



UNIVERSITY
OF WARSAW



March 11-14, 2019

CONFERENCE REPORT

CONTENTS

03

ABOUT
THE ORGANISER

04

ABOUT
THE CONFERENCE

05

SCIENTIFIC COMMITTEE

06

ORGANISING COMMITTEE

07 - 09

CONFERENCE

10

WORKSHOPS

13

ACKNOWLEDGEMENT

14

SPONSORS





ABOUT THE ORGANISER

Interdisciplinary Centre for Mathematical and Computational Modelling (ICM), University of Warsaw was established by a resolution of the Senate of the University of Warsaw dated 29 June 1993, the Interdisciplinary Centre for Mathematical and Computer Modeling (ICM), University of Warsaw, is one of the top HPC centres in Poland.

Supercomputing Frontiers Europe will be held in Warsaw, the capital of Poland. It is the largest city in Poland and the centre of economic, political and cultural events. Warsaw is a metropolis full of life and, at the same time, a city with a unique history. A quarter of its surface is occupied by green spaces, it is also the centre of culture for every taste and every budget.

The conference venues are:

Main conference (11-13 March):
Copernicus Conference Centre



Workshops (14th March):
University of Warsaw Library





ABOUT SCFE19

Supercomputing Frontiers is an annual international conference that provides a platform for thought leaders from both academia and industry to interact and discuss visionary ideas, important visionary trends and substantial innovations in supercomputing.

The first three editions of the conference were organised in Singapore by A*STAR Computational Resource Centre and National Supercomputing Centre Singapore, the fourth — by Interdisciplinary Centre for Mathematical and Computer Modeling (ICM), University of Warsaw, Poland.

Previous editions of the conference showcased a successful scientific programme with eminent plenary keynote speakers such as:

- Thomas Sterling,
- Satoshi Matsuoka,
- Robert Harrison,
- Rick Stevens,
- Jack Dongarra,
- Alan Gara,
- John Gustafson,
- Srinivas Aluru,
- Baroness Susan Greenfield,
- Horst Simon,
- Bronis R. De Supinski,
- Gordon Bell,
- Alessandro Curioni,
- Thom H. Dunning Jr.
- and Haohuan Fu.

Selected papers from the conference were published in the international journal *Supercomputing Frontiers & Innovations* in 2015, 2016 and 2017.

SCIENTIFIC COMMITTEE

PROGRAMME CHAIR

Marek Michalewicz, ICM University of Warsaw

MEMBERS

Pete BECKMAN

Argonne National Laboratory, US

Taisuke BOKU

University of Tsukuba, Japan

Alessandro CURIONI

IBM Research – Zurich, Switzerland

Maciej CYTOWSKI

Pawsey Supercomputing Centre, Australia

Vassil DIMITROV

University of Calgary, Canada

Jack DONGARRA

Innovative Computing Laboratory, University of
Tennessee, US

John FEO

Pacific Northwest National Laboratory, US

Wojciech HELLWING

Center for Theoretical Physics PAS, Poland

John GUSTAFSON

National University of Singapore

David KAHANER

Asian Technology Information Program, Japan

Daniel S. KATZ

University of Illinois Urbana-Champaign, US

Scott KLASKY

Oak Ridge National Laboratory, US

Michael KRAJECKI

University of Reims Champagne-Ardenne, France

Dieter KRANZMULLER

Ludwig-Maximilians-Universität München, Germany

Henryk KRAWCZYK

Gdańsk University of Technology, Poland

Julian Martin KUNKEL

University of Reading, UK

Krzysztof KUROWSKI

Poznań Supercomputing And Networking Center,
Poland

Jysoo LEE

King Abdullah University of Science and Technology,
Saudi Arabia

Maciej MARCHWIANY

Interdisciplinary Centre for Mathematical and
Computational Modelling, University of Warsaw, Poland

Karlheinz MEIER

Heidelberg University, Germany

Richard MURPHY

Micron Technology Inc., US

Gabriel NOAJE

NVIDIA, US

Jacek OKO

Wrocław University of Science and Technology, Poland

Witold RUDNICKI

Interdisciplinary Centre for Mathematical and
Computational Modelling, University of Warsaw, Poland

