics, Switzerland)

Tuesday,
July 3, 2018

15:00 – 15:30
Samarkand Room

The Swiss Personalized Health Network initiative (SPHN) aims to accelerate biomedical research by making clinical-care data collected in multiple hospitals available to scientists. The first phase of the SPHN initiative is supported by the BioMedIT project (SIB), which will establish the necessary IT environment to achieve this goal. Traditionally research involving clinical data requires defining procedures which are developed *ad hoc* for each project (e.g. data access control, data transfer processes, security framework, methods to ensure reproducibility of results). Establishing such procedures is highly time consuming and the associated complexity makes them error prone. BioMedIT will establish a unified set of IT standards and services which will greatly mitigate these difficulties. Considering the heterogeneity of the research landscape (hospitals, academic institutions), and the rapid evolution of technologies, it is essential to define interoperability standards which are agnostic to the underlying technologies. In this context, workflow environments are an essential component to facilitate the development of reproducible and portable methods. Analysis of clinical data creates additional constraints regarding security and data access control that must be considered. In this session we will present the current consensus on implementing workflow environments for clinical research, and identify gaps still to be worked on.

MS22

Fostering Software Engineering Best Practice within Research Teams

Organizer(s): ***Mark Abraham*** *(KTH Royal Institute of Technology, Sweden)*
Anshu Dubey *(Argonne National Laboratory, USA)*

Theory and experiment have long been two equal pillars of science, and many would hope to add simulation as a third pillar. However, the challenges for those writing the simulation software are immense. The development team must (a) encompass strong domain expertise, so that the simulations are fit for their purpose; (b) develop the code in a way that can be sustained even without the original authors; and (c) demonstrate to their user communities through testing, benchmarking and documentation that the software will be useful in the hands of researchers, who will not be able to read the code.

In this minisymposium, we will hear from the developers of large community codes about approaches they have adopted to unite the teams of people around the development and maintenance of shared codebases. These will cover not just the programming languages and development tools that have been shown to work well, but also how to encourage adoption of good software engineering techniques by professionals and students of other disciplines, and the career-development needs of the research software engineers who will execute the bulk of the work.

The Evolution of Software Practice in GROMACS: To Suit Both the Laptop and the Exascale

Mark Abraham (KTH Royal Institute of Technology, Sweden)

Tuesday, July 3, 2018

13:30 – 14:00
Singapore Room

Molecular dynamics simulations are now a widely used investigative technique, often complementing or even serving in place of experiments. The implementation in GROMACS is already one of the most frequently used of all codes in HPC, however needs radical changes in computational efficiency to maximize users' scientific quality. Those changes must act at all scales of parallelism, whether a single laptop or the largest supercomputers, so key algorithms have been redesigned to permit implementations that can be tailored to current and emerging architectures. The different implementations require heavy investment in software development process, so that the global development team can deliver their projects in ways that users will trust. In this talk, I will recount some of the changes we have made, describing approaches that have worked, and why. Developers facing similar challenges will learn how they can benefit from these practices.

Software Process for FLASH, a Code Serving Multiple Scientific Communities

Anshu Dubey (Argonne National Laboratory, USA)

Tuesday, July 3, 2018

14:00 – 14:30
Singapore Room

FLASH is a multiphysics multiscale code that has been in existence for nearly two decades. It was originally developed for simulating astrophysical phenomena, however, investment in designing an extensible architecture has resulted in several science communities adding capabilities and adopting FLASH for their use. Challenges of various kinds have occurred at different stages of the code's evolution, ranging from deeply technical such as interoperating heterogeneous solvers, to sociological such as interdisciplinary interactions and building a community. Many software practices adopted by the FLASH team were ahead of their time compared to the broader scientific communities, therefore, tools such as testing harness, or checking compliance with coding standards were built in-house. This presentation will outline the evolution of FLASH's software process and tools development in response to specific challenges. It will also highlight the benefits of early investment in software design in terms of ongoing scientific productivity of the code.

Challenges in Evolving Software for Cryo-Electron Microscopy: From CPUs to GPUs and Back Again

Erik Lindahl (Stockholm University, Sweden)

Tuesday,
July 3, 2018

14:30 – 15:00
Singapore Room

In a few years, cryo-EM has gone from being the ugly duckling of structural biology to one of the hottest techniques in science, recently recognized by the 2017 Chemistry Nobel Prize. Modern cryo-EM is entirely dependent on advanced computational tools to reconstruct three-dimensional structures from millions of extremely noisy two-dimensional images, and with faster detectors and more advanced processing algorithms the computational step has become a critical bottleneck - some experimental facilities have clusters with tens of thousands of CPUs. Here, I will present how we have managed to reformulate the Bayesian REgularized LIkelihood OptimizatoN algorithm used in the RELION code into data-parallel algorithms that made it possible to move the dominant parts to GPU accelerators using CUDA. I will also describe the work necessary to reformulate these algorithms to benefit from GPUs, general challenges when implementing CUDA parts in large production codes, and show examples of how GPU-specific features such as texture units enabled exceptional performance acceleration. However, all processors benefit from data-parallel algorithms: By porting our CUDA implementations back to C++ with threading and fast math libraries we have now also achieved tremendous speedup on standard x86 processors, with the compiler generating SIMD code instead of manually introducing hardware-specific instructions.

More than Top-Down or Bottom-Up: Fostering Software Engineering Best Practice in Diverse Groups

Neil Chue Hong (University of Edinburgh, UK)

Tuesday,
July 3, 2018

15:00 – 15:30
Singapore Room

We often think of practice being driven by grass-roots change (encouraging participation) or clear leadership (providing direction). In reality, successful application of best practice requires both, particularly in the diverse communities where research software is developed and used. Over the last five years, significant changes related to research software have taken place from training initiatives like Software Carpentry, to the recognition of the role of the Research Software Engineer, to the development and adoption of community guidelines and practices. My talk will look at how each of these is related, and how the drive towards reproducibility, FAIR research outputs, and open science is having an effect on the way that software development is being done in research teams.

MS23

High Performance Graph Algorithms

Organizer(s): ***Olaf Schenk*** *(Università della Svizzera italiana, Switzerland)*
Gerhard Wellein, ***Georg Hager***
(Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany)

Graphs (or networks) are a very powerful abstraction of various phenomena that can be expressed as a relation between entities. For several decades researchers in theoretical computer science and discrete mathematics have been developing a wealth of graph theory and graph algorithms. Recently, however, we see a qualitative change in how graph algorithms are used in practice: The complex structure of graphs in new and emerging applications, the size of typical inputs, and the computer architectures on which graph problems are solved call for novel algorithms and hardware efficient approaches. As computer architectures and memory hierarchies are becoming more complex, increasingly parallel and heterogeneous it is important to develop parallel algorithms and tools with these specific hardware constraints in mind. The minisymposium thus aims to bring together experts in developing, implementing and using modern graph algorithms to address this issue. Existing algorithms and tools will be reviewed in terms of modern HPC architectures and novel hardware efficient approaches will be presented. The application areas for such techniques include sparse matrix partitioning/coloring, and network graph analysis that are traditionally sequential as well as applications that need the extreme performance of emerging hardware architectures.

Tracking Communities in Streaming Graphs

David Bader (Georgia Institute of Technology, USA)

Tuesday,
July 3, 2018

13:30 – 14:00
Sydney Room

A variety of massive datasets, such as social networks and biological data, are represented as graphs that reveal underlying connections, trends, and anomalies. Community detection is the task of discovering dense groups of vertices in a graph. Its one specific form is seed set expansion, which finds the best local community for a given set of seed vertices. Greedy, agglomerative algorithms, which are commonly used in seed set expansion, have been previously designed only for a static, unchanging graph. However, in many applications, new data are constantly produced, and vertices and edges are inserted and removed from a graph. We present an algorithm for dynamic seed set expansion, which maintains a local community over time by incrementally updating as the underlying graph changes. We show that our dynamic algorithm outputs high-quality communities that are similar to those found when using a standard static algorithm. It works well both when beginning with an already existing graph and in the fully streaming case when starting with no data. The dynamic approach is also faster than re-computation when low latency updates are needed.

Parallel Mesh Partitioning with Balanced K-Means

Moritz von Looz (University of Cologne, Germany)

Tuesday,
July 3, 2018

14:00 – 14:30
Sydney Room

Graph partitioning is an indispensable tool for efficient matrix and graph processing in distributed memory, balancing the computational load while minimizing communication. The required methods largely depend on the graph type: Numerical simulation meshes mostly have homogeneous degrees, high diameter and often spatial information, enabling geometric approaches. Complex networks have a low diameter, heterogeneous degrees and no useful spatial information. However, even for numerical simulation meshes, purely geometric approaches often suffer from unsatisfactory solution quality. We discuss two graph partitioners addressing these challenges: (i) ParHIP (Meyerhenke, Sanders, and Schulz), the parallel version of KaHIP, a graph partitioner for complex networks and meshes. In a multilevel process, it performs coarsening and local refinement based on size-constrained label propagation. As an example, using 512 cores, the resulting algorithm produces a high-quality partition of a web graph with 3.3G edges in 16 seconds; (ii) Geographer, the main focus of this presentation, is a new approach for mesh partitioning combining space-filling curves, balanced k-means and combinatorial local refinement. In experiments with meshes on up to 16384 processes, it scales well and relevant quality measures are often better than with ParHIP and ParMeTiS. The core of Geographer is a scalable version of k-means adapted to yield balanced clusters.

Co-Author(s): **Charilaos Tzovas**, **Henning Meyerhenke** (University of Cologne, Germany)

Improvement of Graph Partitions Using the Graph p-Laplacian

Drosos Kourounis (Università della Svizzera italiana, Switzerland)

Tuesday,
July 3, 2018

14:30 – 15:00
Sydney Room

A continuous formulation of the optimal 2-way graph partitioning based on the p-norm minimization of the graph Laplacian Rayleigh quotient is presented, which provides a sharp approximation to the balanced graph partitioning problem, the optimality of which is known to be NP-hard. The minimization is initialized from a cut provided by a state-of-the-art multilevel recursive bisection algorithm, and then a continuation approach reduces the p-norm from a 2-norm towards a 1-norm, employing for each value of p a feasibility-preserving steepest-descent method that converges on the *p*-Laplacian eigenvector. A filter favors iterates advancing towards minimum edge-cut and partition load imbalance. The complexity of the suggested approach is linear in graph edges. The simplicity of the steepest-descent algorithm renders the overall approach highly scalable and efficient in parallel distributed architectures. Parallel implementations of recursive bisection on multi-core CPUs and GPUs are presented for large-scale graphs with up to 1.9 billion tetrahedra. The suggested approach exhibits significant improvements over both METIS and KaHIP for graphs originating from various application domains of graph partitioning, ranging from triangular Delaunay meshes to power networks. Particular emphasis is placed on the benefits of applying the p-Laplacian method on graphs emerging from social networks.

Co-Author(s): **Dimosthenis Pasadakis**, **Olaf Schenk** (Università della Svizzera italiana, Switzerland)

RACE: Recursive Algebraic Coloring Engine

Christie Louis Alappat (Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany)

Tuesday,
July 3, 2018

15:00 – 15:30
Sydney Room

Graph coloring is an important method used to parallelize sparse matrix kernels having inherent data-dependencies. Typical examples range from exact kernels like sparse matrix transpose vector (SpMTV), symmetric sparse matrix vector (SymmSpMV) to iterative solvers like Kaczmarz (KACZ) and Gauss-Seidel (GS). Most of the typical schemes currently available to parallelize such kernels suffer from performance issues on modern hardware or are highly matrix specific or require changes in the entire matrix storage format. We propose a novel method called RACE that achieves high hardware efficiency on modern multi-core architectures and at the same time uses simple storage formats like compressed row storage (CRS). The method used is a recursive level-based method that aims at finding optimal permutations while preserving good data locality. A thorough performance analysis shows that RACE out-performs traditional multi-coloring methods, Intel MKL implementations and even the recursive sparse block (RSB) implementation of SymmSpMV that uses tailored storage format for such operations.

Co-Author(s): **Gerhard Wellein**, **Georg Hager** (Friedrich-Alexander-Universität Erlangen-Nürnberg),
Holger Fehske (University of Greifswald, Germany)

MS24

Plasma I: Exciting Opportunities for Plasma Simulation in the Pre-Exascale Era

Organizer(s): ***Frank Jenko*** *(Max Planck Institute for Plasma Physics, Germany)*

This minisymposium will address exciting opportunities for plasma simulation in the pre-exascale era, addressing many issues which transcend plasma physics and are of interest to a wide audience. Plasma physics offers many opportunities to explore quintessential complex systems characterized by multi-scale and multi-physics problems. The underlying nonlinear integro-differential equations can only be solved with the help of supercomputers. Therefore, not surprisingly, many computational plasma scientists view the upcoming exascale era as the golden age, finally allowing them for the very first time to attack several long-standing fundamental questions, from turbulence to magnetic reconnection to dynamo action - as well as their self-consistent interactions. The goal of the present minisymposium is to present examples of how the computational plasma physics community is preparing for the exascale era. A particularly fascinating aspect of this theme can be put into the formula "big data meets computation." Handling massive amounts of data - before, during, and after a simulation - and combining experimental or observational data with simulation data are two key challenges in this context. Novel ideas along these lines will be presented.

Design and Development of Particle-in-Cell Methods for Emerging Tensor Architectures

Stefano Markidis (KTH Royal Institute of Technology, Sweden)

Tuesday,
July 3, 2018

13:30 – 14:00
Osaka Room

Several companies are developing specialized hardware to boost the performance of dense matrix low-precision computations as the market of AI-based data analytics considerably increased in the last decade. For instance, Google and NVIDIA designed the Tensor Processing Unit (TPU) and Tensor Cores in Volta GPUs [1] respectively. The next pre-exascale machines, such as the Summit and Sierra supercomputers, will feature NVIDIA Volta Tensor Cores. However, it is still unclear how codes for plasma simulations will take advantage of tensor architectures. Widely used massively parallel Particle-in-Cell (PIC) models of plasmas are not yet capable of exploiting these new systems and need to redesigned. Two main aspects have to be considered: first, the PIC algorithms have to be reformulated to use dense matrix multiplications; second, new algorithms have to cope with low-precision calculations, still retaining acceptable accuracy. In this talk, we review the emerging tensor architectures and propose algorithmic changes in PIC codes to exploit tensor hardware. [1] S. Markidis, S.W.D. Chien, E. Laure, I.B. Peng, J.S. Vetter, *NVIDIA Tensor Core Programmability, Performance & Precision*, Accepted for publication in AsHES'18 workshop at IPDPS 2018, 2018.

Co-Author(s): **Chaitanya Prasad**, **Steven Wei Der Chien**, **Erwin Laure** (KTH Royal Institute of Technology, Sweden), **Vyacheslav Olshevsky**, **Giovanni Lapenta** (KU Leuven, Belgium)

Vlasiator – Understanding Near-Earth Space in Six Dimensions

Minna Palmroth (University of Helsinki, Finland)

Tuesday,
July 3, 2018

14:00 – 14:30
Osaka Room

The constant flow of solar wind from our star, the Sun, builds the richest reachable plasma laboratory with spatial and temporal scales not attainable in terrestrial laboratories. Plasma phenomena within the near-Earth space create space weather, referring to harmful effects that can endanger technological systems or human life in space. Space weather predictions are mostly at an empirical stage, while future forecasts will be based on numerical simulations. Up to now, large-scale space weather simulations are based on a very simple theory assuming that plasma is a fluid. Vlasiator is a newly developed large-scale space physics model. Vlasiator modelling targets are immense: to model the entire near-Earth space with a breakthrough resolution, using a description going far beyond the existing large-scale plasma simulations. Therefore, Vlasiator includes advanced high-performance computing techniques available from load-balancing to highly scalable grids to allow massively parallel computations. Due to the unprecedented accuracy at global scales, Vlasiator has been used to discover phenomena that no one thought would exist. The presentation introduces Vlasiator, and some of the recent science results. Future application areas may include space weather, 6D fusion modelling, and spacecraft instrument specification definitions, and as benchmark to test new facilities and architectures.

Co-Author(s): **Urs Ganse**, **Markus Battarbee**, **Brito Thiago**, **Maxime Grandin**, **Yann Pfau-Kempf**, **Lucile Turc** (University of Helsinki, Finland), **Sebastian von Alfthan** (CSC - IT Centre for Science, Finland)

Variable Precision: Making Every Bit Count

Jeffrey A. F. Hittinger (Lawrence Livermore National Laboratory, USA)

Tuesday,
July 3, 2018

14:30 – 15:00
Osaka Room

Decades ago, when memory was a scarce resource, computational scientists routinely worked in single precision and were more sophisticated in dealing with the pitfalls finite-precision arithmetic. Today, however, we typically compute and store results in 64-bit double precision by default even when very few significant digits are required. Many of these bits are representing errors instead of useful information. This over-allocation of resources is wasteful; we communicate and compute on many meaningless bits. At LLNL, we are developing the methods and tools that will enable the routine use of dynamically adjustable precision at a per-bit level depending on the needs of the task at hand. Our goal is to provide more or less precision as needed locally. Acceptance from the community requires that we address three concerns: that we can ensure accuracy, ensure efficiency, and ensure ease of use in development, debugging, and application. In this talk, I will discuss the benefits and the challenges of variable precision computing, highlighting aspects of our ongoing research in data representations, numerical algorithms, and testing and development tools. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Towards a Virtual Fusion Facility on Exascale Supercomputers

Frank Jenko (Max Planck Institute for Plasma Physics, Germany)

Tuesday,
July 3, 2018

15:00 – 15:30
Osaka Room

Building on remarkable advances throughout the last two decades or so, and in expectation of the emergence of exascale supercomputers, the fusion theory community has started to target its ultimate goal: the development of a validated predictive capability aka a virtual fusion facility. This will constitute a milestone in fusion research, providing countless novel opportunities for optimizing the design and operation of fusion experiments and for accelerating the development of fusion energy. One crucial step in this direction is the creation of a backbone for a virtual plasma via the tight coupling of two scalable, cutting-edge gyrokinetic codes, GENE and XGC, addressing the physics in the core and boundary region, respectively. This work is carried out within the U.S. Exascale Computing Project. We will present mathematical and computational aspects of the coupling scheme as well as initial simulations results. Moreover, we will discuss general lessons learned in this ambitious attempt to couple two ab initio codes on pre-exascale machines, which are likely to be relevant to similar efforts in other areas of computational science.

MS25

Scientific Computing in times of MPI+X: Looking at Multiple "X" with regard to Performance and Portability

Organizer(s): ***Sunita Chandrasekaran*** *(University of Delaware, USA)*

Heterogeneity has been very evident among HPC systems and the trend only continues to rapidly evolve. Such a trend involves systems equipped with hierarchical processors along with accelerators and memory/storage components that is expected to facilitate migration of an increasingly diverse set of scientific applications thus meeting the demands of a wide user community. However in order to do so, effectively, we need rich programming models and languages that can tap into the massive potential of these large scale systems. This minisymposium addresses how the widely popular parallel programming paradigms such as CUDA, OpenMP 4.5, OpenACC and Alpaka can be adapted for a variety of applications such as atomistic simulation, turbulence combustion, simulation of smoke propagation, and plasma physics. The talks will explain using these real world applications how to balance performance and portability for a minimum of future work.

Porting Physical Parameterizations from a Climate Model to Accelerators

Thomas Köster (Università della Svizzera italiana, Switzerland),
William Sawyer (ETH Zurich / CSCS, Switzerland)

Tuesday, July 3, 2018

13:30 – 14:00
Nairobi Room

ICON (ICOsahedral Non-hydrostatic) is a climate and numerical weather prediction model developed by the Max Planck Institute for Meteorology (MPI-M) and the German Weather Service (DWD). Together with MPI-M and DWD, MeteoSwiss, the Center for Climate Systems Modeling (C2SM/ETH), and the Swiss National Supercomputing Center are porting ICON to GPUs and many-core architectures. Within the model, physical parameterizations calculate the collective effect of physical phenomena which occur on a sub-grid scale. We suggest multiple directive-based approaches of porting these parameterizations to accelerators, such as OpenACC and the CLAW source-to-source translator. These approaches allow the retention of a single Fortran code, which offers a high degree of performance portability. Using the FortranTestGenerator tool for automatic unit test generation, the turbulence parameterization is isolated in a testbed subset of the model, so that subsequent changes can be easily validated. We describe the challenges of porting CPU cache-bound programs to GPUs. Tool-based analysis of loop kernels is used to estimate attainable performance on various platforms, including x86-based CPUs and GPUs. The validated turbulence parameterization can be integrated into the overall ICON model.

Co-Author(s): **Olaf Schenk** (Università della Svizzera italiana, Switzerland), **Gerhard Wellein** (Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany), **Xavier Lapillonne** (MeteoSwiss, Switzerland)

Zero Overhead Modern C++ for Mapping to Any Programming Model

Axel Huebl (Helmholtz-Zentrum Dresden-Rossendorf, Germany)

Tuesday, July 3, 2018

14:00 – 14:30
Nairobi Room

Modern HPC systems are increasingly heterogeneous and diverse: counting hosts and accelerators, the top 10 supercomputers in 11/2017 alone comprised as many as 11 distinct architectures. Moreover, accompanying programming models can range from directive based, implicit and explicit descriptions up to task-based. With limited development resources but often multi-decade long project lifetimes, maintaining multiple implementations of the same algorithms to widen platform support is unfeasible for most development teams. Alpaka is a standard C++, compile-time meta-programming library providing a unified, explicit, parallel programming model. On typical MPI+X parallelized applications, Alpaka enables developers to describe shared-memory, in-node parallelism. Zero-overhead abstraction is achieved by compile-time specializing C++ templates to native backends (e.g. CUDA, OpenMP, and TBB). Alpaka stays with modern C++ as a standardized, widely supported language without introducing pre-processor or pragma-based annotations to the user directly. It naturally allows inlining, kernel fusion and code-reuse on a single-source programming paradigm. With such, abstractions and control within the final software stack are achievable without duplicating implementations leading to a maintainable code base even for large applications.

Co-Author(s): **Benjamin Worpitz**, **Erik Zenker** (LogMeIn Inc., Germany), **Alexander Matthes**, **René Widera**, **Guido Juckeland**, **Michael Bussmann** (Helmholtz-Zentrum Dresden-Rossendorf, Germany)

Porting Quantum ESPRESSO to GPUs - Lessons Learnt and Remaining Challenges

Pietro Bonfà (CINECA, Italy)

Tuesday, July 3, 2018

14:30 – 15:00
Nairobi Room

Quantum ESPRESSO is a very popular open-source suite of codes for electronic-structure calculations and materials modelling at the nanoscale. PWscf package is the main package and focus of this talk. The aim of this talk is to present the multi-year journey to allow Quantum ESPRESSO to exploit NVIDIA GPU accelerators. The first GPU porting was done in CUDA C back in 2012, and a new version based on CUDA Fortran and CUF kernel was developed and released during the last year by Nvidia and Filippo Spiga. Based on these two experiences, a new effort to embed specific accelerated kernels in the main repository is ongoing. We will present the envisaged strategy for the integration of the accelerated code and discuss the lessons learnt developing and maintaining a continuously evolving complex Fortran community code. The talk will also review how the porting is done, various tricks to integrate in the same code base both the CPU and the GPU code path, integration with libraries and the comparison between CUF kernels and OpenACC directives for selected kernels.

Co-Author(s): **Fabio Affinito**, **Carlo Cavazzoni** (CINECA, Italy)

OpenMP 4.5 Acceleration for Turbulence Simulations on GPUs

Dhawal Buaria (Max Planck Institute for Dynamics and Self Organization, Germany)

Tuesday, July 3, 2018

15:00 – 15:30
Nairobi Room

Optimal use of GPUs in a heterogeneous computing environment requires careful consideration of data movement and (usually) how operations on the CPU and attached GPU(s) can be overlapped as much as possible. Further complexities also arise if the base algorithm requires substantial communication among a large number of MPI parallel processes. We have devised a successful OpenMP 4.5 application that solves 3D advection-diffusion equations for mixing of a scalar (concentration) field transported in a turbulent fluid flow, with a 5X GPU-to-CPU speedup on 8192 nodes of the Cray XK7 Titan machine at Oak Ridge National Laboratory, USA. To minimize data movements between CPU and GPU memory spaces, the entire memory space required for the scalar field is transferred to the GPUs, where most of the computations are performed. Computational loops that perform compact finite difference calculations are accelerated using OpenMP 4.X constructs, where in some instances a change from the default memory layout is beneficial. Scalability is improved by overlapping computation on the GPUs with (i) communication on the CPUs and (ii) data movement between the CPUs and GPUs using the latest tasking capabilities added to the TARGET constructs in OpenMP 4.5 (e.g., the DEPEND and NOWAIT clauses).

Co-Author(s): **Matthew P. Clay** (Air Force Research Laboratory, USA), **P. K. Yeung** (Georgia Institute of Technology, USA)

MS26

Tensor Algebra Computation: Implementations and Applications

Organizer(s): ***Alfio Lazzaro***, ***Juerg Hutter*** *(University of Zurich, Switzerland)*
Edgar Solomonik *(University of Illinois Urbana-Champaign, USA)*

Tensor algebra operations are ubiquitous in domains including data analytics, machine learning, engineering, and science. Moreover, the use of tensor methods in these domains has grown tremendously in the last ten years. High-performance implementations and parallelization of tensor algebra operations underlying these methods require considerations beyond standard techniques used in linear algebra. A further key challenge is the development of effective abstractions and libraries for tensor algebra, as no standard interface or library (like LAPACK) has been established in the scientific community. The presentations in this minisymposium will cover challenges faced in the development of four distinct major tensor software library efforts. These libraries focus on sparse and dense tensor contractions, with three targeting distributed memory: Cyclops, NWChem, and DBCSR, and a fourth, TACO, targeting shared memory. These efforts have been application-driven, in particular by the tensor problems prevalent in electronic structure calculations. The minisymposium will serve as a discussion on the current state of the art of the tensor libraries in order to understand possible synergy between the projects and with the possibility to build a common interface for tensor algebra operations.

Parallel Tensor Computations in Python or C++ Using Cyclops

Edgar Solomonik (University of Illinois Urbana-Champaign, USA)

Tuesday,
July 3, 2018

13:30 – 14:00
Boston 3 Room

Tensor algebra provides a mathematical language prevalent in numerous domains of computing, including computational chemistry, machine learning, and quantum information. The Cyclops library provides high-performance algorithms for fundamental operations (summation, contraction, factorization, slicing, reshaping, etc.) on dense or sparse tensors. The library uses distributed storage via MPI and executes each tensor operation bulk synchronously. Algebraic tensor operations are specified via a high-level Einstein summation syntax, accessible as a standalone library in C++ or Python. We will describe the methods Cyclops uses to achieve high-performance, including communication-avoiding algorithms for (sparse) matrix multiplication, optimized transposition/redistribution routines, and runtime mapping/algorithm selection based on trainable performance models. The library has been used to break computational frontiers in electronic structure calculations and quantum circuit simulation, with further efforts on Cyclops applications for graph analysis, neural networks, and multilevel optimization underway.

Tensor Transposition and Contraction on GPUs

Ponnuswamy Sadayappan (Ohio State University, USA)

Tuesday,
July 3, 2018

14:00 – 14:30
Boston 3 Room

A large variety of contractions involving tensors of different dimensionalities and index combinations represent the most computationally demanding components for several models in the NWChem computational chemistry suite. We discuss a library for implementing arbitrary dense tensor contractions on GPUs, using a lower level library for tensor transposition. For a given tensor contraction, there often exist many different choices of intermediate transposed tensors that enable efficient vendor libraries for matrix multiplication (eg., cuBLAS) to be used to perform the tensor contraction. Performance models are used to enable choice between alternatives. The effectiveness of the library on tensor contractions for the CCSD(T) coupled cluster method will be discussed.

Co-Author(s): **Jinsung Kim**, **Aravind Sukumaran-Rajam** (Ohio State University, USA),
Sriram Krishnamoorthy (Pacific Northwest National Laboratory, USA)

Extending the DBCSR Library to Sparse Tensor Linear Algebra for Electronic Structure Methods beyond Density Functional Theory

Alfio Lazzaro (University of Zurich, Switzerland)

Tuesday,
July 3, 2018

14:30 – 15:00
Boston 3 Room

Advanced algorithms for large-scale electronic structure calculations are mostly based on processing multi-dimensional sparse data. Examples are sparse matrix-matrix multiplications in linear-scaling Kohn-Sham calculations or the efficient determination of the exact exchange energy. When going beyond mean field approaches, e.g. for Moller-Plesset perturbation theory, RPA and Coupled Cluster methods, or the GW methods, it becomes necessary to manipulate higher-order sparse tensors. Very similar problems are also encountered in other domains, like signal processing, data mining, computer vision, and machine learning. Our project is concerned with the development of such a tensor library. The starting point of the project is the realization that most tensor operations can be mapped on matrix multiplications. We can therefore base the development on the already existing domain library DBCSR, a distributed block sparse matrix multiplication library. DBCSR has a multi-layered structure that automatically takes care of and optimizes several computational aspects like parallelism (MPI, OpenMP, CUDA), data (cache) locality and on-the-fly filtering. In this presentation, we describe the status of the library development, the implemented functionalities, and the API in Fortran and C/C++. Then, we report on a comparison with other libraries available in the community, in terms of functionalities and performance.

Co-Author(s): **Juerg Hutter**, **Patrick Seewald**, **Ilia Sivkov** (University of Zurich, Switzerland)

The Tensor Algebra Compiler

Saman Amarasinghe (Massachusetts Institute of Technology, USA)

Tuesday,
July 3, 2018

15:00 – 15:30
Boston 3 Room

The Tensor Algebra Compiler (taco) automatically generates kernels to compute tensor and linear algebra expressions on both dense and sparse data. This frees application and library developers from hand-coding these kernels. The generated sparse kernels have excellent performance and match hand-coded kernels where these are available, while generalizing to an uncountable number of other kernels. Ref: www.tensor-compiler.org.

Co-Author(s): **Fredrik Kjolstad**, **Stephen Chou** (Massachusetts Institute of Technology, USA), **Shoaib Kamil** (Adobe Research, USA), **David Lugato** (CEA, France)

MS27

Actionable Health Intelligence: From Precision Medicine to Population Health

Organizer(s): ***Georgia Tourassi*** *(Oak Ridge National Laboratory, USA)*

The healthcare sector is clearly experiencing a data revolution. With advances in digital health records, genomic sequencing, burgeoning growth of social networks and media for community-health, and the emerging "App" market for health-related mobile and web-enabled applications – there is tremendous access and availability of private and public data for advancing precision medicine and improving population health. This ability to leverage these datasets for translational value, across the continuum of basic, preclinical, and clinical science will be critical for addressing in an effective and timely manner emerging personalized and population healthcare challenges. At the same time, artificial intelligence is making continuing advances in biomedicine. However, there are outstanding questions of how AI can provide actionable clinical insights. We will bring together a community of biomedical researchers and computer scientists to present the latest advances as well as discuss successes, challenges, and next frontiers in health intelligence. The overarching goal of the symposium is to highlight health informatics applications and related methodological advances that involve heterogeneous biomedical data (e.g., imaging, genomic, text, and sensor data) while emphasizing the current and emerging challenges of ensuring their translational value and broad population impact.

Radiogenomics in the Era of Precision Medicine

Constantinos Pattichis (University of Cyprus, Cyprus)

Tuesday,
July 3, 2018

16:00 – 16:30
Sydney Room

Precision medicine aspires to leverage new knowledge emanating from heterogeneous genomic, environmental, and clinical data analysis, facilitating increased understanding of disease progression, treatment efficacy, and prevention, towards developing new, personalized therapies and interventions. Transforming clinical practice for precision medicine dictates fundamental advances that range from new big data analytics tools development and research groups formation, to standardization of acquisition and sharing of research data and electronic health records (EHR). This talk highlights the potential of exploiting radiogenomics approaches in the era of precision medicine. It provides a brief overview of selected studies that have been conducted under the precision medicine initiative in the U.S.A, exploiting the joint processing of imaging, genomics, and clinical data. More precisely, quantitative imaging studies that have been performed for breast invasive carcinoma and glioblastoma, describing key technologies and outcomes, while highlighting future directions will be presented. Acknowledgement: This study is partially funded by the H2020-WIDESPREAD-04-2017-Teaming Phase 1, Grant Agreement 763781, *Integrated Precision Medicine Technologies (IPMT)*.

Co-Author(s): **Costas Pitris**, **Andreas Panayides** (University of Cyprus, Cyprus)

Deep Multi-Omics to Predict Clinical Cancer Phenotypes

Georgia Tourassi (Oak Ridge National Laboratory, USA)

Tuesday,
July 3, 2018

16:30 – 17:00
Sydney Room

Computational phenotyping is expected to play a critical role in advancing precision medicine and improving healthcare quality. Such endeavor relies on integration of heterogeneous data from different modalities such as electronic health records, imaging, and genomics. There are several existing efforts leveraging statistical and conventional machine learning techniques to model relationships between multi-omics data and clinical outcomes. Recently deep learning has emerged as a powerful alternative bypassing the laborious curation and feature engineering steps of conventional techniques. In precision oncology, the integration of radiomics and genomics is highly promising for predictive modeling of cancer risks and outcomes and for data-driven discovery of high precision phenotypes to support genome-phenome association studies. To date radiogenomics studies have mainly focused on investigating correlations between genomic and radiomic features, or selecting salient features to determine clinical tumor phenotypes. In this presentation, I will discuss our efforts to apply scalable deep learning on radiogenomic data to predict clinical phenotypes of invasive breast cancer. Based on 3D full MRI imaging data from The Cancer Imaging Archive and gene expression data from The Cancer Genome Atlas, we show that deep radiogenomics produces better prediction of breast cancer pathological state and molecular receptor status.

Co-Author(s): **Hong-Jun Yoon** (Oak Ridge National Laboratory, USA)

Explainable-AI: From Human Systems Biology to the 3D Interactome and Precision Medicine

Daniel Jacobson (Oak Ridge National Laboratory, USA)

Tuesday,
July 3, 2018

17:00 – 17:30
Sydney Room

This project takes place in the context of our efforts to understand human biology as a complex system. We are using 1000 fully re-sequenced human genomes to calculate co-evolutionary relationships between genes. In addition, we are using publicly available genome wide association results, gene expression, protein-protein interaction and human interactome databases combined with deep learning algorithms in order to predict protein complexes. Publicly available protein structures of human proteins (as determined by X-Ray crystallography), as well as predicted proteins structures (as determined by homology modeling or *ab-initio* prediction), are being used in combination with co-evolution information and molecular dynamics for *in-silico* protein-protein docking in order to predict the structures of protein complexes responsible for cellular functions that lead to organismal phenotypes. Small molecule docking of compounds will be performed *in-silico* against the structures of individual proteins as well as the assembled protein complexes resulting from protein-protein docking. Drug interaction information extracted from the DrugBank database as well as those inferred via explainable-AI from the Veterans Administration polypharmacy data are being layered on top of the small-molecule to protein/protein complex docking results in order to better understand the 3D interactions responsible for drug interactions.

Drug Response Prediction in Cancer Cell Lines and Patient-Derived Xenografts

Fangfang Xia (Argonne National Laboratory, USA)

Tuesday,
July 3, 2018

17:30 – 18:00
Sydney Room

Predicting tumor cell response to drug treatments is a critical challenge for accomplishing the promise of precision medicine in oncology. As part of the joint project between DOE and NCI to develop advanced computing solutions for cancer, we are developing a deep learning based framework for modeling tumor-drug interaction and predicting dose response in pre-clinical screening. We are integrating information across thousands of cell lines, tens of thousands of compounds, NCI's in vitro drug screen results amassed over decades, and new experimental biological data derived from patient-derived xenografts. In this talk, we will present our dose response predictions on single and paired drugs. We will discuss the challenges in working with biological data and our preliminary work on applying semi-supervised methods to learn representations of various molecular assay types and the drug compound space. We will also present our work on combining uncertainty quantification with parallel inference runs to guide new biological experiments.

MS28

Advances in Automation and Efficiency for the Exascale Era – Experiences from the Biomolecular Sciences

Organizer(s): ***Rossen Apostolov*** *(KTH Royal Institute of Technology, Sweden)*

Life Sciences have become crucially dependent on software for analysis of experimental data, systems modelling and simulation, data integration across various repositories and databases, etc. The dramatic increase of available tools has enabled scientists to perform ever more complex studies while taking advantage of modern high-end (HPC and HTC) compute facilities. Experimental facilities are producing staggering amounts of data which led to the rapid development of novel data analytics and machine learning techniques. We are at a stage where there is a need for additional focus on improving the interoperability of software applications, enabling better coupling of tools with data sources and devising efficient workflows and libraries for the upcoming Exascale era in HPC. Such advances will considerably improve the productivity of researchers and allow them to address novel scientific problems. This minisymposium brings together invited experts from two leading institutions in the field – BioExcel, the European Center of Excellence for Computational Biomolecular Research (www.bioexcel.eu), and MolSSI, Molecular Sciences Software Institute (www.molssi.org) in the US to discuss advances in this important field. After the talks we will hold a round-table discussion on "Simulations at Exascale - myth or reality?".

Building Blocks for Adaptive Workflows

Shantenu Jha (Rutgers University, USA)

Tuesday,
July 3, 2018

16:00 – 16:30
Samarkand Room

Next-generation exascale systems will fundamentally expand the reach of biomolecular simulations and the resulting scientific insight, enabling the simulation of larger biological systems (weak scaling), longer timescales (strong scaling), more complex molecular interactions, and robust uncertainty quantification (more accurate sampling). Solving biological problems that require longer timescales, involve more complex interactions and robust uncertainty quantification will require significant algorithmic improvements that incorporate high-level parallelism and leverage the statistical nature of molecular processes. Interestingly, many such simulation algorithms require adaptive workflows. We argue the need for workflow-systems using a building blocks approach to support adaptive workflows on extreme-scale heterogeneous and dynamic resources. We discuss RADICAL-Cybertools as an implementation of the building block concept, and discuss how RADICAL-Cybertools are being used to support a wide range of adaptive workflows in biomolecular simulations.

Facing Compute Platform Portability Challenges with Scientific Workflows - Experiences from Common Workflow Language

Stian Soiland-Reyes (University of Manchester, UK)

Tuesday,
July 3, 2018

16:30 – 17:00
Samarkand Room

Scientific Workflow systems are well established for computational analysis in all science domains, following the rapid development of workflow technology and community practices spanning the two recent decades, the *eScience era*. Workflow systems have gained traction in the era of Big Data Science due to their "ASAP properties": Automation over repetitive pipelines and simulation sweep campaigns; Scaling over computational infrastructure and handling large data; Abstraction to shield users and programs from complexity and incompatibilities; and Provenance to auto-document execution logs and data lineage for future analysis. A major hindrance for wider adaptation and reuse of workflows, even when open source, is that they are written for specific workflow systems or infrastructures. *Common Workflow Language* (CWL) has emerged as a community initiative with support across a range of existing workflow engines, using a language specification that focus on the common denominator of command line tools exchanging files. Support for CWL on HPC expanded in the recent months, such as IBM's CWLEXEC on LSF, or Toil with Singularity. In this talk we will present the challenges of moving CWL workflows towards Exascale, while retaining key features of workflows such as *reproducibility, interoperability, usability* and *provenance*.

Workflow Automation and Efficiency for Macromolecular Simulations and Screening

Adam Hospital Gasch (Institute for Research in Biomedicine, Spain)

Tuesday,
July 3, 2018

17:00 – 17:30
Samarkand Room

Life science is one of the largest and fastest growing communities in terms of needs for high-end computing. Biological studies usually require an integration of different computational approaches, defining complex, automated multi-step analysis workflows with inter-dependent steps, including CPU-intensive tasks generating large amounts of data. This number and diversity of tasks to be integrated, together with the short lifetime and fast turnover of computer codes and life sciences-related methods, make standardization of these workflows an extremely challenging task. BioExcel CoE has been working, together with Elixir project, on putting forward a set of best practices to develop, document and describe life sciences workflows, following the FAIR principles: Findability, Accessibility, Interoperability and Reproducibility. Examples of the first workflow prototypes implemented following this approach (Automatic modeling of protein mutations and Virtual Screening), illustrating the benefits of the introduced best practices, will be presented.

Round-Table Discussion: Simulations at Exascale - Myth or Reality?

Rossen Apostolov (KTH Royal Institute of Technology, Sweden)

Tuesday,
July 3, 2018

17:30 – 18:00
Samarkand Room

Exascale supercomputers seem to be around the corner. Producing them will be a real challenge, no doubt, considering issues with processor design, power-consumption and so on but engineers are confident about their delivery within a few years. Life science (and not only) software applications are capable of running at peta-scale in HPC/HTC regime, but are they ready for the next level push? When the Exa-machines come, will there be simulation engines and job dispatchers able to orchestrate billions of cores? Will researchers be able to tackle major scientific problems and deliver amazing discoveries that are unattainable at lower computing scale? How well prepared are the communities? We have invited Prof. Erik Lindahl, Lead Scientist of BioExcel, the European Center of Excellence for Computational Biomolecular Research (www.bioexcel.eu) and Prof. Daniel Crawford, Director of MolSSI, the Molecular Sciences Software Institute (www.molssi.org), USA, together with leading experts in the field (Shantenu Jha, MolSSI) and Adam Hospital, Stian Soiland-Reyes (BioExcel) to address these questions and try to understand what is needed to improve the interoperability of software applications, enable better coupling of tools with data sources, develop efficient libraries and devise user-friendly and extensible workflows/pipelines for the upcoming Exascale era in HPC.

Co-Author(s): **Erik Lindahl** (Stockholm University, Sweden), **Daniel Crawford** (Virginia Tech, USA), **Shantenu Jha** (Rutgers University, USA), **Adam Hospital** (Institute for Research in Biomedicine, Spain), **Stian Soiland-Reyes** (University of Manchester, UK)

MS29

Advances in Computational Geosciences, Part II

Organizer(s): ***Ebru Bozdag*** *(Colorado School of Mines, USA)*
Dimitri Komatitsch *(CNRS, France)*

Recent advances in theory and numerical methods in parallel to the availability of high-quality massive data sets and high-performance computing provide unprecedented opportunities to improve our understanding of Earth's interior and its mechanism. The goal of this session is to bring computational and Earth scientists together to form a platform to discuss the current status, challenges and future directions in computational geosciences highlighting numerical simulations, the state-of-the-art HPC applications and their scientific outcomes. Contributions include, but are not limited to, the areas of earthquake engineering, passive and active-source seismic imaging, geodynamical modelling, magneto-fluid dynamics, etc. in conjunction with computational approaches such as numerical solvers, large-scale workflow, big data, optimisation strategies, etc. on HPC systems.

Simulating the Solid Earth and Planets over Billions of Years: From Magma Oceans to Plate Tectonics to Exoplanets

Paul J. Tackley (ETH Zurich, Switzerland)

Tuesday,
July 3, 2018

16:00 – 16:30
Darwin Room

The coupled system of plate tectonics and convection of the Earth's solid mantle is the driver of geological change on our planet, including continental drift, volcanoes, earthquakes, crustal production, atmospheric degassing and recycling, and cooling of the core, which drives the geodynamo. Modelling this process is challenging due to the wide range of length scales (from faults to continents) and time scales (seconds to billions of years) and the complex rheology of rocks, which exhibit visco-elasto-plastic behaviour with strongly temperature-dependent viscosity varying by orders of magnitude over short length scales. Nevertheless, it is now routine to perform global-scale 3-D spherical simulations that span the 4.5 billion year age of our planet and contain complex effects such as partial melting and crustal production and solid-solid phase transitions. StagYY is one of the leading codes for performing such simulations. It uses a finite-volume discretization on a yin-yang spherical grid. Both built-in geometric multigrid, and the range of solvers available through PETSc, can be used. Here, technical details will be discussed, including recent enhancements made with PASC funding such as use of hybrid (GPU-CPU) architectures. Some recent scientific results published in Science and Nature will then be summarised.

Dynamic Viability of Earthquake Rupture Cascades on Complex Fault Systems

Alice-Agnes Gabriel (Ludwig Maximilian University of Munich, Germany)

Tuesday,
July 3, 2018

16:30 – 17:00
Darwin Room

Puzzling features of earthquake dynamics are inferred from recent well recorded events. A prominent example is the 2016 Mw7.8 Kaikoura, New Zealand earthquake, considered the most complex rupture observed to date and causing surface rupture of at least 21 segments of the Marlborough fault system. High-quality observations suggest a large gap separating surface rupture traces, the possibility of significant slip on the subduction interface, and slow apparent rupture speed. I will present a comprehensive 3D dynamic model of the Kaikoura earthquake unraveling the event's riddles in a physics-based manner. High resolution modeling is enabled by the open-source software SeisSol (www.seissol.org) that couples seismic wave propagation with frictional fault failure and off-fault inelasticity with high-order accuracy in space and time (minimal dispersion errors). SeisSol exploits unstructured tetrahedral meshes to account for complex geometries, e.g. high resolution topography and bathymetry, 3D subsurface structure, and fault networks. The achieved degree of realism and accuracy is enabled by recent computational optimizations targeting strong scalability on many-core CPUs and a ten-fold speedup owing to an efficient local time-stepping algorithm. Understanding the physical conditions that allow rupture cascades will advance our ability to quantify earthquake hazard, especially regarding the possibility of extreme events on real fault networks.

Imaging of the Italian Lithosphere Based on Adjoint Tomography

Emanuele Casarotti (INGV, Italy)

Tuesday,
July 3, 2018

17:00 – 17:30
Darwin Room

During the PRACE project IMAGINE_IT (3D full-wave tomographic IMAGINg of the Entire ITalian lithosphere) we iteratively create a 3D tomographic model of the Italian lithospheric structure. Our goal was to build a new reference 3D seismic velocity model for the region constrained by a large number of observed full seismic waveforms. We used recorded data of dense seismological networks (400 seismic stations) together with extremely efficient numerical techniques and an enormous computational power (36 million core hours) provided by European Tier-0 system CURIE (GENCI, FR). We exploited the powerful combination of a spectral-element method (code SPECFEM3D) and an adjoint method, for tomographic inversion and imaging based on misfit reduction between observed data (associated to 163 regional earthquakes) and synthetic full waveforms. We performed 25 tomographic iterations, moment tensor inversions and some point spread function resolution analysis. We are able to constrain Vp and in particular Vs at unprecedented resolution and interesting structural and tectonic features start to be accurately modelled. Creating a refined geological model of the lithosphere in Italy will enhance the capability of analysing seismic effects. This has consequences for the assessment of seismic hazard and for planning effective measures based on rapid scenarios.

Co-Author(s): **Federica Magnoni** (INGV, Italy), **Dimitri Komatitsch** (CNRS, France), **Jeroen Tromp** (Princeton University, USA)

Full-Waveform Inversion of the Solid Earth from Crust to Core

Ebru Bozdag (Colorado School of Mines, USA)

Tuesday,
July 3, 2018

17:30 – 18:00
Darwin Room

Accurate and high-resolution images of Earth's interior are crucial to improve our understanding of the inner dynamics of our planet. Global adjoint tomography is one of the extreme projects in seismology due to the intense computational requirements and vast amount of data that can potentially be assimilated in inversions. The first-generation global adjoint tomography model, GLAD-M15, was constructed using data from 253 earthquakes with transverse isotropy confined to the upper mantle, inverting crust and mantle simultaneously. We now perform inversions for next-generation global adjoint models with more complete parameterisations including surface-wave azimuthal anisotropy, anelasticity, etc. while increasing the database in complementary inversions. The GPU version of SPECFEM3D_GLOBE is used for forward and adjoint simulations on the Oak Ridge Leadership Computing Facility's Cray XK7 Titan system, a computer with 18,688 GPU accelerators. We will perform 9 s simulations (currently 17 s) on Oak Ridge's next generation supercomputer "Summit". The ultimate aim is to go down to 1 Hz in global simulations to perform whole-Earth inversions including the core and assimilate all available seismic data from all global CMT earthquakes within the magnitude range of 5.5 to 7.0 in the construction of global models.

Co-Author(s): **Matthieu Lefebvre**, **Wenjie Lei**, **Jeroen Tromp**, **Youyi Ruan**, **James Smith** (Princeton University, USA), **Ridvan Orsvuran** (University of Cote d'Azur, France), **Daniel Peter** (King Abdullah University of Science and Technology, Saudi Arabia), **Dimitri Komatitsch** (CNRS, France)

MS30

Efficient Parallel Methods in High-Dimensional Approximation and Beyond

Organizer(s): ***Helmut Harbrecht***, ***Peter Zaspel***
(University of Basel, Switzerland)

The aim of this minisymposium is to discuss research at the intersection of high-dimensional approximation and parallel computing. High-dimensional approximation drives e.g. uncertainty quantification, optimization, machine learning and big data as well as simulations of complex physics models. It is well-known that approximation of functions of growing dimension has the challenge of the curse of dimensionality. Over the last decades, many powerful mathematical tools have been developed to do weaken or overcome this. These include, but are not limited to (quasi) Monte Carlo, multi-level / multi-fidelity techniques, sparse grids, low-rank and tensor product approximations, hierarchical matrices, compressed sensing and meshfree methods.

There is a growing interest to solve high-dimensional approximation problems at large scale. While many of the discussed methods have good or even optimal approximation properties and complexities for larger dimensions, some have been primarily developed for sequential execution. However, to solve large scale approximation problems, it becomes necessary to develop fast, scalable and parallel numerical methods. This minisymposium invites contributions in high-dimensional approximation ranging from initial studies for the use of parallel techniques up to full scale parallel methods that run on large HPC clusters. Speakers will showcase both algorithmic-oriented and application-centered research.

Portable Distributed Sparse Grid Density Estimation for Big Data Clustering

David Pfander (University of Stuttgart, Germany)

Tuesday, July 3, 2018

16:00 – 16:30
Nairobi Room

Gaining knowledge from vast amounts of data is a challenge in science and engineering. While most conventional learning algorithms scale significantly worse than linear in the number of data points, we follow a different approach and discretize the underlying feature space. This way, the algorithmic complexity depends on the number of discretization points invested and is only linear in the number of data points. Using sparse grids, we mitigate the curse of dimensionality that otherwise renders conventional grid-based approaches unfeasible. Altogether, we obtain an algorithm that is well-suited for massive numerical datasets in moderate dimensions. Nevertheless, the big data era requires the use of HPC systems to keep the overall training time within reasonable bounds. We present new parallel and distributed algorithms that scale on HPC clusters and systems, even with heterogeneous hardware, for the example of clustering based on density estimation on sparse grids. The implementations are performance portable and suitable for processors and accelerators of different vendors. Depending on the hardware, we obtain between 33% and 76% of the achievable peak performance on a single node. Furthermore, we show results on Piz Daint in a heterogeneous and distributed setting.

Co-Author(s): **Dirk Pflüger** (University of Stuttgart, Germany)

Scalable Solvers for Meshless Methods on Many-Core Clusters

Peter Zaspel (University of Basel, Switzerland)

Tuesday, July 3, 2018

16:30 – 17:00
Nairobi Room

Our goal is to solve large-scale stochastic collocation problems in a high-order convergent and scaling fashion. To this end, we recently discussed the radial basis function (RBF) kernel-based stochastic collocation method. In this meshless method, the higher-dimensional stochastic space is sampled by (quasi-)Monte Carlo sequences, which are used as centers of radial basis functions in a collocation scheme. Preliminary applications for this uncertainty quantification framework were (elliptic) model problems and incompressible two-phase flows. One challenge of the discussed approach is the solution of a well-structured large dense linear system to compute the quadrature weights. Linear systems of similar type arise in Gaussian process regression and several machine learning approaches. Classical direct factorization techniques to solve the above linear system for a large to huge kernel sample count are barely tractable, even on large parallel computers. Therefore, we discuss iterative approaches to solve such linear systems on large parallel many-core clusters by hierarchical matrices. The presentation will cover the latest results with respect to numerical methods and applications. This work is partly based on joint work with Michael Griebel, Helmut Harbrecht and Christian Rieger.

Inducing Input and Hyperparameter Optimization for Large Scale Sparse Gaussian Process Regression

Jannik Schürg (University of Bonn, Germany)

Tuesday, July 3, 2018

17:00 – 17:30
Nairobi Room

This talk presents how hyperparameters and inducing inputs as used in Sparse Gaussian Processes can be optimized in a parallel and distributed fashion. Further, empirical results are shown that a good choice through this method can lead to significant predictive improvement. Classic Gaussian Process Regression is a common method for regression but it is limited to small datasets due to its cubic runtime in the size of the input. Several techniques have been proposed in the past to overcome this. One class of methods known as Sparse Gaussian Processes is using so called inducing inputs to obtain a low-rank representation. The talk presents how one of the recent Sparse Gaussian Process methods can be adapted, in order to tune the performance with optimized inducing inputs. More precisely, an algorithm is presented to compute the gradient of the variational bound for the a priori likelihood w.r.t. inducing inputs. An implementation was tested on a cluster with up to 2048 Xeon cores, and on the many-integrated-cores architecture Xeon Phi.

A Highly Scalable, Fault-Tolerant Implementation of the Sparse Grid Combination Technique

Michael Obersteiner (TU Munich, Germany)

Tuesday, July 3, 2018

17:30 – 18:00
Nairobi Room

In this talk we give an overview on the current advancements in the ExaHD project. We focus on the implementational aspects of the Sparse Grid Combination Technique for massively parallel simulations and the integration into a fault-tolerant framework. The Combination Technique introduces a second level of parallelism, replacing a single simulation run by several independent runs. A crucial step of the method is the combination of the single grids which results in a sparse grid representation of the result. Due to its hierarchical nature, the Sparse Grid Combination Technique provides a unique approach to realize algorithm-based fault tolerance without the need of checkpoint-restart by calculating alternative combination schemes that exclude failed resources. A key aspect with respect to scalability is the efficient implementation of the combination step which is crucial to achieve good parallel performance for exascale simulations. Several optimizations based on a manager-worker approach and an efficient communication scheme demonstrate excellent scaling on up to 180k cores. Convergence and scaling results are shown for the application code GENE which solves the gyrokinetic equations to simulate hot plasma in a fusion reactor.

Co-Author(s): **Hans-Joachim Bungartz** (TU Munich, Germany), **Dirk Pflüger** (University of Stuttgart, Germany)

MS31

How Can We Escape the Data Avalanche in Climate Science?

Organizer(s): ***Joachim Biercamp*** *(German Climate Computing Centre, Germany)*
Oliver Fuhrer *(MeteoSwiss, Switzerland)*
Christoph Schär *(ETH Zurich, Switzerland)*

Increasing ensemble size, grid resolution and complexity of Earth system models have driven data volumes to a point where scientists are forced to give some of it up. Estimates for total data volumes of the Climate Model Intercomparison Project 6 (CMIP6, Eyring *et al.* 2016) range from 15 to 30 PBytes. Traditional disk or tape storage bandwidth and capacity is not keeping pace with increases in computational capacity, and data storage and movement constraints are already limiting scientific analysis of these enormous simulation efforts.

With significant momentum towards global convection-resolving climate simulations, traditional compute-store-analyze workflows still common in climate science today will break down and research into alternative data analysis workflows is urgently needed. This minisymposium will highlight and critically discuss the potential and challenges of several pathways to alleviate the data bottleneck in climate science: data compression, in-situ/in-transit data analysis and processing, recomputation.

Beating the Data Bottleneck - Write Less and Use Tiered Storage with Smart Middleware!

Bryan Lawrence (NCAS-CMS, UK)

Tuesday, July 3, 2018

16:00 – 16:30
Rio Room

With increasing resolution and larger ensembles, model output data volumes are becoming prohibitive, particularly for ensembles (where more of the compute is used on data production than on the shorter time steps needed for higher resolution). The situation is exacerbated by users who want data at grid point resolution instead of at the resolution resolved by the mathematics. Existing compression approaches do not play nicely with storage (the phrase "chunk fighting" having been recently coined to describe dealing with compression chunk size and storage-system chunk size mismatches). There is no simple solution to these problems: user education, smarter "cross-ensemble" compression, and better use of tiered storage and smarter workflows are all necessary - but far from sufficient. We describe here some of our observations of real world data handling problems at scale, and some of the tools and approaches being developed by the ESIWACE project to provide part of the solution. We introduce the Earth System Data Middleware and the Semantic Storage Tools, as well as new work on cross-ensemble analysis and compression - in-flight and in post-processing. We will conclude with some of our plans for the future, including dedicated analysis systems and "earth system middleware aware active storage".

Co-Author(s): **Julian Kunkel** (University of Reading, UK)

Lossy Data Compression for Climate Simulation Data: Reducing Data Volume while Preserving Information

Allison H. Baker (National Center for Atmospheric Research, USA)

Tuesday, July 3, 2018

16:30 – 17:00
Rio Room

High-resolution climate model simulations generate enormous data volumes that strain computing center storage resources at institutions such as the National Center for Atmospheric Research (NCAR). Further, storage limitations are negatively impacting science objectives by forcing scientists to run fewer or shorter simulations and/or output data less frequently. Therefore, NCAR has been investigating using data compression to reduce data volumes from the widely used Community Earth System Model (CESM). Striking a balance between meaningfully reducing data volume and preserving the integrity of the simulation data is non-trivial, particularly given the large and diverse set of climate variables. In this talk, we first discuss the challenges of compressing climate data. We then describe our efforts thus far to evaluate the effects of data compression on the original data, which we believe should, at a minimum, not be distinguishable from the natural variability of the climate system. The ultimate goal is that the reconstructed and original climate simulation data are indistinguishable during post-processing analyses, which vary widely according to climate scientists' interests.

Co-Author(s): **Dorit M. Hammerling**, **Haiying Xu**, **John Clyne**, **John Dennis**, **Shaomeng Li** (National Center for Atmospheric Research, USA)

In-Situ to the Rescue?

Jan Frederik Engels (German Climate Computing Centre, Germany)

Tuesday,
July 3, 2018

17:00 – 17:30
Rio Room

The maturing of high-resolution simulation models signifies the end of the classic post-processing/post-visualization workflow. With growing data sizes and simulation complexity, the processes of visualization and analysis become technically more difficult, but at the same time also more important. The human observer needs assistance to perceive and comprehend these large amounts of data, as well as guidance to find the important information within. Fortunately, a large variety of different approaches exist to handle such simulation data. One of which is in-situ visualization, which analyzes and visualizes the data alongside the simulation process. A choice of different setups exist, that allow either a loose or a tight coupling between the simulation model and the visualization software, and as such, the data can be visualized on-the-fly, or images, animations and even geometry stored to disk for later analysis. The biggest benefit of in-situ visualization is also its largest drawback: The data size and complexity is vastly reduced and can not be recovered to its full complexity after the simulation and visualization are finished. This presentation discusses both benefits and drawbacks of the most common in-situ setups and presents such an implementation based on the ICON model using ParaView/Catalyst.

Co-Author(s): **Niklas Röber** (German Climate Computing Centre, Germany)

SimFS: A Simulation Data Virtualizing File System Interface

Salvatore Di Girolamo (ETH Zurich, Switzerland)

Tuesday,
July 3, 2018

17:30 – 18:00
Rio Room

Scientific simulations, such as climate predictions, often produce petabytes of data to be stored in parallel filesystems or large-scale databases. This data is then analyzed, often by thousands of researchers, over the course of decades. However, storing these volumes of data for long time periods is not cost effective and, in some cases, practically impossible. SimFS virtualizes the simulation output: it only stores a small part and the missing data is re-simulated on-demand. SimFS decides which data to store according to the observed analysis access patterns. The data virtualization is invisible to the analysis applications because SimFS intercepts all calls to standard I/O libraries. SimFS enables a trade-off between on-disk solutions, where all the simulation data is stored on disk, and in situ, where no data is stored and analyses are always coupled with simulations. This trade-off is driven by the amount of storage and compute resources assigned to SimFS. Overall, by exploiting the growing computing power and relaxing the storage capacity requirements, SimFS offers a viable path towards exascale simulations.

Co-Author(s): **Pirmin Schmid**, **Torsten Hoefler** (ETH Zurich), **Thomas Schulthess** (ETH Zurich / CSCS)

MS32

Increasing Credibility of Simulation and Analytic Software for Science

Organizer(s): ***Anshu Dubey*** *(Argonne National Laboratory, USA)*
Michael Heroux *(Sandia National Laboratories, USA)*
Mark Abraham *(KTH Royal Institute of Technology, Sweden)*

Recent years have seen a rise in both scrutiny and attention being paid to scientific software because of the corresponding rise in scientific discovery through simulations and data analysis. To produce high quality science it is critical to employ high quality tools and processes. The vast majority of projects engaged in using computations for advancing scientific understanding have yet to attain sufficient robustness in either their software or their process. What has changed is the realization by these projects that they must strive for robustness if they want credibility. However, given the fundamentally multidisciplinary nature of computational science, and scarcity of resources, training and experience, research teams struggle. One of the recently launched community efforts for addressing this gap is the Better Scientific Software website (BSSw.io), a resource for scientific software developers and users which includes curated and contributed content, discussion forums and many other helpful features. Through this minisymposium we seek to introduce the scientific software community to BSSw, highlighting some of the content offered. We also seek to broaden our reach with presentations from practitioners in several scientific communities highlighting what they do and their plans for the sustainability of their software.

Software Engineering for Simulation Neuroscience

Felix Schuermann (EPFL, Switzerland)

Tuesday,
July 3, 2018

16:00 – 16:30
Singapore Room

Compared to other scientific fields, neuroscience only more recently picked up on the capabilities offered by simulation. While this "late start" in principle allows learning sustainable practices from other fields, it also is leaving the field with substantial technical debt. Selected neuro simulator projects are leading good computational science and engineering practices, but for simulation neuroscience an ecosystem of model generation, simulation and analysis software is needed. Big science projects such as the Swiss Blue Brain Project or the European Human Brain Project offer the organizational environment for sustainable science/software co-development. On the example of selected software projects, this talk will illustrate domain specific challenges and recent progress.

Reproducibility in Scientific Software

Michael A. Heroux (Sandia National Laboratories, USA)

Tuesday,
July 3, 2018

16:30 – 17:00
Singapore Room

Reproducibility is at the heart of the scientific method. However, reproducible computational science has historically been a challenge. Difficulties in fully capturing and preserving the computational environment details, challenges of working with floating point arithmetic, and the lack of software practices that enable confident retrieval of specific computational data and software are among the challenges. In addition, our incentive systems do not adequately reward investment in best practices for reproducibility in relation to other demands. In this presentation, we discuss recent changes in expectations for publications, funding and community recognition that provide improved incentives for better reproducibility, and highlight a few of the technical improvements that make reproducible computation more achievable.

Outreach for Better Scientific Software

David E. Bernholdt (Oak Ridge National Laboratory, USA)

Tuesday,
July 3, 2018

17:00 – 17:30
Singapore Room

Historically, there have been many impediments to widespread improvement in the software development practices in the scientific and HPC software communities, a situation which has clear implications for the credibility of the resulting software. Professional reward systems tend to place far more emphasis on scientific results than on software; workshops and journals that welcome discussions of software and the software development experience are still very limited. Researchers are not trained in software engineering, and there is limited and sometimes conflicting information about which software engineering practices are useful in scientific computing and how to implement them in environments which may differ significantly from the typical commercial software settings in which they were typically developed. In this talk, I will discuss some of the strategies the IDEAS Productivity project has been pursuing to try to enhance awareness and bring relevant resources to the community through a broad range of outreach activities. And I hope to engage the audience in discussions of how they can contribute as well.

General Discussion and Community Input

Anshu Dubey (Argonne National Laboratory, USA)

Tuesday,
July 3, 2018

17:30 – 18:00
Singapore Room

We seek to broaden the reach of the BSSw portal by engaging with a broader community. Therefore we will follow the formal presentations with an open discussion session. Given the emergent nature of this field, we believe that gathering input from the community, and brainstorming for ideas, is extremely important.

MS33

Machine Learning Schemes with High Extrapolation Accuracy for Materials Discovery

Organizer(s): ***Stefan Goedecker*** *(University of Basel, Switzerland)*

Materials discovery requires one to explore many possible chemical compositions as well as many possible structures for a given composition. Even though highly powerful computers allow us to perform density functional calculations quite rapidly for a limited number of configurations of moderate size, the huge number of force evaluations required for structural and stochiometric explorations of a very large number of materials is too expensive on the density functional level.

Machine learning schemes may come to our assistance in this context. They have been shown to deliver density functional accuracy at a greatly reduced numerical cost for limited test sets. Since all machine learning schemes are intrinsically interpolation schemes it is however not clear how well these schemes can extrapolate to discover for instance entirely new materials that were not contained in the fitting database.

The minisymposium will therefore focus on problems and systems that are difficult to treat with present day machine learning schemes and present methods that might overcome some of the current limitations.

Structure and Dynamics of Au Nanoclusters Using ANN Based Interatomic Potentials

Satya Bulusu (IIT Indore, India)

Tuesday,
July 3, 2018

16:00 – 16:30
Boston 3 Room

Here we demonstrate that artificial neural network (ANN) based interatomic potentials are accurate in describing interatomic interactions in metal nanoclusters and also make the calculations affordable. In terms of computational speed, ANNs are very fast and hence it allows us to run molecular dynamic simulations up to time scales of 3 ns on a single CPU for a medium sized nanocluster. ANN potentials are explored for bare, doped ((AgAu)55) and thio protected Au nanoclusters to study the effect of size, composition on structure and dynamical properties. For bare Au nanoclusters, molecular dynamics simulations were performed on Au17, Au34, Au58, Au147 and Au309. The study shows that there is a dynamical coexistence of solidlike and liquidlike phases near melting transition. For (AgAu)55, using c-T phase diagram, surface area, surface charge, probability of isomers and Landau free energies, we show enhancement of catalytic property of Ag-Au nanoalloys by incorporation of Ag up to 24% by composition in Au nanoparticles. We show, using ANN, the effect of composition of SH for different sizes of thio protected nanoclusters. UV-visible spectra were utilized to probe the structure of nanoclusters.

Co-Author(s): **Shweta Jindal**, **Siva Chiriki** (IIT Indore, India)

On Creating Databases for Machine Learned Interatomic Potentials

Gabor Csanyi (University of Cambridge, UK)

Tuesday,
July 3, 2018

16:30 – 17:00
Boston 3 Room

The last few years have seen fervent activity in applying machine learning techniques to create interatomic potentials and force fields, with a wide variety of approaches being tried: kernel methods, neural networks, scattering transforms, symmetrised polynomials, etc. Much less attention has been devoted to thinking about the databases that these approximants are fit to. My talk is an attempt to begin to address this, using the case study of a very extensive database of periodic structures of silicon, and a correspondingly extensive suite of benchmark tests that a materials modeller would use to judge the quality of a potential. I will argue that alternatives to Boltzmann sampling to generate such databases will be important in the future.

Materials Modeling Using Neural Networks

Matti Hellström (University of Göttingen, Germany)

Tuesday,
July 3, 2018

17:00 – 17:30
Boston 3 Room

The potential energy surface (PES) of a system lies at the heart of many problems in materials modeling, theoretical condensed matter physics and computational chemistry. Neural networks can be parameterized to provide a computationally inexpensive way of predicting the PES, by training to reference data obtained from quantum mechanical calculations. In this way, neural networks can be used together with molecular simulation techniques to model materials in a realistic environment. Here, the strategies for parameterizing the neural networks, as well as some results from simulations for electrolyte solutions and solid/liquid interfaces, are presented.

Co-Author(s): **Jörg Behler** (University of Göttingen, Germany)

Using Machine Learning Interatomic Potentials for Crystal Structure Prediction

Seyed-Alireza Ghasemi (Institute for Advanced Studies in Basic Sciences, Iran)

Tuesday,
July 3, 2018

17:30 – 18:00
Boston 3 Room

Fueled by the rapid development of computer power, crystal structure predictions using density functional calculations have become viable based on powerful global optimization techniques. Nevertheless, structure predictions are still computationally quite expensive and can become prohibitive for complex systems such as surfaces, large systems, and large crystalline unit cells as well as in cases where various different external parameters such as pressure or particle numbers have to be examined. Due to their reduced numerical cost, recently developed interatomic potentials based on machine learning are able to tackle such problems, provided the interatomic potentials are sufficiently transferable. This transferability is hard to achieve since it is difficult to prepare a set of structures that are diverse enough to probe the entire low energy configuration space during the training process. I will present our approach to overcome these problems as well as results obtained by our machine learning scheme called CENT. It is based on a charge equilibration process where the neural network is used to interpolate an intermediate physical quantity, namely atomic electronegativity. I will present several applications done so far, including TiO2 sheets, ZnO low-density phases, CaF2 and alkali-halide compounds.

MS34

Plasma II: Frontiers in Gyrokinetic Turbulence Simulation on New and Emerging HPC Platforms

Organizer(s): ***Stephan Brunner***, ***Laurent Villard*** *(EPFL, Switzerland)*

Turbulence in plasmas confined by a background magnetic field stand out as one of today's great challenges in physics. To describe such systems, the most appropriate approach is based on gyrokinetic theory, which takes advantage of the small characteristic frequency of turbulence as compared to the cyclotron frequency. It results in a reduction of phase space from 6D to 5D, and eliminates the fastest time scales from the description. Despite this reduction in complexity, solving the set of gyrokinetic equations, consistently with the electromagnetic field equations, remains a formidable task. It is an intrinsically nonlinear, multi-scale problem and thus requires major HPC resources, for which advanced numerical techniques are a necessity in order to obtain a reasonable time-to-solution. In this minisymposium, progress on state-of-the-art codes based on all three major computational approaches (Eulerian, semi-Lagrangian and Lagrangian-PIC) will be presented, as well as possible innovative alternatives.

How to Prepare the Gyrokinetic Code GYSELA for Future Exascale Machines

Virginie Grandgirard (CEA, France)

Tuesday, July 3, 2018

16:00 – 16:30
Osaka Room

Understanding turbulence and heat transport in fusion plasma is a key challenge for fusion devices like ITER. Non-linear 5D gyrokinetic codes are the most accurate framework to tackle this problem. They are extremely HPC challenging. The petascale code GYSELA [V. Grandgirard *et al.*, Comp. Phys. Com. 2017 (2016) 35] is one of them based on a Backward Semi-Lagrangian scheme. During the 2013-2016 period, we obtained good strong scaling on different supercomputers (60% efficiency at 65k cores on Sandy-Bridge machine, 87% on a BlueGene/Q machine at 32k cores). Since then, three main physics developments have been performed: (i) The trapped electrons are no longer assumed adiabatic but kinetic; (ii) a more general collision operator has been implemented to study synergies between neoclassical and turbulent transport; and (iii) GYSELA has been upgraded to model a Scrape-Off-Layer like transition between closed and open field lines. These modifications have had a considerable impact on the numerical cost of the code. We will present here the numerical and parallel optimizations performed to minimize this impact and to increase the global performance of the code to prepare for future exascale simulations.

Co-Author(s): **Yuuichi Asahi**, **Julien Bigot**, **Elisabetta Caschera**, **Guilhem Dif-Pradalier**, **Peter Donnel**, **Xavier Garbet**, **Philippe Ghendrih**, **Guillaume Latu**, **Yanick Sarazin** (CEA, France), **Nicolas Bouzat** (INRIA, France), **Michel Mehrenberger** (University of Strasbourg, France)

Advances and Optimizations of Gyrokinetic Turbulence Code GKV towards Exascale Computing

Masanori Nunami (National Institute for Fusion Science, Japan)

Tuesday, July 3, 2018

16:30 – 17:00
Osaka Room

Understanding of the physics in the plasma turbulent transport is one of the critical issues in fusion researches. The gyrokinetic simulation, which solves time evolution of five-dimensional plasma distribution function, is a promising approach for that, but computationally challenging. In order to establish the predictive turbulence simulations with the multiple scale fluctuations and the multiple particle species, the local flux-tube gyrokinetic code GKV has been developed. Towards exascale supercomputing, a fully non-blocking optimized communication-computation overlap technique using assistant cores (AC), which are independent from the calculation cores, is proposed for the application to GKV with spectral (FFT) and finite-difference schemes. The effects of optimization are examined in Fujitsu FX100 (with 32 computing cores and 2 assistant cores/node), where AC enables us to employ the fully non-blocking communications overlapped by the OpenMP thread-parallelizations with much less overhead. It is clarified that the combination of the non-blocking techniques by AC and the thread-parallelization scheduling leads to not only reduction in OpenMP overhead, but also improved load/store and cache performance. In this talk, we would like to discuss recent advances and the optimization techniques of GKV towards exascale computing.

Co-Author(s): **Motoki Nakata** (National Institute for Fusion Science, Japan), **Tomo-Hiko Watanabe**, **Shinya Maeyama** (Nagoya University, Japan), **Akihiro Ishizawa** (Kyoto University, Japan), **Yasuhiro Idomura** (Japan Atomic Energy Agency, Japan)

CPU and GPU Parallelization of Spectral Particle Methods

Jakob Ameres (TU Munich, Germany)

Tuesday,
July 3, 2018

17:00 – 17:30
Osaka Room

Lagrangian particle methods suffer from noise in form of variance which increases with the degrees of freedom of the field discretization. Using a spectral basis e.g. Fourier modes for the fields reduces the degrees of freedom while retaining high order convergence. Additionally, such a Fourier spectral field discretization yields an energy and momentum conserving scheme for the Vlasov-Poisson and also Vlasov-Maxwell equations, which is very attractive from a physical point of view. Contrary to the standard Particle in Cell every particle contributes to every spectral basis function yielding a dense charge assignment. The evaluation of such an orthogonal basis can be implemented in various ways ranging from decompositions into expensive work packages with tiny memory footprint and vice-versa. This means that the underlying algorithm can be tuned between a compute- or memory-bound regime. Hence benchmarks between different methods and kernels using OpenCL on CPU and GPU are discussed.

Co-Author(s): **Eric Sonnendrücker** (TU Munich, Germany)

Porting a Legacy Global Lagrangian PIC Code on Many-Core and GPU-Accelerated Architectures

Noé Ohana (EPFL, Switzerland)

Tuesday,
July 3, 2018

17:30 – 18:00
Osaka Room

Modern supercomputer architectures are evolving towards embedding more and more cores per compute node, often making use of accelerators such as GPUs, in which thousands of threads can be executed concurrently. To make legacy codes profit efficiently from such resources usually requires a major refactoring effort. I will present the strategy that we adopted for the production code ORB5, a global gyrokinetic Particle-In-Cell (PIC) code for studying turbulence in tokamak plasmas, developed by many physicists over a period of 20 years, which clearly exceeds the timescale of HPC architecture evolution. Among others, the code now includes multiple kinetic species, electromagnetic effects, and collisions. The present refactoring work includes the restructuring of the main kernels, changing the data structure, multithreading with OpenMP on CPUs or OpenACC on GPUs, and optimization on different architectures. The modularity of the resulting code makes it more "future-proof", i.e. extensible to new physics features or computing architectures, and easier to maintain and develop in a collaborative fashion.

Co-Author(s): **Andreas Jocksch**, **Claudio Gheller** (ETH Zurich / CSCS, Switzerland), **Emmanuel Lanti**, **Aaron Sheinberg**, **Stephan Brunner**, **Laurent Villard** (EPFL, Switzerland)

MS35

Gravitational-Wave Data Analysis with the Current Generation of Advanced Detectors

Organizer(s): ***Maria Haney***, ***Philippe Jetzer***
(University of Zurich, Switzerland)

In the last two years, the field of gravitational wave astronomy has seen a breakthrough: In February 2016, the ground-based Laser Interferometer Gravitational-Wave Observatory (LIGO) announced the first direct detection of gravitational waves and the first observation of a binary black hole merger. Very recently, the LIGO-Virgo collaborations and their partner observatories announced the first joint observation of a neutron star merger in gravitational waves and the electromagnetic spectrum, marking a new era in multi-messenger astronomy. This minisymposium is intended to give an overview over certain important aspects of gravitational-wave data analysis with the current generation of ground-based detectors, highlighting applications of high-performance computing, machine learning and citizen science. The four presentations of the minisymposium will address the following topics: searches for gravitational waves in the detector data, computational aspects of source parameter estimation and physics implications of the LIGO-Virgo data, numerical relativity and its applications for the modelling of gravitational waves, aspects of data quality and data cleaning for the LIGO-Virgo data.

The LIGO/Virgo Search for Gravitational Waves

Alexander Nitz (Max Planck Institute for Gravitational Physics, Germany)

Wednesday, July 4, 2018

11:15 – 11:45
Osaka Room

The LIGO and Virgo detectors have completed a prolific observation run. We are now observing gravitational waves from both the mergers of binary black holes and neutron stars. We'll discuss how these discoveries were made and look into what the near future of searching for gravitational waves from compact binary mergers will look like.

Methods and Challenges in the Characterization of Gravitational-Wave Sources

Salvatore Vitale (Massachusetts Institute of Technology, USA)

Wednesday, July 4, 2018

11:45 – 12:15
Osaka Room

Advanced gravitational-wave detectors have so far detected signals emitted by the coalescence of two neutron stars or two black holes. Their astrophysical (masses, spins) and extrinsic (position, orientation) parameters have been estimated using stochastic samplers that efficiently explore a 15D parameter space. Although theoretically straightforward, parameter estimation can become non trivial due to the time required to calculate waveform models, the behavior of the noise, and the correlation between parameters. In this talk I will review the main methods results, and challenges associated with the characterization of compact binaries detected by advanced LIGO and Virgo. I will also underline which new challenges will arise as the sensitivity of gravitational-wave detectors improves by a factor of few or a factor of 10.

Numerical Relativity and its Applications for the Modelling of Gravitational Waves

Sascha Husa (University of the Balearic Islands, Spain)

Wednesday, July 4, 2018

12:15 – 12:45
Osaka Room

The optimally efficient detection of gravitational wave events, and the robust identification of the sources of such events rely on accurate models of the gravitational wave signals for astrophysically plausible events. Such models are synthesized from numerical solutions of the Einstein equations and perturbative models. This talk will review the current status and open challenges in solving the Einstein equations numerically as a system of partial differential equations as applied to the coalescence of compact binaries in the context of gravitational source modelling.

Data Quality for Gravitational-Wave Detectors

Andrew P. Lundgren (University of Portsmouth, UK)

Wednesday, July 4, 2018

12:45 – 13:15
Osaka Room

Gravitational-wave detectors are extremely complicated and unprecedently sensitive machines. We monitor many thousands of sensors and control systems to detect any hardware problems, electronics failures, noise from the environment, and all the many other subtle issues which degrade our sensitivity. I'll talk about many of the ways that we mine this large data set, including methods from signal processing and statistics, machine learning, and citizen science, to diagnose and solve the issues obstructing our searches for gravitational waves.

MS36

HPC for HEP: Enabling Big Data from Large Instruments on Leadership Class HPC Infrastructures

Organizer(s): ***Frank Wuerthwein*** *(UC San Diego, USA)*
Kaushik De *(The University of Texas at Arlington, USA)*

This minisymposium will present the latest advances in using HPC systems worldwide for physics results from large experiments in High Energy Physics and Particle Astrophysics. While HPC usage worldwide will be described, the experiences from the use of leadership class computing facilities in the US by the Large Hadron Collider experiments will be highlighted. LHC experiments have traditionally used the infrastructure provided by the Worldwide LHC Computing Grid. This has been supplemented in recent years by incorporating traditional HPC systems into the production and analysis computing systems at the LHC. Primarily CPU intensive simulation workflows are executed at HPCs - though other workflows are also being tested. The experiences gained by the LHC experiments over the past few years have opened HPC usage for other experimental and data intensive sciences. Four talks at this minisymposium will summarize the state of the art and the future wishlist for HPC usage for current and future experiments.

Running ATLAS Simulations on HPCs

Kaushik De (The University of Texas at Arlington, USA)

Wednesday, July 4, 2018

11:15 – 11:45
Sydney Room

Experiments at the Large Hadron Collider require data intensive processing and traditionally do not use HPCs. Till a few years ago, the ATLAS experiment at the LHC was using less than 10 million hours of walltime at HPCs annually, while over an exabyte of data was processed annually on the grid. A large increase in data volume and data complexity at the LHC in 2016 created a shortage of computing cycles, and HPC systems stepped in to help the LHC achieve its physics goals. Currently, ATLAS is on schedule to utilize about half a billion hours of walltime usage on HPCs during the past 12 months. This is a huge increase in usage over a few years - requiring numerous innovations and improvements. This talk will describe the use of HPCs worldwide by ATLAS, primarily for simulations, and specifically focus on how the HPCs are integrated with the workflow management and data management systems, and the lessons learned during this integration.

Big Data on HPC via HEPCloud

Dirk Hufnagel (Fermilab, USA)

Wednesday, July 4, 2018

11:45 – 12:15
Sydney Room

The higher energy and luminosity from the LHC in Run2 has put increased pressure on CMS computing resources. Extrapolating to even higher luminosities (and thus higher event complexities and trigger rates) beyond Run3, it becomes clear that simply scaling up the current model of CMS computing alone will become economically unfeasible. High Performance Computing (HPC) facilities, widely used in scientific computing outside of HEP, have the potential to help fill the gap. Here we describe the USCMS efforts to integrate US HPC resources into CMS Computing via the HEPCloud project at Fermilab. We describe the HEPCloud project, its goal and the driving vision behind it, to function as a portal to an ecosystem of diverse computing resources commercial or academic. We then will focus on how CMS has been using HPC for CMS workflows, mainly at the NERSC Cori facility. We will discuss our experiences in running data intensive workflows on HPC, especially how the IO requirements (both local storage and network) have mapped to HPC resources. We will describe past and current challenges and future plans for HPC use in CMS.

Perspective - Lessons from Titan, Looking to the Future

Jack Wells (Oak Ridge National Laboratory, USA)

Wednesday, July 4, 2018

12:15 – 12:45 Sydney Room

OLCF has seen the growing use of Titan for data sciences emerging from a variety of disciplines. The goals of experimental and observational data-intensive (EOD) science, like the ATLAS and ALICE experiments at the Large-Hadron Collider (LHC), are joining the goals of simulation studies in their requirements for access to computing at the largest scales. The BigPanDA project, along with other data projects, has served as a driver for innovation at OLCF. These innovations have included the opportunistic backfill in Titan's scheduled compute nodes with the large, malleable workload available from distributed-computing projects; as sand may fill the gaps between rocks packed into a jar. The OLCF has also deployed multiple container strategies to (1) automate deploying containers (Kubernetes/OpenShift) as a framework for providing user-required services and applications, and (2) HPC container runtimes (Singularity/Shifter) focused on use within in a batch submission system. Moreover, Titan's GPU-accelerated architecture has attracted a surge in machine-learning and artificial-intelligence workloads. With the advent of the Summit supercomputer in 2018 with over 27,000 machine-learning-optimized GPUs, high-bandwidth data movement, and large node-local memory, the volume of data analysis and machine-learning workloads is expected to grow significantly into the future.

HPC Systems and the Integration Challenges of Large Instruments

Frank Wuerthwein (UC San Diego, USA)

Wednesday, July 4, 2018

12:45 – 13:15 Sydney Room

Large Instruments like the four LHC detectors, LIGO, Virgo, IceCube, LSST, DUNE, BelleII, and so forth, are designed, built, and operated by large international collaborations. To bring all the resources to bear that these collaborations have access to requires globally distributed Cyberinfrastructure that leadership class HPC systems need to integrate into seamlessly in order to be maximally useful. We discuss an architectural wish list for future large scale HPC systems in order to better support science with large instruments. We give some examples for how these features have been integrated into existing machines.

MS37

HPUQ: Current Challenges in Uncertainty Quantification for Mechanistic Models, Part II: Theory, Methods and Tools

Organizer(s): ***Ritabrata Dutta*** *(Università della Svizzera italiana, Switzerland)*
Nikos Karathanasopoulos *(ETH Zurich, Switzerland)*
Bastien Chopard *(University of Geneva, Switzerland)*

The HPUQ minisymposium focuses on uncertainty quantification (UQ) of mechanistic models for natural sciences (eg. Engineering, Life and Aquatic Sciences) using high performance computing (HPC). The statistical inference (e.g. calibration) of models for complex mechanistic models, in the abundance of data arriving from heterogeneous sources, poses a methodological and computational challenge for scientists. In the first session, the minisymposium highlights cutting edge frameworks for rigorous and robust UQ as ABCpy, ∏4U, PyMLMC, SPUX to address these issues, with a focus towards optimal algorithmic performance and efficient utilization of HPC resources. In the second session of the minisymposium, we shift the focus to the applications of UQ methodologies in several important scientific domains spanning from Biomedicine and Biomechanics to Aerospace Engineering and Fluid Dynamics.

ABCpy: Benchmarking ABC Algorithms from HPC Perspective

Ritabrata Dutta (Università della Svizzera italiana, Switzerland)

Wednesday, July 4, 2018

11:15 – 11:45
Singapore Room

ABCpy is a highly modular scientific library for Approximate Bayesian Computation (ABC) written in Python. Our main contribution is to illustrate a software engineering effort that enables domain scientists to easily apply ABC (for likelihood-free Bayesian uncertainty quantification of Mechanistic models) to their research without being ABC experts; using ABCpy they can easily run large parallel simulations without much knowledge about parallelization, even without much additional effort to parallelize their code. Further, ABCpy enables ABC experts to easily develop new inference schemes and evaluate them in a standardized environment and to extend the library with new algorithms. These benefits come mainly from the modularity of ABCpy. We give an overview of the design of ABCpy and provide a performance evaluation concentrating on parallelization. This points us towards the inherent imbalance in some of the ABC algorithms. We develop a dynamic scheduling MPI implementation to mitigate this issue and classify ABC algorithms according to their adaptability towards high-performance computing.

Co-Author(s): **Marcel Schoengens** (ETH Zurich / CSCS, Switzerland), **Avinash Ummadisinghu**, **Antonietta Mira** (Università della Svizzera italiana, Switzerland), **Nicole Widerman** (ETH Zurich, Switzerland), **Jukka-Pekka Onnela** (Harvard University, USA)

The Hierarchical Bayesian Framework Applied to Molecular Dynamics

Georgios Arampatzis (ETH Zurich, Switzerland)

Wednesday, July 4, 2018

11:45 – 12:15
Singapore Room

The Hierarchical Bayesian (HB) framework for the quantification of uncertainty and model selection in the presence of heterogeneous data will be presented as well as an efficient algorithm for the sampling of the posterior, high dimensional distribution. The framework is then applied to two problems from Molecular Dynamics. In the first problem we revisit the exponent related to the repulsion force in the Lennard-Jones potential. Using experimental data from the radial distribution function of argon in various thermodynamic conditions we show that the exponent should be close to approximately 6.5. We show that the proposed 6-p potential applies to a wider range of thermodynamic conditions, than the classical 6-12 potential. In the second problem we address the question of the best coarse-grained model for liquid water. Typically, the level of coarse-graining and the model complexity are preselected based on physical intuition. These assumptions are rarely systematically addressed even though the model's accuracy, efficiency and transferability critically depends on them. We propose the HB framework as a mean for the rigorous selection of coarse-graining level and show its validity conditional on macroscopic quantities on various thermodynamic conditions.

Co-Author(s): **Petros Koumoutsakos** (ETH Zurich, Switzerland)

PyMLMC + SPUX: Uncertainty Quantification Using Multi-Level and Particle Filtering Techniques

Jonas Sukys (Swiss Federal Institute of Aquatic Science and Technology, Switzerland)

Wednesday, July 4, 2018

12:15 – 12:45
Singapore Room

Evolution of complex systems such as hydrodynamic flows and ecological networks can be modeled using differential equations and individual based models. Examples include Saint-Venant and Navier-Stokes equations for lakes, rivers and tsunamis, multi-phase Euler equations for cavitation, Darcy's law for porous flows and predator-prey foodwebs for mesocosm dynamics. Many of such dynamical systems strongly depend on alleatorically uncertain input data, such as initial data, sources and model coefficients, with additional epistemic uncertainty directly influencing the evolution trajectories. In this talk I will introduce two parallel uncertainty quantification frameworks: PyMLMC and SPUX. PyMLMC propagates uncertainty in model input using optimal fidelity multi-level Monte Carlo sampling, which significantly accelerates standard Monte Carlo method by clever variance reduction relying on a series of coarse resolution simulations used as control variates. A significantly more challenging task is probabilistic model parameter estimation incorporating prior expert knowledge and observed experimental data. The Python framework SPUX employs Particle Markov Chain Monte Carlo for efficient marginal likelihood approximations by iteratively evolving and adaptively filtering multiple state particles in parallel while completely avoiding heavy filesystem access. Modularity and efficient use of computational resources makes SPUX accessible for domain scientists interested in Bayesian inference for complex stochastic models.

Low-Rank Tensor Approximations for Sensitivity Analysis of Complex Models with High-Dimensional Input

Katerina Konakli (COWI, Denmark)

Wednesday, July 4, 2018

12:45 – 13:15
Singapore Room

This work presents a computationally efficient method for conducting sensitivity analysis of complex, expensive-to-evaluate computer models. The focus is set on the so-called Sobol' sensitivity indices, which represent the fraction of the total variance of a response quantity of interest (QoI) that can be attributed to a random input variable or a group thereof. The proposed method for computing these indices is based on substituting the original model with a low-rank tensor approximation (LRA) meta-model. The LRA meta-model provides a statistically equivalent representation of the QoI as a sum of rank-one tensors, the parameters of which can be determined from a relatively small number of runs of the original model. Because the number of unknown parameters in a LRA meta-model grows only linearly with the dimension of the random input, LRA can be particularly efficient in high-dimensional problems. It is demonstrated that the Sobol' indices can be computed *analytically* in terms of the LRA parameters, thus enabling efficient analysis of computationally heavy models. The accuracy and efficiency of the approach is manifested in example applications related to structural mechanics, heat conduction and hydrogeology.

MS38

Mass and Energy Transport Phenomena in Solid State

Organizer(s): ***Ivano Tavernelli***, ***Matthieu Mottet***
(IBM Research, Switzerland)

Modeling energy and mass transport phenomena in the solid state accurately and efficiently in computer simulations is of critical importance to the discovery and design of new materials in important applications, such as lithium-ion solid-state electrolytes, proton-conducting fuel cells and better thermoelectrics. Progress in recent years enabled the modeling community to elucidate transport mechanisms in several materials of importance, but to achieve further progress major obstacles need to be tackled.

Firstly, the computational cost of the current methods limits the time and lengths scales of the systems that can be modeled. We will present several efforts that are targeted either at reaching larger system sizes and time scales in simulations, or to improve the estimates of transport coefficients from finite simulation lengths using more advanced statistical tools.

Secondly, with increases in computer power and novel methods for modeling and analysis it has become possible to push for large-scale screening of structural databases for desired transport properties. Such high-throughput approaches require new concepts of automatization, data reproducibility and results dissemination in the field of materials' simulation, and we will present the developed concepts, applications and results.

The Materials Genome in Action

Seyed Mohamad Moosavi (EPFL, Switzerland)

Wednesday, July 4, 2018

11:15 – 11:45
Boston 3 Room

It is now possible to make an enormous spectrum of different nonporous materials simply by changing the building blocks in the synthesis of Metal Organic Frameworks (MOFs) or related materials. This unique chemical tenability allows us to tailor-make materials that are optimal for a given application. The promise of finding just the right material seems remote, however: because of practical imitations, we can only ever synthesise, characterise and test a tiny fraction of all possible materials. To take full advantage of this development, therefore, we need to develop alternative techniques, collectively referred to as Materials Genomics, to rapidly screen large numbers of materials and obtain fundamental insights into the chemical nature of the ideal material for a given application. In this lecture we illustrate this approach by suggesting how to obtain optimal materials for gas separations and gas storage.

Co-Author(s): **Berend Smit** (EPFL, Switzerland)

High-Throughput Screening for New Solid-State Electrolyte Candidates

Leonid Kahle (EPFL, Switzerland)

Wednesday, July 4, 2018

11:45 – 12:15
Boston 3 Room

Extensive computational screening of structural databases for solid-state ionic conductors can lead to novel candidate materials for next-generation solid-state lithium-ion batteries, while deepening our understanding of the microscopic processes governing ionic diffusion in the solid state. Such a task is ambitious because no current modeling technique is both unbiasedly predictive for chemically diverse systems and computationally affordable. In order to reach this goal of efficiency and accuracy, we simplify the potential-energy surface that would be provided by density-functional theory with physically motivated approximations. The result is a novel hybrid quantum/empirical model that can be used to perform molecular dynamics simulations of solid-state diffusion. The efficiency of the model allows one to adopt a high-throughput screening approach, that here is deployed using the AiiDA materials informatics platform. In this talk I will present the different screening protocols and show how high-level workflows can automatize and streamline the calculation of transport coefficient. Last, the full provenance of the data calculated is fully preserved by AiiDA, allowing one to search in the data atomistic descriptors that are predictive for ionic diffusion.

Co-Author(s): **Aris Marcolongo** (IBM Research, Switzerland), **Nicola Marzari** (EPFL, Switzerland)

Doping Solid-State Electrolytes: Classical Modelling and Insights

Matthieu Mottet (IBM Research, Switzerland)

Wednesday, July 4, 2018

12:15 – 12:45
Boston 3 Room

Doping and elemental substitution are essential tools for the optimisation of the conductivity of solid-state electrolytes (SSE). Therefore, the ability to model doped structure is an important step to discover and optimise new SSE in silico. The complexity and disorder introduced by doping make it computationally unrealistic to use first-principle frameworks for high throughput studies. Instead, the use of classical force-fields allows to overcome this limitation. In this talk, I will present a systematic approach to the training of polarisable force-fields for solid-state electrolytes and demonstrate a methodology allowing to account for the disorder introduced by doping. This approach is similarly well suited for structure with partial ionic occupancy. In particular, I will showcase detailed results for W-doped garnet-type SSE, with particular attention to the impact of the change in the carrier concentration and in the local potential introduced by the impurities on the dynamics of the system. The study underlines the challenges presented by doping and the complex interplay between its thermodynamic and kinetic effects.

Co-Author(s): **Aris Marcolongo**, **Ivano Tavernelli**, **Teodoro Laino** (IBM Research, Switzerland)

Accurate Thermal Conductivities from Optimally Short Molecular Dynamics Simulations

Loris Ercole (SISSA, Italy)

Wednesday, July 4, 2018

12:45 – 13:15
Boston 3 Room

The evaluation of thermal transport coefficients in extended systems is known to require impractically long simulations, thus calling for a paradigm shift that would allow to deploy state-of-the-art quantum simulation methods. We recently introduced a new method[1] to compute these coefficients from optimally short molecular dynamics simulations, based on the Green-Kubo theory of linear response and the cepstral analysis of time series. Information from the *full* sample power spectrum of the current for a *single* and relatively short trajectory is leveraged to evaluate and optimally reduce the noise affecting its zero-frequency value, whose expectation is proportional to the corresponding conductivity. Our method is unbiased and consistent, in that both the resulting bias and statistical error can be made arbitrarily small in the long-time limit. A simple data-analysis protocol is proposed and validated in some paradigmatic cases (liquid-Ar and H2O, crystalline Mgo and amorphous-SiO2), showing that simulation times of one to a few hundred picoseconds are sufficient in these systems to achieve an accuracy of the order of 10% on the estimated thermal conductivities. Finally, we present a first application of this method to the *ab initio* simulation of heat transport in amorphous silica. [1]Ercole, Marcolongo, and Baroni, *Sci. Rep.* 7, 15835 (2017)

Co-Author(s): **Aris Marcolongo** (IBM Research, Switzerland), **Stefano Baroni** (SISSA, Italy)

MS39

Scalable Solvers for Forward and Inverse Problems in Geophysics

Organizer(s): ***Christian Boehm***, ***Václav Hapla***
(ETH Zurich, Switzerland)

Nearly all fields in geophysics combine numerical models and data measurements to predict the future state of a dynamical system and/or infer unknown parameters. Such models may produce highly nonlinear systems with extremely large numbers of unknowns.

The ever-increasing power and widespread availability of massively parallel supercomputers offers researchers the opportunity to continually increase both the spatio-temporal resolution and the physical complexity within their numerical models. However, this requires access to solvers that can harness the resources of high-performance computing clusters efficiently and which scale for problem sizes with billions of degrees of freedom.

In this minisymposium, we discuss numerical and algorithmic approaches, scientific libraries and coding practices to develop and maintain scalable solvers for various applications in geophysics. Examples include, but are not limited to, seismic wave propagation and imaging, geodynamics, and hydro-mechanical processes.

Extreme Scale Seismic Wave Propagation Simulation for Mars

Vaclav Hapla (ETH Zurich, Switzerland)

Wednesday, July 4, 2018

11:15 – 11:45
Darwin Room

In 2018, the NASA InSight mission will place a highly sensitive broadband seismometer on Mars' surface to investigate its deep interior structure. As Mars has strong 3D features such as topography and crustal thickness, elastic wave propagation simulations are a crucial component for data interpretation. At the highest frequencies that are typically used, the resulting system has trillions of spatial degrees of freedom and requires hundreds of thousands of time steps. As there will only be a single receiver, reciprocity of the wave equation can be used to swap sources and receivers, so a total of three numerical simulations allows computing seismograms for any number of seismic sources. To solve problems of that scale, we develop the Salvus software suite for full waveform modelling and inversion. Salvus makes use of the well-known PETSc toolkit. Its DMPlex module represents a mesh by a graph whose vertices represent cells, faces, edges and nodes uniformly. Discretization methods can be used unchanged for meshes of different shapes and dimensions. For such large simulations, loading the whole mesh onto a single processor must be avoided. Hence, we employ parallel I/O, on-the-fly partitioning, and load-balancing techniques, working with a distributed DMPlex representation throughout the whole simulation.

Co-Author(s): **Martin van Driel**, **Michael Afanasiev**, **Christian Boehm**, **Lion Krischer** (ETH Zurich, Switzerland)

Seismic Wave Propagation on Complex Topographies Applied in the Alpine Area Using the ExaHyPE Hyperbolic PDE Engine

Leonhard Rannabauer (TU Munich, Germany)

Wednesday, July 4, 2018

11:45 – 12:15
Darwin Room

ExaHyPE is a Horizon 2020 EU project to develop a high-performance engine to solve hyperbolic systems of PDEs using the high-order discontinuous Galerkin finite element method. The project goals are to develop an engine with flexible support for various applications which shall be tailored towards expected exascale architectures. The end-user is provided with an abstraction of the complicated algorithms to implement the ADER-DG numerical scheme and of the issues related to scalability and parallel adaptive mesh refinement (AMR), which are handled internally by the Peano framework. In our presentation we will give an introduction on how to implement scalable seismic wave propagation algorithms on complex topographies in the ExaHyPE engine. We will show and compare time-to-solution results for two different approaches for simulations in the alpine area. By remaining internally on a Cartesian mesh both methods allow modeling of the topography without the cumbersome meshing process. First is a newly developed curvilinear mesh approach which we are able to implement by only transforming flux and source terms. Second is a diffuse interface method which extends the PDE by a parameter handling the transition from solid to surface to a non-linear system.

Co-Author(s): **Kenneth Duru**, **Alice-Agnes Gabriel** (Ludwig Maximilian University of Munich, Germany), **Michael Bader** (TU Munich, Germany)

StagBL: A Scalable, Portable, High-Performance Discretization and Solver Layer for Geodynamic Simulation

Patrick Sanan (ETH Zurich, Switzerland)

Wednesday, July 4, 2018

12:15 – 12:45
Darwin Room

StagBL is an open-source parallel solver and discretization library for geodynamic simulation, encapsulating and optimizing operations essential to staggered-grid finite volume Stokes flow solvers. These form the basis for highly-efficient application codes for long-term mantle convection and lithospheric dynamics. StagBL prevents common bottlenecks to improving scalability, swapping solvers, adapting to new architectures, and optimizing performance. The StagBL project addresses these issues by providing a streamlined library to provide a path to performance from toy codes to quality, scalable implementations. It provides a parallel staggered-grid abstraction in C and Fortran, and an interface (DMStag) for PETSc. Tools are available to define boundary conditions, interact with particle systems, and efficiently solve Stokes systems in small (direct solver), medium (simple preconditioners), and large (block factorization and multigrid) model regimes. By implementing common kernels beneath a uniform abstraction layer, StagBL enables optimization for modern hardware, thus reducing community barriers to large-scale parallel simulation on modern architectures, and a platform to develop innovative new tools. By working directly with leading application codes StagYY, I3ELVIS, and LaMEM, and providing an API and examples for others, StagBL aims to become a community tool supplying scalable, portable, reproducible performance to novel science in regional- and planet-scale geodynamics.

HPC Solution Methods for Simulation of Hydro-Mechanical Processes in Geo-Environment

Radim Blaheta (Institute of Geonics CAS, Czech Republic)

Wednesday, July 4, 2018

12:45 – 13:15
Darwin Room

First, the contribution describes a hydro-mechanical model which combines Richards' model for variably saturated flow and nonlinear elasticity. If the flow and deformation are coupled by deformation dependent permeability and retention relation and saturation dependent elastic moduli, then the coupled hydro-mechanical model is suitable e.g. for simulation of processes in bentonite based engineering barriers in deep geological repository for the high-level radioactive waste. A validation of such type model was done within the international DECOVALEX project. The second part of the contribution deals with discussion about suitable iterative solution methods which combine iterations for solving coupled hydro-mechanical system with iterations for solving decoupled problems. Finally, we describe HPC solution methods for the systems arising from linearization.

Co-Author(s): **Martin Hasal**, **Jakub Kruzik**, **Tomas Luber**, **Zdenek Michalec**, **Jiri Stary** (Institute of Geonics CAS, Czech Republic)

MS40

Towards Weather and Climate Simulations at 1-km Resolution

Organizer(s): ***Peter Dominik Dueben*** *(ECMWF, UK)*
Carlos E. Osuna *(MeteoSwiss, Switzerland)*

Reliable weather predictions and climate projections are of vital importance for society and for the creation and preservation of prosperity. An increase in horizontal resolution in global simulations of the atmosphere to ~1 km would enable us to represent large cloud-systems and deep convection explicitly within simulations. This would reduce much of the uncertainty in predictions of weather and climate. Unfortunately, it is not possible yet to run operational simulations with state-of-the-art global weather and climate models at this level of resolution using today's supercomputing facilities. This minisymposium will provide an update on the progress of global weather and climate models when running at 1 km horizontal resolution. The most significant challenges towards cloud-resolving simulations will be assessed (including scalability, simulations far away from peak-performance, I/O volume and efficient time-stepping schemes) and possible ways to overcome these barriers will be discussed.

At the Edge of Resolution: Earth System Modelling at ECMWF

Nils P. Wedi (ECMWF, UK)

Wednesday, July 4, 2018

11:15 – 11:45
Rio Room

Global km-scale simulations of weather and climate are at the frontier to become technically feasible as a research tool for exploring the added value of resolving rather than parametrising the vertical redistribution of momentum and heat at globally uniform resolutions of O(1km). Handling and processing the global data is still pushing the (design) limits of current infrastructures, but it is expected to inspire the future design process, and pave the road for the next generation of weather and climate services. For example, ECMWF plans to run an ensemble of forecasts at 5km resolution in 2025. This talk will discuss and compare simulations at 18km, 9km, 5km, 2.5km and 1.25km resolution, by analysing the statistical differences of horizontal and vertical motions, and expected ensemble statistics for extreme events at 5km, and by limited investigations into the relative effects of different choices for the numerics and (remaining) physical parametrizations. These illustrate that resolution is important but not the only aspect delivering continuous improvements in global numerical weather prediction.

Using Global Cloud-Resolving Models for Weather Predictions and for Studies of Clouds in the Climate System

Shian-Jiann Lin (NOAA, USA)

Wednesday, July 4, 2018

11:45 – 12:15
Rio Room

The FV3 group at the NOAA/GFDL and the RCEC at the Academia Sinica, Taiwan, is developing a new type of Global Cloud-Resolving Model (GCRM) based on an integrated dynamics-physics concept, in which several fast-acting physics (e.g., cloud microphysics) are incorporated into a new FV3 (nu-FV3) framework. This new model is still being actively developed. This new model improves the dynamics-physics interaction and increases in computational efficiency due to the separation of the fast-acting physics from the slow-physics, allowing a near tenfold increase in overall time step. We have also built some of the SubGrid Orographical processes into the nu-FV3 dynamics, which unavoidably breaks the traditional boundary between "dynamics" and "physics". We believe the boundary between the "dynamics" and "physics" set by the traditional modeling framework is one reason that limits modeling advancements. A preliminary version of this new type of GCRM is used for the DYAMOND project. We will carry out several 40-day "convective-parameterization-free" experiments across the gray-zone at three different horizontal resolutions: 13, 6.5, and 3.25 km. As a potential tool for sub-seasonal predictions, we shall analyze the hindcast skill (first 10 days) as well as the systematic "climate basis" for the last 30 days.

Near-Global RCM Simulations to Establish a Baseline for Global 1 km GCM Simulations

Oliver Fuhrer (MeteoSwiss, Switzerland)

Wednesday, July 4, 2018

12:15 – 12:45
Rio Room

Reducing the horizontal resolution of global weather and climate models to the kilometer-scale holds the promise of reducing some of the long-standing biases and uncertainties. At these resolutions, some of the key processes such as deep convection, gravity wave drag and ocean eddies can be resolved explicitly on the model grid and thus much closer to first principles. But how far are we from achieving this goal? The presentation will show results from scaling a regional climate model (COSMO) to cover almost the entire Earth at increasing resolutions of up to 1 km. COSMO has been systematically adapted to make use of hybrid compute node designs with accelerators such as graphics processing units (GPUs) and thus can make efficient use of all of Europe's currently largest supercomputer, Piz Daint. To our knowledge this represents the first complete atmospheric model being run entirely on accelerators at this scale. At a grid spacing of 930 m (1.9 km), we achieve a simulation throughput of 0.043 (0.23) simulated years per day and an energy consumption of 596 MWh per simulated year. We discuss the implications of these simulations as a baseline for what is achievable when systematically adapting our codes to make use of emerging computer architectures.

Co-Author(s): **Tarun Chadha**, **Torsten Hoefler**, **Grzegorz Kwasniewski**, **Daniel Luethi**, **David Leutwyler**, **Christoph Schär** (ETH Zurich, Switzerland), **Xavier Lapillonne**, **Carlos E. Osuna** (MeteoSwiss, Switzerland), **Thomas Schulthess**, **Hannes Vogt** (ETH Zurich / CSCS, Switzerland)

ESCAPE: Energy-Efficient Scalable Algorithms for Weather Prediction on Exascale Supercomputers

Andreas Mueller (ECMWF, UK)

Wednesday, July 4, 2018

12:45 – 13:15
Rio Room

In the simulation of complex multi-scale flow problems, such as those arising in weather and climate modelling or in engineering, one of the biggest challenges is to satisfy operational requirements in terms of time-to-solution and available energy without compromising the accuracy and stability of the solution. These two competing factors require extreme computational capabilities in conjunction with state-of-the-art algorithms that can optimally suit the targeted underlying hardware while improving the convergence to the desired solution. The European Centre for Medium Range Weather Forecasts (ECMWF) is leading the H2020 FET-HPC project ESCAPE (Energy-efficient SCalable Algorithms for weather Prediction on Exascale supercomputers). The ESCAPE project includes the development of new algorithms that are specifically designed for better energy efficiency, testing and optimisation of different numerical techniques and improved portability through domain specific languages. The project incorporates through ECMWF's project partners the expertise of leading European regional forecasting consortia, university research, experienced high-performance centres and hardware vendors. This talk gives an overview of the ESCAPE project and summarises some of the key results obtained so far.

Co-Author(s): **Willem Deconinck**, **Nils P. Wedi**, **Peter Bauer** (ECMWF, UK)

MS41

Use of AI to Analyze Complex Biological Systems

Organizer(s): ***Daniel Jacobson*** *(Oak Ridge National Laboratory, USA)*
Ben Brown *(Lawrence Berkeley National Laboratory, USA)*
Georgios Gkoutos *(University of Birmingham, UK)*

The cost of generating biological data is dropping exponentially, a decrease that has far outstripped predictions based on Moore's Law. This has ushered in a new era of systems biology in which there are unprecedented opportunities to gain insights into complex biological systems. Integrated biological models need to capture the higher order complexity of the interactions among cellular components. Solving such complex combinatorial problems will give us unprecedented levels of understanding of biological systems. However, this leads to a combinatorial explosion in the search space of biological data. These exponentially increasing volumes of data, combined with the desire to model more and more sophisticated sets of relationships within a cell and across an organism (or in some cases even ecosystems), have led to an unmet need for computational resources and sophisticated algorithms that can make use of them. Thus the bottleneck in biological science is no longer data generation but is in fact computational analysis. A full model of all of the higher order interactions of cellular and organismal components is one of the ultimate grand challenges of systems biology. The use of machine and deep learning algorithms provide some of the methodologies with which to achieve this goal.

Explainable AI and the Discovery of Complex Genetic Architectures: Plants, Insects and Humans: Systems Biology and the 3D Interactome

Daniel Jacobson (Oak Ridge National Laboratory, USA)

Wednesday, July 4, 2018

11:15 – 11:45
Samarkand Room

Phenotypes emerge from the interactions of genes in the context of life history, environmental exposure and physiology. We use algorithms that enable the discovery of interactions amongst sets of genes and across defined phenotypes, conditional on different exposomes, which collectively constitute strong predictors of physiological outcomes, while providing insight into molecular mechanisms. Thus, we complement our application of traditional QTL mapping techniques with the newly developed iterative Random Forests (iRF) procedure that we have scaled up to run on OLCF supercomputing architectures. For each phenotype, we regress scores against genotype using iRF. Importantly, iRF generates scores for main effects (individual variants / QTLs) and interactions (epistatic networks) on the same scale, enabling the direct comparison of effect size and statistical significance. The resulting collection of QTLs can be represented and interrogated as a hypergraph and community detection analysis on the resulting hypergraphs enables the detection of larger epistatic modules. These techniques, along with SNP correlation algorithms are used to find epistatic interactions and predict protein complexes. Those predictions become the basis for the construction of the 3D interactome of a given species proteome.

Scalable Deep Learning for Extracting Cancer Phenotypes from Unstructured Clinical Text

Georgia Tourassi (Oak Ridge National Laboratory, USA)

Wednesday, July 4, 2018

11:45 – 12:15
Samarkand Room

Pathology reports are a primary source of information for cancer registries which process high volumes of free-text reports annually. Information extraction and coding is a manual, labor-intensive process. In this talk I will discuss the latest deep learning technology, presenting both theoretical and practical perspectives that are relevant to natural language processing of clinical pathology reports for computational cancer phenotyping. Using different deep learning architectures, I will present benchmark studies for various information extraction tasks and discuss their importance in supporting a comprehensive and scalable national cancer surveillance program.

DeepPVP: Phenotype-Based Prioritization of Causative Variants Using Deep Learning

Georgios Gkoutos (University of Birmingham, UK)

Wednesday, July 4, 2018

12:15 – 12:45
Samarkand Room

Prioritization of variants in personal genomic data is a major challenge. Recently, computational methods that rely on comparing phenotype similarity have shown to be useful to identify causative variants. In these methods, pathogenicity prediction is combined with a semantic similarity measure to prioritize not only variants that are likely to be dysfunctional but those that are likely involved in the pathogenesis of a patient's phenotype. Here we present DeepPVP, a variant prioritization method that combines automated inference with deep neural networks to identify the likely causative variants in whole exome or whole genome sequence data. We demonstrate that DeepPVP performs significantly better, both in terms of speed and accuracy, than existing methods, including phenotype-based methods that use similar features.

Interpretable Density Estimation in Genomics Data

Ben Brown (Lawrence Berkeley National Laboratory, USA)

Wednesday, July 4, 2018

12:45 – 13:15
Samarkand Room

Recent advances in unsupervised learning, including VAEs and GANs, train generative models with implicit densities parametrized by neural networks. Because such methods do not present an explicit density, they cannot tractably be used as surrogate models for optimization or search in auxiliary objectives, which reduces their utility in science and engineering applications. Further, these methods perform poorly on high-dimensional data, particularly in settings where the intrinsic dimension is much less than the ambient dimension. As a result, such techniques are ineffective in biological contexts, particularly in the case of genetics and functional genomics data, where we work almost exclusively in the $N << P$ regime. We introduce a novel adversarial training algorithm for explicit density models, which iterates between fitting a discriminator random forest, and approximating a density using feature interactions discovered in the forest. This method, which we term dueling iterative Random Forests (diRFs), can find high-order interactions amongst features, is completely parallelizable, performs well in the high-dimensional regime, and produces density models that can be used in constrained optimization problems. We apply this algorithm to human and plant population genetics and genomics data, and present findings.

Co-Author(s): **Daniel Jacobson** (Oak Ridge National Laboratory, USA)

MS42

Coupling Strategies Towards Exascale for Complex Earth System Modelling

Organizer(s): ***Willem Deconinck*** *(ECMWF, UK)*
Katherine Evans *(Oak Ridge National Laboratory, USA)*

The algorithms underlying numerical weather prediction (NWP) and climate models that have been developed in the past few decades face an increasing challenge to adapt to paradigm shifts imposed by new hardware developments. The emerging diverse and complex hardware solutions have a large impact on the programming models traditionally used in NWP and climate modelling software, triggering a rethink of design choices for future software frameworks and how Earth system model (ESM) components interact. Furthermore there is a drive to increase the model complexity to include ever more processes of the whole Earth system. These complexities will inevitably break NWP modelling infrastructures that did not take these aspects into consideration 30+ years ago. As upcoming hardware solutions evermore push the boundaries of parallel execution, each ESM component may be reaching a limit in parallel scaling efficiency.

Already, coupling infrastructure exists that enables concurrent execution of model components. In order to make use of increasing parallelism, it may be required to take a step backwards, and redesign the ESM components to become more modular and enable more concurrent execution. This minisymposium will provide an update on progress both in infrastructure developments and developments in concurrent model component execution.

Flexible Earth System Modelling on Multiple Grids

Willem Deconinck (ECMWF, UK)

Wednesday, July 4, 2018

14:15 – 14:45
Rio Room

There is a drive to increase the model complexity to include ever more processes of the whole Earth system. Some of these processes may require computations on grids of different type or resolution than the atmospheric grid. Multiple grid structures may be required as part of the numerical filtering strategy for atmospheric wave motions or to simply save computational cost of selected physical processes. These different grids may have different domain decompositions for parallel computations, and different parallelisation strategies. Moreover, the internal memory layout for a field that is optimal for a specific combination of hardware architecture and numerical algorithm may not be optimal for another combination. These complexities will inevitably break NWP modelling infrastructures that did not take these aspects into consideration 30+ years ago. In this talk, we demonstrate how Atlas, a new library developed at ECMWF, is used to complement ECMWF's Integrated Forecasting System (IFS) model to enable a number of physical processes to be implemented on multiple grids. Atlas helps to accommodate flexibility in hardware and software choices as well as increasing model complexity.

Co-Author(s): **Michail Diamantakis** (ECMWF, UK)

Comodels: A New Approach for Coupling Models for the [Tera,Exa]Scale

George Mozdzynski (ECMWF, UK)

Wednesday, July 4, 2018

14:45 – 15:15
Rio Room

ECMWF's IFS spectral model has been using a hybrid MPI/OpenMP parallelization approach since about 2002 when just 2 OpenMP threads per MPI task (OMP_NUM_THREADS=2) was used on an IBM Power4 system. The number of threads per task have gradually increased over the years to today where 12 or 18 achieves the best performance on a CRAY XC-40 (Broadwell) depending on the specific IFS model resolution. Further, it is realized that simply increasing the threads per task does not deliver improved performance. Likewise, the MPI communications cost of the IFS spectral model imposes severe demands on the switch fabric and scaling high resolution model cases beyond o(100k) cores yields minimal performance gains. Given that future supercomputer architectures are expected to support o(100) computational threads per MPI task (or socket) a codesign effort involving ECMWF, EPCC, and CRAY was started in 2015 to explore how such an extreme thread count could be achieved in practice. This effort followed on from this team's successful collaboration in the EU funded CRESTA project (2012-2014). In this talk we will present an evolutionary OpenMP parallelization approach for IFS which we call comodels which is an ongoing development at ECMWF and present some new results of this work and future plans.

Modeling Systems at the End of Dennard Scaling

Venkatramani Balaji (Princeton University, USA)

Wednesday, July 4, 2018

15:15 – 15:45 Rio Room

Conventional computational hardware has reached some physical limits: the phenomenon known as 'Dennard scaling' gave rise to Moore's Law, and many cycles of exponential growth in computing capacity. The consequence is that we now anticipate a computing future of increased concurrency and slower arithmetic. Earth system models, which are weak-scaling and memory-bandwidth-bound, face a particular challenge given their complexity in physical-chemical-biological space, to which mapping single algorithms or approaches is not possible. A particular aspect of such 'multi-scale multi-physics' models that is under-appreciated is that they are built using a combination of local process-level and global system-level observational constraints, for which the calibration process itself remains a substantial computational challenge. In this talk, we examine approaches to Earth system modeling in the post-Dennard era. The possibilities include following the industry trend toward machine learning and build models that learn; stochastic methods and emulators for fast exploration of uncertainty; using fewer bits of precision, among others. The talk will present ideas and challenges as we prepare for a post-Dennard future.

Making the Expensive Affordable: Running a Chemistry Model in the UKESM Climate Model

Richard Hill (Met Office, UK)

Wednesday, July 4, 2018

15:45 – 16:15 Rio Room

The Met Office/NERC UKESM coupled climate model employs a chemistry model (UKCA) embedded within the Unified Model (UM) atmosphere. At high resolutions this is prohibitively expensive due to computational costs of the chemistry. Consequently, we have been developing a mechanism to run the chemistry model at a lower resolution than the "main" atmospheric model. Since the chemistry code is embedded within the UM, it may not be run it as a separate stand-alone entity or at a different resolution from the atmosphere code. We believe it would be possible to achieve the necessary performance by creating a "hybrid" coupled model: a configuration featuring two concurrently run copies of the same UM code at different resolutions. One component would run at high resolution, without chemistry, the other component at lower resolution featuring full chemistry. This approach requires the coupling exchange of 3D fields between the two components. We are putting together a coupled system featuring a full resolution UM atmosphere, a reduced resolution UM atmosphere-UKCA chemistry, a NEMO ocean and a CICE sea ice model coupled using OASIS3-MCT. I aim to give an overview of the plans, progress made and issues which may yet prove to be major stumbling blocks.

MS43

Distributed Asynchronous Parallel Computing: Progress and Challenges for Multi-Physics Applications on Heterogeneous Architectures

Organizer(s): ***Hemanth Kolla*** *(Sandia National Laboratories, USA)*
Jacqueline Chen *(Sandia National Laboratories, USA)*

In the push towards exascale computing the current paradigm of bulk-synchronous distributed computing is giving way to a more asynchronous paradigm, where it is becoming increasingly important to have the inherent asynchrony in an application be exploited for performance and scalability. Whether the asynchrony is enabled algorithmically, expressed explicitly in the user program, discovered by a programming model and/or runtime there are critical research questions that need to be addressed. In this paradigm, *asynchronous many task* (AMT) programming models have made great progress in demonstrating the concept and paving the way forward. While asynchronous task based programming models for shared-memory systems have been around for a long time serious challenges remain in extending them to a distributed memory setting. This minisymposium is targeted at providing a platform to convey the progress and assess the challenges in distributed asynchronous computing. The talks included in the minisymposium will cover the range of relevant topics; resilience and fault tolerance for distributed AMT programming models, asynchrony-tolerant numerical schemes and algorithms for stencil-based applications, task-based parallel simulations of multi-phase flows, and task-based runtimes utilizing directed acyclic graphs and a domain specific language for multi-physics simulations of turbulent reacting flows.

Towards Exascale Simulations of Particle-Laden Turbulence in a Radiation Environment: The PSAAP Program at Stanford

Hilario Torres (Stanford University, USA)

Wednesday, July 4, 2018

14:15 – 14:45
Darwin Room

In the framework of the Predictive Science Academic Alliance Program (PSAAP) the US Department of Energy is funding a Multidisciplinary Simulation Center at Stanford University to explore exascale computing strategies for multiphysics simulations. Stanford Center's research portfolio blends efforts in computer science, uncertainty quantization, and computational physics to tackle a challenging physical problem: the transfer of radiative energy to a turbulent mixture of air and solid particles. The context is provided by a relatively untested and poorly understood method of harvesting solar energy. The talk will describe the Center's effort to develop and validate a computational environment to simulate this challenging multi-physics problem emphasizing the strategies employed to carry out high-fidelity simulations and how uncertainty quantification techniques can be used to assess the overall performance of the system. A novel task-based programming system (Legion) is being deployed to tackle heterogeneous compute systems and retain portability and performance. Details of the implementation challenges and results obtained on various architectures will be discussed. The integration of large scale simulations and multi-level sampling for uncertainty analysis within the Legion framework will also be summarized.

Co-Author(s): **Gianluca Iaccarino** (Stanford University, USA)

A Scalable Asynchronous Computing Approach for Solving PDEs at Extreme Scale

Aditya Konduri (Sandia National Laboratories, USA)

Wednesday, July 4, 2018

14:45 – 15:15
Darwin Room

Synchronization overheads pose a major challenge as applications advance towards extreme scales. In this work, we present an asynchronous computing algorithm based on finite difference schemes for PDEs where no synchronization between processing elements (PEs) is enforced. PEs are allowed to continue computations regardless of messages status and are thus asynchronous. We show that accuracy of commonly used finite difference schemes is degraded when they are used asynchronously. Since message arrivals at PEs is essentially a random process, so is the behavior of the error. Within a statistical framework we show that average errors drop always to first-order regardless of the original scheme. The value of the error is found to depend on both grid spacing as well as characteristics of the computing system including number of PEs and statistics of the delays. We propose new asynchrony-tolerant schemes that maintain their accuracy under relaxed synchronization. We present results from the simulations of linear and non-linear PDEs, including reacting flow simulations, to demonstrate the feasibility of the method.

Co-Author(s): **Emmet M. Cleary** (California Institute of Technology, USA), **Diego A. Donzis** (Texas A&M University, USA), **Jacqueline H. Chen** (Sandia National Laboratories, USA)

Fault Tolerance in Asynchronous Many-Task (AMT) Programming Models and Runtimes

Hemanth Kolla (Sandia National Laboratories, USA)

Wednesday, July 4, 2018

15:15 – 15:45
Darwin Room

Among the many challenges faced by asynchronous many-task (AMT) programming models and runtimes, fault-tolerance is particularly daunting. The ability to allow asynchronous progress might appear to be favourable from a fault-tolerance perspective, since the overhead of recovering from a failed task could potentially be hidden. However, correctness and coherence requirements can become overwhelming for a poorly designed AMT runtime, more than off-setting any potential advantage of exploiting asynchrony. In this talk we present results from a systematic study of fault-tolerance for AMT systems. We establish that graph based analytical models are not tractable for the task-graphs of even the simplest applications. Accordingly, we present the design of, and results from, a task-graph simulator in which various aspects of an AMT system and its fault-tolerance are carefully parametrized. Simulator results of a stencil application task-graph are presented for various scenarios involving overdecomposition, failure rate, task scheduling and fault-tolerance strategy. The focus is particularly on two fault-tolerance strategies: task replication and task replay. Mock-up implementations of the stencil application, along with task replay and replication, in the shared-memory AMT system, HabaneroC++, are studied and compared with the simulator results.

Co-Author(s): **Keita Teranishi**, **Jackson Mayo**, **Rob Armstrong**, **Nicole Slattengren** (Sandia National Laboratories, USA)

Tools and Techniques to Enable Multiphysics Applications on Heterogeneous Architectures

James C. Sutherland (University of Utah, USA)

Wednesday, July 4, 2018

15:45 – 16:15
Darwin Room

Deploying multiphysics applications on heterogeneous architectures is particularly challenging because of the complexity and volume of code that must be maintained as well as the complex logic associated with the interplay between numerical algorithms and hardware. In this talk, we will explore some of the abstractions that we have found useful in developing simulation tools for turbulent reacting flows. This includes task-based runtime systems that utilize directed acyclic graphs together with a domain-specific language that provides simple syntax while preserving performance.

Co-Author(s): **Tony Saad** (University of Utah, USA)

MS44

Emerging Trends in Statistical Mechanics Applications to Nanostructured Materials

Organizer(s): ***Irina Paci***, ***Jeffrey Paci*** *(University of Victoria, Canada)*

The miniaturisation of devices towards the molecular scale is a dynamic field of research that has evolved dramatically in recent years. Its applications span a broad field, from the fabrication of smaller, more powerful computer chips, to charge storage, robotics, solar cells or biosensors. Complex functional materials are being envisioned and developed with these applications in mind, in an unprecedented effort towards rational design, through a combination of theory, computation and experiment.

In this context, the challenge for theory is to formulate laws and uncover patterns when the interactions that drive the relevant processes are as varied as the systems being investigated. Potential energy surfaces in nanoscale materials are complex, with multiple competing minima and steep barriers. The properties of the complex materials are also a challenge for computation: quantum effects in nanoscale interactions are reflected in the behaviour of the material as a whole, a statistical entity. This minisymposium will bring together researchers who have made significant contributions to the two essential challenges in materials simulations: sampling complex potential energy surfaces in structural prediction and bridging the relevant length scales in materials fabrication and properties. Current challenges and perspectives in method development and new applications will be discussed.

Bridging the Electronic, Atomistic and Mesoscopic Scales Using Machine Learned Models

Subramanian Sankaranarayanan (Argonne National Laboratory, USA)

Wednesday, July 4, 2018

14:15 – 14:45 Boston 3 Room

The ability to perform accurate calculations efficiently is crucial for computational materials design. In this talk, we will discuss our streamlined approach to force field development using first principles density functional theory training data and machine learning algorithms. Our objective has been to develop new, first-principles based, more accurate and more robust inter-atomic potentials for accurate simulations of dynamical processes at reactive interfaces and low dimensional systems such as clusters and molecules. The procedure involves several steps including (a) generation and manipulation of extensive fitting data sets through electronic structure calculations, (b) defining functional forms, (c) formulating novel highly optimized fitting procedures, (d) dual-Hamiltonian optimization to leverage classical FFs with more accurate approaches, and (e) subsequently coding and implementing these algorithms on high performance computers (HPCs). We will also discuss the validation of this approach on several diverse material systems ranging from precious metal nanocatalysts to newly discovered two dimensional materials such as stanene and silicene.

Metal and Metal-Oxide Clusters at Realistic Conditions: Beyond the Static, Monostructure Description

Luca M. Ghiringhelli (Fritz Haber Institute, Germany)

Wednesday, July 4, 2018

14:45 – 15:15 Boston 3 Room

The processes occurring at surfaces play a critical role in the manufacture and performance of advanced materials, e.g., electronic, magnetic, and optical devices, sensors, and catalysts. A prerequisite for analyzing and understanding the electronic properties and the function of surfaces is detailed knowledge of the atomic structure, i.e., the surface composition and geometry under realistic gas-phase conditions. The key quantity for studying the structure and function of surface/cluster in reactive atmospheres is the Gibbs free energy, as function of number of particles, pressure, and temperature. Here, I present a set of methods for the sampling of the configurational space of metal and metal-oxide clusters in reactive (e.g., O2, H2) atmosphere, in the canonical and grand-canonical ensembles, aiming at the unbiased determination of the phase diagrams as function of temperature and partial pressure of the reactive gas. Applications to gold, magnesium-oxide, and titanium-oxide nanoclusters, with first-principles potential-energy surfaces, will demonstrate the insight gained by the direct access to observables at finite temperature.

From Computational Spectroscopy to Artificial Water Splitting

Sandra Luber (University of Zurich, Switzerland)

Wednesday, July 4, 2018

15:15 – 15:45
Boston 3 Room

Spectroscopy is extremely helpful for the analysis of materials and chemical processes. In addition to experimental data, (spectroscopic) calculations provide important insight and allow the targeted study of specific structures, their dynamics and interactions. This is also true for artificial photosynthesis, which permits the splitting of water into molecular hydrogen and oxygen by means of solar light and is therefore a promising strategy to meet the increasing worldwide need for clean energy. Detailed investigation of the underlying mechanisms and factors determining the activity of catalysts is a prerequisite for the design of more efficient systems. We present our recent research for the development of forefront methods in spectroscopy based on high-performance dynamic ab initio methods as well the in-depth study and design of water splitting catalysis.

(i) Massively-Parallel Simulation of Self-Assembled Diblock-Copolymer Nano-Materials; (ii) Ab-Initio Quantum Monte Carlo Simulations for Single Vacancy Graphene and Isotropically-Strained Graphene

Ludwig Schneider (University of Göttingen, Germany),
Tomonori Shirakawa (SISSA, Italy)

Wednesday, July 4, 2018

15:45 – 16:15
Boston 3 Room

(i) Polymeric materials exhibit a rich equilibrium phase diagram, qualifying them for applications in electronic devices, filters, battery materials. Self-assembly of these materials rarely results in the equilibrium structures. Instead, configurations are trapped in long-lived meta-stable states. The properties of these structures, such as percolation and mechanical stability, can deviate from those of corresponding equilibrium phases. Investigating meta-stable states is challenging, due to finite size effects. SOMA, our massively-parallel implementation of the Single-Chain-in-Mean-Field algorithm, enables study of systems with billions of particles, unravelling the percolation characteristics of self-assembled diblock-copolymers as a function of volume fraction. (ii) Employing an ab-initio quantum Monte Carlo scheme (QMC), we will discuss the electronic structures of graphene in conditions where strong electron correlation plays an important role. Experimental studies of graphene have uncovered emerging spin-half free moments around vacancies. In QMC simulations of single vacancy graphene, we found a localized spin, composed of dangling sigma-orbitals around the vacancy. A model for the spin structure will be discussed. For isotropically strained graphene, we found evidence of two insulating phases before mechanical failure: a dimer phase where the structural dimerization implies the opening of a charge gap, and an antiferromagnetic phase induced by strong on-site Coulomb repulsion.

Co-Author(s): **Marcus Müller** (University of Göttingen, Germany)

MS45

Evolution of Knowledge Management in Astrophysics

Organizer(s): ***Roland Walter*** *(University of Geneva, Switzerland)*
Claudio Gheller *(ETH Zurich / CSCS, Switzerland)*

The volume of data generated by astrophysics observatories, and the complexity of knowledge management, has increased constantly and is now on the verge to reach even higher levels with observatories planned to generate 10-100s PB per year. This increase in volume, mostly driven by ground based telescopes with high resolution detectors and/or extreme time sampling, requires new computing infrastructures and new ways for scientists to interact with data and computing.

Observatory interfaces will need to move from data to analysis, interpretation and knowledge and from analysis to synthesis. Data mining, driven by the needs of science and education will require an integration of archive, pipeline and interpretation, which are still largely conceived as disconnected services. This evolution will benefit from the increased computing power and artificial intelligence which will help interpreting data flow exceeding human insight and will likely transform the astronomy business model.

The goal of this minisymposium is to clarify this evolution by comparing the needs of current and planned observatories and study how this implementation could take place both at the technical level and in the way scientists collaborate. The specific opportunities for Switzerland will also be discussed.

(i) Introduction; (ii) Gravitational-Wave Detector Data and Analysis Session Summary

Roland Walter (University of Geneva, Switzerland), **Maria Haney** (University of Zurich, Switzerland)

Wednesday, July 4, 2018

14:15 – 14:45
Osaka Room

(i) Introduction to the minisymposium from the organizers.
(ii) Following the breakthrough of the first direct detection of gravitational waves in 2015, the current generation of ground-based gravitational-wave detectors has observed gravity waves emitted by merging black holes and neutron stars. This summary talk will give an overview of the status of gravitational-wave data analysis and data management for the LIGO and Virgo instruments, and will provide an outlook to data and computing challenges for future generations of gravitational-wave observatories.

(i) Challenges in the Gaia Mission of the European Space Agency (ESA); (ii) Data Management at ESO and ALMA

Laurent Eyer (University of Geneva, Switzerland), **Felix Stoehr** (ESO/ALMA, Germany)

Wednesday, July 4, 2018

14:45 – 15:15
Osaka Room

(i) Gaia is a major mission of ESA whose aim is to survey our Galaxy. Gaia determines the position, distance, motion and derived physical properties of more than 1 billion stars with a growing set of measurements that already reached 120 billion after the first 22 months of observations. The data processing of this mission is one of the major challenges. A large part of the complexity of Gaia data comes from its intertwined nature. An international consortium of 500 people was established with the goal of providing astronomical catalogues publicly available. One of the duties of the consortium is to detect and analyse the variable celestial objects; this effort is coordinated by the University of Geneva with an associated data processing center, and is composed of 60 people. The tasks of this coordination unit is to classify the sources, and to analyse the different time series. Further specific analysis is done on a subset of sources to provide their astrophysical properties. In April 2018, we will output about 0.5 million variable sources and associated times series in the Gaia Data Release 2. We present the techniques developed for the variability classification and specific analyses, and solutions describing the data handling, workflow and soft knowledge management.
(ii) As technology as well as astronomy itself evolve, so do the challenges and opportunities related to Science Data Management. We give a high-level overview of some of the major challenges and outline some possible solutions and opportunities that emerge out of those challenges and place both into a larger context.

Co-Author(s): **Krzysztof Nienartowicz** (SixSQ, Switzerland), **Marc Audard** (University of Geneva, Switzerland)

(i) Data Management for the Cherenkov; (ii) Neutrino Telescope Data Management and Analysis

Etienne Lyard, **Teresa Montaruli** (University of Geneva, Switzerland)

Wednesday, July 4, 2018

15:15 – 15:45
Osaka Room

(i) The Cherenkov Telescope Array (CTA) will be the world's largest and most sensitive high-energy gamma-ray observatory. Composed of more than 100 telescopes of different sizes between 4 and 23 meters in diameter, it will operate from two separate sites in Chile and at the Canary Islands. It will generate up to 10PB of raw data per year that will be stored in a distributed archive. This talk will outline the current status, plans and challenges that we face to implement the analysis and data management pipeline of CTA.
(ii) We will describe the data flow structure of onsite DAQ to filtered steams for various physics scopes of IceCube and ANTARES and the plans for KM3NeT. The Data formats and data management software will also be described as well as plans for making data public.

Discussion Forum

Roland Walter (University of Geneva, Switzerland)

Wednesday, July 4, 2018

15:45 – 16:15
Osaka Room

Open discussion forum on knowledge management in astrophysics.

MS46

HPC beyond HEP: Opening Doors for New Data Intensive Sciences at Leadership Class HPCs Using BigPanDA

Organizer(s): ***Kaushik De*** *(The University of Texas at Arlington, USA)*
Alexei Klimentov *(Brookhaven National Laboratory, USA)*
Torre Wenaus *(Brookhaven National Laboratory, USA)*

The ATLAS experiment at CERN's Large Hadron Collider depends on the Worldwide LHC Computing Grid, the WLCG, for its remote computing infrastructure. PanDA, the workload management system used by ATLAS, annually processes over an exabyte of data using an average of 250,000 distributed batch slots, to enable hundreds of new scientific results. An effort was launched to extend PanDA, called BigPanDA, to access HPC resources, funded by the US Department of Energy (DOE-ASCR). Through this successful effort, ATLAS today uses about 20 million hours monthly on the Titan supercomputer at Oak Ridge National Laboratory. Many other supercomputers have also been integrated into ATLAS computing. This minisymposium will explore the software and operational lessons learned in integrating HPCs with traditional grid computing, and describe recent efforts to use BigPanDA for many other scientific domains. Three talks will summarize the state of the art and the future wishlist for HPC usage for current and future experiments, while a concluding expert panel discussion will focus on the future.

Enabling Biology, Chemistry and Other Sciences on Titan through BigPanDA

Danila Oleynik (The University of Texas at Arlington, USA)

Wednesday, July 4, 2018

14:15 – 14:45
Sydney Room

The Oak Ridge Leadership Computing Facility (OLCF) is one of the most powerful HPC centers available to researchers from different scientific fields to solve some of the world's most challenging scientific problems. Small scientific groups often need to develop expertise to optimize their applications for running on Titan, and to fit the usage policies of such big machines. We have installed the BigPanDA workload management system at OLCF to simplify the submission of user tasks to Titan. In this talk we will present results of an R&D project to execute workloads from different scientific groups at OLCF. We will describe all steps: starting from deployment of PanDA server as service on demand at OLCF in OpenShift containers, to the adaptation of PanDA client tools for new users. Examples from some of the different scientific fields using this service will include biology/genomics, molecular dynamics, LQCD, solid-state and neutrino physics, and different data science experiments: nEDM, LSST, and IceCube.

Co-Author(s): **Ruslan Mashinistov** (The University of Texas at Arlington, USA), **Pavlo Svirin**, **Sergey Panitkin** (Brookhaven National Laboratory, USA)

BigPanDA Experience on Titan for the ATLAS Experiment at the LHC

Alexei Klimentov (Brookhaven National Laboratory, USA)

Wednesday, July 4, 2018

14:45 – 15:15
Sydney Room

The PanDA software is used for workload management on distributed grid resources by the ATLAS experiment at the LHC. An effort was launched to extend PanDA, called BigPanDA, to access HPC resources, funded by the US Department of Energy (DOE-ASCR). Through this successful effort, ATLAS today uses over 25 million hours monthly on the Titan supercomputer at Oak Ridge National Laboratory. Many challenges were met and overcome in using HPCs for ATLAS simulations. ATLAS uses two different operational modes at Titan. The traditional mode uses allocations - which require software innovations to fit the low latency requirements of experimental science. New techniques were implemented to shape large jobs using allocations on a leadership class machine. In the second mode, high priority work is constantly sent to Titan to backfill high priority leadership class jobs. This has resulted in impressive gains in overall utilization of Titan, while benefiting the physics objectives of ATLAS. For both modes, BigPanDA has integrated traditional grid computing with HPC architecture. This talk will summarize the innovations to successfully use Titan for LHC physics goals.

Co-Author(s): **Kaushik De**, **Danila Oleynik** (The University of Texas at Arlington, USA), **Jack Wells** (Oak Ridge National Laboratory, USA), **Sergey Panitkin** (Brookhaven National Laboratory, USA)

BigPanDA: Blue Brain and Beyond

Shantenu Jha (Rutgers University, USA), **Fabien Delalondre** (EPFL, Switzerland)

Wednesday, July 4, 2018

15:15 – 15:45 Sydney Room

BigPanDA is a Workload Management System designed to support the execution of experimental workloads and workflows on distributed resources. In the first part of this talk, we discuss a "proof of concept" project that was stated in 2017 and conducted jointly by the BigPanDA team and the Blue Brain Project (BBP) of the Ecole Polytechnique Federal de Lausanne (EPFL). This proof of concept project showed the efficient application of the BigPanDA system to support the complex scientific workflow of the BBP using a mix of desktop, cluster and supercomputers to reconstruct and simulate accurate models of brain tissue. In the second part of this talk, we will discuss how the next generation task execution layer of BigPanda, known as the NGE (Next Generation Executor) is being enhanced for upcoming pre- and full exascale systems. We will discuss the design and integration of NGE with the BigPanda system (including Harvester), demonstrate its native support for MPI tasks (and thereby of immediate relevance to the BBP) and characterize its performance.

Panel: BigPanDA Experience at Oak Ridge - Learning from the LHC, Going Far Beyond

Torre Wenaus (Brookhaven National Laboratory, USA)

Wednesday, July 4, 2018

15:45 – 16:15 Sydney Room

In recent years, the ATLAS experiment at the Large Hadron Collider has been very successful in using HPCs as an integral part of their computing operations. Projects pioneered on Titan have led the way for a new pattern of computing usage for the LHC. The success of the LHC program has also led to new initiatives to try the same tools and practices for other data sciences. A panel of experts in HPC and Distributed Computing will lead a discussion about what the success at Titan means for the future of Exascale HPCs and for data science communities.

Co-Author(s): **Kaushik De** (The University of Texas at Arlington, USA), **Frank Wuerthwein** (UC San Diego, USA), **Jack Wells** (Oak Ridge National Laboratory, USA), **Alexei Klimentov** (Brookhaven National Laboratory, USA), **Vladimir Korenkov** (JINR, Russia), **Simone Campana** (CERN, Switzerland), **Shantenu Jha** (Rutgers University, USA)

MS47

HPUQ: Current Challenges in Uncertainty Quantification for Mechanistic Models, Part II: Applications in Life Sciences and Engineering

Organizer(s): ***Jonas Šukys*** *(Swiss Federal Institute of Aquatic Science and Technology, Switzerland)*
Panagiotis Hadjidoukas *(ETH Zurich, Switzerland)*
Antonietta Mira *(Università della Svizzera italiana, Switzerland)*

The HPUQ minisymposium focuses on uncertainty quantification (UQ) of mechanistic models for natural sciences (eg. Engineering, Life and Aquatic Sciences) using high performance computing (HPC). The statistical inference (e.g. calibration) of models for complex mechanistic models, in the abundance of data arriving from heterogeneous sources, poses a methodological and computational challenge for scientists. In the first session, the minisymposium highlights cutting edge frameworks for rigorous and robust UQ as ABCpy, ∏4U, PyMLMC, SPUX to address these issues, with a focus towards optimal algorithmic performance and efficient utilization of HPC resources. In the second session of the minisymposium, we shift the focus to the applications of UQ methodologies in several important scientific domains spanning from Biomedicine and Biomechanics to Aerospace Engineering and Fluid Dynamics.

Combining Clinical Observations, Mathematical Modeling and HPC Approximate Bayesian Computation for Developing New Diagnosis Techniques

Bastien Chopard (University of Geneva, Switzerland)

Wednesday, July 4, 2018

14:15 – 14:45
Singapore Room

Platelets are transported by blood and play an important role in several physiological processes. Dysfunction of platelets deposition on a vessel wall can reveal several pathologies. In this talk we present a new mathematical model describing how platelets deposit. By combining clinical observations, our mathematical model and HPC Approximate Bayesian computations, we are able to obtain patient specific data that cannot be otherwise measured. These data can provide clinicians with new and more significant information than current platelet function tests.

Co-Author(s): **Ritabrata Dutta** (Università della Svizzera italiana, Switzerland), **Karim Zouaoui Boudjeltia** (Université libre de Bruxelles, Belgium)

Reverse Engineering of Tendons: The Data Conundrum and Current Computing Challenges

Nikolaos Karathanasopoulos (ETH Zurich, Switzerland)

Wednesday, July 4, 2018

14:45 – 15:15
Singapore Room

Tendons are hierarchical, collagen-based tissues that transfer mechanical load from muscles to bones. Their malfunction or failure affects the mobility of millions of people. Tendons exhibit a multiscale inner architecture which plays an important role in their overall mechanical behavior. Yet, existing experimental observations commonly restrain to a single inner scale, providing highly uncertain and often contradictory information. In order to infer the most probable tendon inner structure, we construct computational schemes that are descriptive of the composite, multiscale tendon inner architecture, which we couple with data through a Bayesian analysis framework. Thereupon, we compute inner material designs that optimally fit the elastic and volumetric behavior of tendons. Carrying out the inference process requires the use of sophisticated enough models to adequately describe the tendon mechanics, a prerequisite that can bring current high performance computing capabilities to its limits.

Combined Error and Uncertainty Bound Estimates with Application to CFD Problems

Timothy Barth (NASA, USA)

Wednesday, July 4, 2018

15:15 – 15:45
Singapore Room

Hydrodynamic realizations often contain numerical error arising from finite-dimensional approximation (e.g. numerical methods using grids, basis functions, particles, etc) and statistical uncertainty arising from incomplete information and/or statistical characterization of model parameters and random fields. We present a general framework for uncertainty and error quantification that includes the combined effects of statistical uncertainty and numerical error so that uncertainty statistics with *a posteriori* error bound estimates are provided. For problems containing no sources of uncertainty, a standard error bound estimate is obtained. For problems containing no numerical error, a standard uncertainty estimate is obtained. Specifically, we consider error bounds for moment statistics given realizations containing finite-dimensional numerical error (Barth, 2013). The error in computed output statistics contains contributions from both realization error and the error resulting from the calculation of statistics integrals using a numerical method. We then devise computable *a posteriori* error bounds by numerically approximating all terms arising in the error bound estimates for a variety of standard UQ methods: Multi-level dense quadrature basis methods (Tatang, 1994); Multi-level sparse quadrature basis methods (Smolyak, 1963) utilizing node-nested hierarchies; Multi-level sampling methods (Mishra and Schwab, 2010) for high-dimensional random variable spaces. HPC CFD calculations are shown to demonstrate features of the framework.

Using One Thousand GPUs to Understand the Euler Equations

Kjetil Lye (ETH Zurich, Switzerland)

Wednesday, July 4, 2018

15:45 – 16:15
Singapore Room

The question of well-posedness for the compressible Euler equations is still open. Recent theoretical and numerical evidence have indicated that solutions of the Euler equations exhibit random behavior, even with deterministic initial data. We use the framework of statistical solutions to model this inherit randomness. We review the theory of statistical solutions for conservation laws. Afterwards, we introduce a convergent numerical method for computing the statistical solution of conservation laws. The proposed numerical method needs a fast numerical simulation of the Euler equations to be feasible, which is why we present the Alsvinn simulator: A fast, modular, CPU- and GPU-based finite volume simulator using modern programming techniques. The simulator will be available as free open source software. Using the Alsvinn simulator, together with resources from CSCS, we have performed large scale experiments using the strength of the CSCS Daint cluster. In the last part of the talk, we show the results of the numerical simulations to highlight why high performance computing can be useful for mathematical understanding of the Euler equations.

MS48

Unconventional Methods for Partial Differential Equations

Organizer(s): ***Wesley P. Petersen*** *(ETH Zurich, Switzerland)*

This minisymposium will explore some unconventional methods for numerical simulations of partial differential equations. Often, discretizations of PDEs produce regular error patterns, some of which can be both quantified and smoothed by adding small stochastic terms. In fact, stochastic processes can be intimately connected with PDEs. For example, the heat equation is just a PDE form for the distribution of Brownian random motion. More general frameworks can be built around the Feynman-Kac formula (e.g. Mark Freidlin's book on functional integrals, Princeton Univ. Press, 1985). Such formulations are both flexible and intrinsically parallelizable (e.g. W. Petersen and P. Arbenz, Oxford Univ. Press, 2004). In high spacial dimensions, such techniques, when formulated as Monte-Carlo methods, are much more accurate than might be expected and make many difficult high dimensional simulations even possible. In, say, 3-D, certain interactions (e.g. foams, bubbles) can be modelled phenomenologically where details about these interactions are not well understood. In addition, we will discuss balancing methods or systems in external fields, and the connections between particle methods and the PDEs which describe the distributions.

High-Order Well-Balanced Finite Volume Methods for Euler Equations with Gravity

Luc Grosheintz-Laval (ETH Zurich, Switzerland)

Wednesday, July 4, 2018

14:15 – 14:45
Nairobi Room

In this talk I will present high-order, well-balanced finite volume methods (FVM) for the Euler equations with gravity. The Euler equations are a system of hyperbolic PDEs commonly used to describe inviscid compressible gases. They admit stationary steady-state solutions, known as hydrostatic equilibria. These equilibria arise from a force balance of gravity forces and pressure gradients. A variety of interesting natural phenomena occur approximately in hydrostatic equilibrium. Examples include numerical weather/climate prediction on earth or exoplanets and convection in starts, among others. Textbook FVM generally do not preserve these equilibria exactly (i.e. to machine precision). Motivated by the given examples I will present high-order FVM which preserve the hydrostatic equilibria to machine precision without making any assumptions on the equation of state (EOS), which could be a tabulated EOS, or the gravitational potential, which could be the numerical solution of a Poisson equation and only given at certain point-values.

Splitting Methods for ODEs, PDEs, and SDEs - with Examples

Wesley P. Petersen (ETH Zurich, Switzerland)

Wednesday, July 4, 2018

14:45 – 15:15
Nairobi Room

Evolution equations are usually written with the time derivative of a process on the left hand side and multiple additive pieces on the right. Often, each of the additive pieces would permit relatively easy to compute approximations if taken alone. Taken altogether, only low order or less stable approximations are possible. Splitting methods use the approximations for the solutions of individual pieces with compositions of these to construct higher order methods with desired stability properties. Such splittings are very general: ordinary differential equations (Yoshida methods), partial differential equations (Godunov dimensional splittings), and Ito stochastic differential equations. This talk will show some formal compositions along with several numerical examples. These examples will be a Trotter-type anharmonic oscillator approximation, the solution of the Fisher/KPP equation on a terrestrial map, and the simulation of an oscillating stochastic differential equation. Two important issues will emphasized: the connections between stochastic differential equations and diffusion processes, and the parallel computing aspects of these simulations.

Mutual Impact of Bubbles and Waves Studied with an Efficient Finite Volume Solver

Fabian Wermelinger (ETH Zurich, Switzerland)

Wednesday, July 4, 2018

15:15 – 15:45
Nairobi Room

While the finite volume method is a well established technique for solving PDEs, most of the existing implementations suffer from low utilization of computing resources limiting its application for large-scale problems. Finite volume solver Cubism-MPCF implements a modern algorithm for two-phase compressible flows based on a Godunov-type scheme with WENO reconstruction. Designed for high performance and scalability, the solver reached performance of 11 PFLOP/s on Sequoia supercomputer. Present work combines applications of Cubism-MPCF to various phenomena including cavitation, shock-induced collapse and acoustics of bubbly liquids. Study of cavitation considers the collapse of a cluster of gas bubbles due to increased pressure. The large number of bubbles of about 12000, for the first time allowed to describe the collective behaviour such as propagation of the collapse front together with the evolution of microjets formed near individual bubbles. Detailed simulations of the shock-induced collapse cover a wide range of physical parameters revealing cases for which the effects of viscosity and surface tension become significant but ignored in other studies. Another application demonstrates Anderson localization of acoustic waves in bubbly liquids. Moreover, interaction with a standing wave leads to rearrangement and deformation of bubbles. These effects are unavailable for linearized models commonly used in acoustics.

Co-Author(s): **Petr Karnakov** (ETH Zurich, Switzerland)

Fight Uncertainty with Randomness: Stochastic Particle Methods for Microfluidics

Lucas Amoudruz (ETH Zurich, Switzerland)

Wednesday, July 4, 2018

15:45 – 16:15
Nairobi Room

In a porous medium, a fluid pushed inside a more viscous one propagates in a finger-like structure. This phenomenon is of great interest in the fields of oil recovery or pollutant spreading in ground water. These so-called "fingering instabilities" depend on the viscosities, densities and surface tensions of the fluids. Such flows have been widely studied inside Hele-Shaw cells, which serve as an experimental platform for studying fundamental flow patterns in constricted geometries. It consists in two parallel plates separated by a small gap, in which the fluids are flowing under gravity or applied pressure gradient. Recent experiments exhibited surprising stability conditions in the limit of zero surface tension between the fluids. They report an extended stability region in terms of viscosity ratios, which is not predicted by linear stability analysis. Reproducing these results numerically is a crucial step towards understanding such complex flow patterns. We present numerical simulations of multicomponent fluids flowing inside Hele-Shaw cells. The simulations employ the Dissipative Particle Dynamics method, a stochastic particle method widely used in microfluidic applications.

Poster Sessions

Flash Poster Session

Tuesday, July 3, 2018
10:15 – 11:00
Montreal Room
Chair: **Maria Grazia Giuffreda** (ETH Zurich / CSCS, Switzerland)

Poster Session & Reception

Tuesday, July 3, 2018
19:30 – 21:30
Foyer 2nd Floor

Chemistry & Materials

CHM-01
Accurate and Efficient Molecular Dynamics with Nuclear Quantum Effects
Venkat Kapil, Michele Ceriotti (EPFL, Switzerland), Joost VandeVondele (ETH Zurich / CSCS, Switzerland)

CHM-02
AiiDA: A Simulation Platform with Full Provenance Support and Flexible Workflows
Spyros Zoupanos, Leonid Kahle, Sebastiaan Huber, Martin Uhrin, Nicolas Mounet, Rico Andreas Häuselmann, Snehal Kumbhar, Leopold Talirz, Fernando Gargiulo, Nicola Marzari, Giovanni Pizzi (EPFL, Switzerland), Andrea Cepellotti (UC Berkeley, USA), Andrius Merkys (Vilnius University, Lithuania), Boris Kozinsky (Harvard University, USA)

CHM-03
Bridging the Gap between Atomistic and Macroscopic Models of Homogeneous Nucleation
Bingqing Cheng, Michele Ceriotti (EPFL, Switzerland)

CHM-04
The Crucial Role of the Hydrogen Bonding Network in Water Oxidation Catalyzed by a Cobalt-Cubane
Mauro Schilling, Sandra Luber (University of Zurich, Switzerland)

CHM-05
Datamining of Magnetic Double Perovskites
Michele Visciarelli, Thor Wikfeldt, Anna Delin (KTH Royal Institute of Technology, Sweden)

CHM-06
Development of a Modular API for Computation of Non-Bonded Interactions in Particle Simulations
Prashanth Kanduri, Victor Holanda Rusu (ETH Zurich / CSCS, Switzerland)

CHM-07
DFT+U Gamma-Surfaces of UO2
Monica Kosa, Raoul Ngayam Happy, Matthias Krack (Paul Scherrer Institute, Switzerland), Sébastien Groh (University of Basel, Switzerland)

CHM-08
Improving the Performance of the DBCSR Library for Sparse Matrix Multiplication for Many-Core and GPU Computing Systems
Andreas Gloess, Alfio Lazzaro, Juerg Hutter, Tiziano Mueller, Patrick Seewald, Ilia Sivkov (University of Zurich, Switzerland)

CHM-09
Materials Cloud: A Platform for Open Materials Science
Giovanni Pizzi, Leopold Talirz, Snehal Kumbhar, Fernando Gargiulo, Marco Borelli, Elsa Passaro, Aliaksandr Yakutovich, Nicola Marzari (EPFL, Switzerland), Ole Schütt (Empa, Switzerland), Joost VandeVondele, Thomas Schulthess (ETH Zurich / CSCS, Switzerland)

CHM-10
A Symmetry-Adapted Approach to Machine Learning of Tensors
David M. Wilkins, Andrea Grisafi, Michele Ceriotti (EPFL, Switzerland)

Climate & Weather

CLW-01
Automatic Optimization of Domain Specific Languages for Weather and Climate Models
Tobias F. Wicky, Fabian Thuering, Torsten Hoefler (ETH Zurich, Switzerland), Carlos E. Osuna, Oliver Fuhrer (MeteoSwiss, Switzerland)

CLW-02
ESiWACE: Performance Predictions for Storm-Resolving Simulations of the Climate System
Joachim Biercamp, Philipp Neumann (German Climate Computing Center, Germany)

CLW-03
Experiments with Containerising a State-of-the-Art Weather and Climate Model for Application in HPC
Simon Wilson, Bryan Lawrence (NCAS-CMS, UK)

CLW-04
Performance Study of Climate and Weather Models: Towards a More Efficiently Scalable Model
Mario Acosta, Xavier Yepes, Oriol Tinto, Kim Serradell, Miguel Castrillo (Barcelona Supercomputing Center, Spain)

CLW-05
Porting Physical Parameterizations from a Climate Model to Accelerators
Thomas Köster, Olaf Schenk (Università della Svizzera italiana, Switzerland), Gerhard Wellein (Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany), William Sawyer (ETH Zurich / CSCS, Switzerland), Xavier Lapillonne (MeteoSwiss, Switzerland)

CLW-06
Towards More Efficient Adjustment of Free Parameters in a Global Climate Model?
Doris Folini, Martin Wild (ETH Zurich, Switzerland)

Computer Science & Applied Mathematics

CSM-01
Accelerating Life Science Notebook Applications: Architectural Issues and Use Cases
Antonio Maffia, Helmar Burkhart (University of Basel, Switzerland), Gang Mu (Roche, Switzerland)

CSM-02
Adaptive Grid Refinement Techniques for Particulate Flow Simulations with the Lattice Boltzmann Method
Christoph Rettinger, Ulrich Rüde (Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany)

CSM-03
Are Smooth Particle Hydrodynamics Applications Inherently Resilient to Faults?
Aurelien Cavelan, Florina Ciorba, Rubén Cabezon (University of Basel, Switzerland)

CSM-04
Balanced Graph Partition Refinement Using the Graph p-Laplacian
Dimosthenis Pasadakis, Drosos Kourounis, Toby Simpson, Olaf Schenk (Università della Svizzera italiana, Switzerland), Kohei Fujita, Takuma Yamaguchi, Tsuyoshi Ichimura (The University of Tokyo, Japan)

CSM-05
BioMedIT: Enabling Interoperable Biomedical Analysis
Jaroslaw Surkont, Thierry Sengstag (University of Basel, Switzerland), Kevin Sayers, Ioannis Xenarios, Torsten Schwede (Swiss Institute of Bioinformatics, Switzerland), Bernd Rinn (ETH Zurich, Switzerland), Marcel Riedi (University of Zurich, Switzerland)

CSM-06
A Distributed Parallel Approach for Large Scale Optimal Power Flow with Security Constraints
Juraj Kardos, Drosos Kourounis, Olaf Schenk (Università della Svizzera italiana, Switzerland)

CSM-07
Evaluating OpenACC on a Large Scale Particle Simulation
Samuel Adolfo Cruz Alegría, Alessandra Martha de Felice, Hrishikesh Gupta, Juraj Kardos (Università della Svizzera italiana, Switzerland)

CSM-08
Evaluating TensorFlow Optimization Techniques for Solving Elliptic Boundary Control Problems
Manav Choudhary, Olaf Schenk (Università della Svizzera italiana, Switzerland)

CSM-09
High Performance Topology Optimization
Ezekiel Barnett, Sameer Rawat, Sumeet Gyanchandani, Dimosthenis Pasadakis (Università della Svizzera italiana, Switzerland)

CSM-10
HPC-as-a-Service for Driving Artificial Intelligence for Drug Discovery
Vojtech Cima, Jan Martinovi, Vaclav Svaton (IT4Innovations National Supercomputing Center, Czech Republic), Nina Jeliazkova, Vedrin Jeliazkov (Ideaconsult Ltd., Bulgaria), Vladimir Chupakhin (Janssen Pharmaceutica NV, Belgium)

CSM-11
Importance of Rank Reordering for Advanced Polar Decomposition Algorithms
Aniello Esposito (Cray Inc., UK), David Keyes, Hatem Ltaief, Dalal Sukkari (King Abdullah University of Science and Technology, Saudi Arabia)

CSM-12
Learning Method for Asynchronous, Distributed Reinforcement Learning
Timon Willi, Vihang Patil, Radim Janalik (Università della Svizzera italiana, Switzerland)

CSM-13
Neuronal Network Simulation Code for the Exascale Era
Susanne Kunkel (Norwegian University of Life Sciences, Norway), Jakob Jordan (University of Bern, Switzerland), Tammo Ippen, Moritz Helias, Markus Diesmann (Forschungszentrum Jülich, Germany), Itaru Kitayama, Mitsuhisa Sato, Jun Igarashi (RIKEN, Japan)

CSM-14
A New Community-Driven Resource for Scientific Software Improvement Exchange
Anshu Dubey, Lois C. Mclnnes (Argonne National Laboratory, USA), David E. Bernholdt (Oak Ridge National Laboratory, USA), Michael A. Heroux (Sandia National Laboratories, USA)

CSM-15
ORCA and Cut-and-Solve: A Potential High-Performance Solution to Learning Genetic Causes of Complex Diseases
Michael Chan, Sharlee Climer, Sanjiv K. Bhatia (University of Missouri - St. Louis, USA), Daniel Jacobson (Oak Ridge National Laboratory, USA), Carlos Cruchaga (Washington University School of Medicine, USA)

CSM-16
Parallelization of the Boundary Element Method
Michal Merta, Jan Zapletal, Michal Kravcenko (IT4Innovations National Supercomputing Center, Czech Republic)

CSM-17
Performance and Implementation of a Geometric Multigrid Solver with Trilinos
Matthias Frey, Andreas Adelmann (Paul Scherrer Institute, Switzerland)

CSM-18
Performance Evaluation of Dynamic Loop Scheduling Techniques Using MPI Passive RDMA on Distributed Memory Systems
Ahmed Eleliemy, Florina Ciorba (University of Basel, Switzerland)

CSM-19
Performance Evaluation of Nodal Discontinuous Galerkin Methods on Many-Core Architecture
Shivasubramanian Gopalakrishnan, Mandar Gurav (Indian Institute of Technology Bombay, India)

CSM-20
Practical Communication-Optimal Algorithm for Dense Matrix-Matrix Multiplication
Marko Kabic, Thibault Notargiacomo, Joost VandeVondele (ETH Zurich / CSCS, Switzerland)

CSM-21
Practical Experience with Task-Based Programming Techniques for Quantum Chemistry Software
Heike Jagode, Anthony Danalis, Jack Dongarra (University of Tennessee, USA)

CSM-22
Redesigning Numerical Modelling Algorithms for Efficient, Large-Scale Cloud Deployment
James W. D. Hobro, Anindya Sharma (Schlumberger, UK)

CSM-23
Software-Defined Events through PAPI for In-Depth Analysis of Application Performance
Anthony Danalis, Heike Jagode, Jack Dongarra (University of Tennessee, USA)

CSM-24
A Study of the Performance of Scientific Applications with Dynamic Loop Scheduling under Perturbations
Ali Mohammed, Florina Ciorba (University of Basel, Switzerland)

CSM-25
Towards an Exascale-Ready Mini-App for Smooth Particle Hydrodynamics
Danilo Guerrera, Florina Ciorba, Rubén Cabezon, Aurélien Cavelan (University of Basel, Switzerland), Lucio Mayer, Darren S. Reed (University of Zurich, Switzerland), David Imbert (Nextflow Software, France), Jean-Guillaume Piccinali (ETH Zurich / CSCS, Switzerland), Ioana Banicescu (Mississippi State University, USA), Domingo Garciá-Senz (Universitat Politècnica de Catalunya, Spain), Thomas R. Quinn (University of Washington, USA)

CSM-26
Towards Whole Program Generation for Ocean Modeling
Sebastian Kuckuk, Harald Köstler (Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany)

CSM-27
Using Data Analysis Techniques to Detect Ransomware
Sushma Yellapragada, Upasna Sharma, Abhishek Barry (The Northcap University, India)

CSM-28
Utopia: A High Performance C++ Embedded Domain Specific Language for Scientific Computing
Nur Aiman Fadel, Andreas Fink (ETH Zurich / CSCS, Switzerland), Patrick Zulian, Rolf Krause (Università della Svizzera italiana, Switzerland), Dimitrios Karvounis (ETH Zurich, Switzerland)

CSM-29
Validation of the Self-Adaptive Navigation System by Enhanced HPC Traffic Simulator
Vit Ptosek, Jan Martinovič, Jiri Sevcik, Katerina Slaninova (IT4Innovations National Supercomputing Center, Czech Republic), Radim Cmar (Sygic, Slovakia)

Emerging Application Domains

EAD-01
On the Solution of Macroeconomic Models with Distributional Channels and Default Risk
Luca Mazzone (University of Zurich, Switzerland)

Engineering

ENG-01
Adaptive Particle Representation: A Novel Framework for Adaptive-Resolution Simulations
Suryanarayana Maddu, Bevan L. Cheeseman (Max Planck Institute for Molecular Cell Biology and Genetics, Germany), Pietro Incardona, Ivo F. Sbalzarini (Centre of System Biology Dresden, Germany)

ENG-02
Assessment of Detached Eddy Simulation in Predicting Separated Flow over Airfoils at a Moderate Reynolds Number
Ramesh Balakrishnan, Jun Fang (Argonne National Laboratory, USA), Kenneth Jansen (University of Colorado Boulder, USA), Philipp Schlatter, Ricardo Vinuesa (KTH Royal Institute of Technology, Sweden), Michel Rasquin (Cenaero, Belgium)

ENG-03
Large Eddy Simulation of Tsunami Triggered Coastal Inundation in the Presence of Mitigation Parks
Simone Marras (New Jersey Institute of Technology, USA), Jenny Suckale, Yilang Xu (Stanford University, USA), Beatriz Eguzkitza, Guillaume Houzeaux, Mariano Vázquez (Barcelona Supercomputing Center, Spain)

ENG-04
A Low-Mach Simulation of Flow and Heat Transfer in a Motored Internal Combustion Engine Using the Spectral Element Method
Saumil Patel, Misun Min (Argonne National Laboratory, USA), Georgios Giannakopoulos, Christos Frouzakis, Konstantinos Boulouchos (ETH Zurich, Switzerland), Ananias Tomboulides (Aristotle University of Thessaloniki, Greece), Paul Fischer (University of Illinois Urbana-Champaign, USA)

ENG-05
OpenFPM for Scalable Particle-Mesh Simulations on Distributed-Memory Computers
Pietro Incardona, Ivo Sbalzarini (Centre of System Biology Dresden, Germany)

ENG-06
Wavelet Based Data Compression Strategies for Exascale CFD Simulations
Patrick Vogler, Ulrich Rist (University of Stuttgart, Germany)

Life Sciences

LIF-01
AI-GWAPA: Explainable-AI and Genome Wide Association Phytobiome Analysis
Piet Jones, Ka Hung Lee (University of Tennessee, USA), Benjamin Garcia, Stephan Irle, Udaya Kalluri, Wellington Muchero, Jay Chen, Gerald Tuskan, Daniel Jacobson (Oak Ridge National Laboratory, USA)

LIF-02
The Bromodomain-Peptide (Un)Binding Network
Cassiano Langini, Marco Bacci, Andreas Vitalis, Amedeo Caflisch (University of Zurich, Switzerland)

LIF-03
Explainable Machine Learning for Systems Biology: Tensor Iterative Random Forests
Jonathon Romero, Ashley Cliff (University of Tennessee, USA), Doug Hyatt, Daniel Jacobson (Oak Ridge National Laboratory, USA), Ben Brown (Lawrence Berkeley National Laboratory, USA)

LIF-04
iReceptor: A Platform for Exploring and Analyzing Antibody/B-cell and T-cell Receptor Repertoire Data across Federated Repositories
Brian Corrie, Felix Breden (Simon Fraser University, Canada)

LIF-05
Modeling Biological Networks with Exponential Random Graph Models
Alex Stivala (Swinburne University of Technology, Australia), Maksym Byshkin, Antonietta Mira, Alessandro Lomi (Università della Svizzera italiana, Switzerland), Garry Robins (The University of Melbourne, Australia)

LIF-06
Parallel Kraken for Meta-Omic Microbiome and Phytobiome Classification
Benjamin Garcia, Doug Hyatt, Daniel Jacobson (Oak Ridge National Laboratory, USA), Piet Jones (University of Tennessee, USA), Ian Hodge (Stanford University, USA)

LIF-07
SAPPHIRE Basin Recognition: An Unsupervised Algorithm to Identify and Project Metastable and Transition States in High-Dimensional Time Series Data
Francesco Cocina, Marco Bacci, Andreas Vitalis, Amedeo Caflisch (University of Zurich, Switzerland)

Physics

PHY-01
AFiD-GPU: A Versatile Navier-Stokes Solver for Wall-Bounded Turbulent Flows on GPU Clusters
Xiaojue Zhu, Vamsi Spandan, Detlef Lohse, Richard Stevens (University of Twente, Netherlands), Everett Phillips, Massimiliano Fatica, Gregory Ruetsch, Josh Romero (NVIDIA Inc., USA), John Donners (SURFsara, Netherlands), Rodolfo Ostilla-Mónico (University of Houston, USA), Yantao Yang (Peking University, China), Roberto Verzicco (University of Rome Tor Vergata, Italy)

PHY-02
Hydrodynamical High Performance Simulations of an Accretion Disk surrounding a Supermassive Black Hole and its Interactions
Fabian Klein, Rainer Spurzem, Andreas Just (University of Heidelberg, Germany), Rolf Kuiper (Universität Tübingen, Germany)

PHY-03
A Performance Model for Quantum ESPRESSO's PWscf
Pietro Bonfà, Fabio Affinito, Carlo Cavazzoni (CINECA, Italy)

PHY-04
Training Deep Learning Models on Many-Core Processors
David Ojika (University of Florida, USA)

Solid Earth Dynamics

SED-01
Extreme Scale Global Convection Models for Flow-Induced Topography
Simon Bauer, Hans-Peter Bunge, Siavash Ghelichkhan, Marcus Mohr (Ludwig Maximilian University of Munich, Germany), Markus Huber, Barbara Wohlmuth (TU Munich, Germany), Ulrich Rüde (Friedrich-Alexander-Universität Erlangen Nürnberg, Germany)

Exhibitors

PASC18 includes exhibition stands from local and global organizations engaged in activities related to computational science and HPC. PASC18 participants have the opportunity to get acquainted with research and services from the following organizations:

- **BioExcel** – Centre of Excellence for Computational Biomolecular Research
- **CADMOS (UNIGE-EPFL-UNIL)** – Centre for Advanced Modeling Science
- **CSZ** – Computational Science Zurich
- **E-CAM** – European Centre of Excellence
- **HPE** – Hewlett Packard Enterprise
- **ICS** – Institute of Computational Science of the Università della Svizzera italiana
- **Microsoft Switzerland Ltd.**
- **Swiss National Supercomputing Centre (CSCS) of ETH Zurich**


NAVIER-STOKES EQUATION
Lugano
Switzerland
Association for
Computing Machinery


7
6-28 June 2017
ghpc
cscs

Industry sponsors

Platinum sponsorship

Gold sponsorship

Silver sponsorship

Institutional partners

Media partners

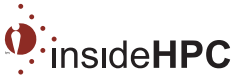

pasc18.pasc-conference.org

NAVIER-STOKES EQUATION

$$\rho\left(\frac{\partial \mathbf{v}}{\partial t} + \mathbf{v}\cdot\nabla\mathbf{v}\right) =$$

$$= -\nabla p + \mu\nabla^2\mathbf{v} + \mathbf{f}$$

$\mathbf{f}$ = body forces (gravity or centrifugal)

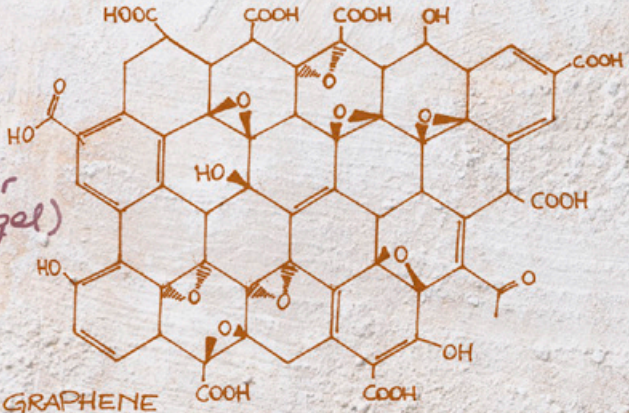

GRAPHENE OXIDE

METROPOLIS ALGORITHM

initialize x_1, n and s
for $i = 1 : (n-1)$ do
  while x_{i+1} not assigned do
    draw $z \in [0,1]$ and $u_i \in [-1,1]^d$
    $x_{new} = x_i + u_i s$
    if $f(x_{new})/f(x_i) \geq z$ then $x_{i+1} = x_{new}$
  end while
end for