Sven-Bodo SCHOLZ

Heriot-Watt University, UK

John SHALF

Lawrence Livermore National Laboratory, US

Horst SIMON

Lawrence Berkeley National Laboratory, US

Happy SITHOLE

Council for Scientific and Industrial Research, South
Africa

Sean SMITH

Australian National University

Thomas STERLING

Indiana University Bloomington, US

Rick STEVENS

Argonne National Laboratory, US

Bronis DE SUPINSKI

Lawrence Livermore National Laboratory, US

Vladimir VOEVODIN

Research Computing Center, Moscow State University,
Russia

Roman WYRZYKOWSKI

Czestochowa University of Technology, Poland

ORGANISING COMMITTEE



GENERAL CHAIR

Marek Michalewicz

ICM University of Warsaw

MEMBERS

Agnieszka WIECHECKA

Wojciech DYŻEWSKI



CONFERENCE DAY 1

■ Keynote Speaker

■ Invited Speaker

■ Sponsors

8:00 — 9:00	Registration		
9:00 — 9:05	Opening Remarks from the Organiser Marek Michalewicz, ICM UW, Poland	OPENING	Chair: Marek Michalewicz
9:05 — 9:15	Welcome Address Rector UW, His Magnificence Prof. Marcin Pałys		
9:15 — 9:25	Welcome Address Ministry of Science and Higher Education		
9:25 — 10:35	Advanced Computing at NASA Rupak Biswas, NASA Ames Research Center, USA	SUPERCOMPUTING EDUCATION	
10:35 — 10:45	Supercomputing invaders: Advances of Warsaw Team in Student Cluster Competitions around the world Warsaw Team, University of Warsaw, Poland		
10:45 — 10:55	Building a quantum computing awareness Paweł Góra, University of Warsaw, Poland		
10:55 — 11:15	BREAK		
11:15 — 12:00	Supercomputing and AI: Impact and Opportunities Anne C. Elster, Norwegian University of Science and Technology, Norway	SUPERCOMPUTING AND AI	Chair: Jacek Majewski
12:00 — 12:30	AI at scale with Cray Urika-XC – Recent achievements and future perspectives Alessandro Rigazzi, Cray Inc.		
12:30 — 13:00	New Technologies Needed for HPC, Exascale, AI Bill Mannel, Vice President & General Manager, HPC & AI Solutions Segment, Hewlett Packard Enterprise		
13:00 — 13:50	LUNCH		
13:50 — 14:35	Uncertainty at the Exa-Scale: the development of validated multi-scale workflows that quantify uncertainty to produce actionable results David Coster, Max Planck Institute for Plasma Physics, Germany	APPLICATIONS	
14:35 — 14:50	Time dependent density functional theory and supercomputing: new prospects for modelling superfluidity in neutron stars Gabriel Wlazłowski, Warsaw University of Technology, Poland		
14:50 — 15:05	Supercomputing with FPGAs Grzegorz Korcyl and Piotr Korcyl, Jagiellonian University, Poland		
15:05 — 15:20	Application of the Approximate Bayesian Computation to create the data-consistent model of the Forbush decrease of galactic cosmic ray intensity Anna Wawrzyńczak-Szaban <i>et al.</i> , Siedlce University of Natural Sciences and Humanities, Poland	APPLICATIONS	
15:20 — 15:35	Impact of WC/Co/diamond sample with peridynamics (view abstract) Eligiusz Postek <i>et al.</i> , Institute of Fundamental Technological Research Polish Academy of Sciences, Poland		
15:35 — 15:55	BREAK		
15:55 — 16:20	Sol: Transparent Neural Network Acceleration Nicolas Weber, NEC Research Laboratories	APPLICATIONS	
16:20 — 17:05	Chemistry with supercomputers Wojciech Grochala, University of Warsaw, Poland		
17:05 — 17:20	Atomistic spin model simulations of magnetic properties of ferromagnetic (Ga,Mn)N layers Dariusz Szeńkiel <i>et al.</i> , Institute of Physics Polish Academy of Sciences, Poland		
17:20 — 17:35	Spin-spin magnetic interactions from high-level Density Functional Theory calculations Dominik Kurzydłowski and Wojciech Grochala	APPLICATIONS	Chair:
17:35 — 17:50	Simulating Incompressible Fluid Mechanical Equations with Fourier Spectral Methods Brandon Cloutier <i>et al.</i> , University of Michigan, USA		
17:50 — 18:05	Towards convection-resolving Numerical Weather Prediction on GPUs Zbigniew Piotrowski <i>et al.</i> , Institute of Meteorology and Water Management, Poland		

CONFERENCE DAY 2

■ Keynote Speaker

■ Invited Speaker

■ Sponsors

8:00 — 9:00	Registration		
9:00 — 10:10	Memristor – Remembrance of things past Leon Chua, University of California Berkeley, USA	ARCHITECTURES	Chair: Krzysztof Kurek
10:10 — 10:55	"In-Memory Computing": Accelerating AI Applications Evangelos Eleftheriou, IBM Research Labs Zurich, Switzerland		
10:55 — 11:15	BREAK		
11:15 — 12:00	Optical Interconnect Technology Alex Wright-Gladstein, Ayar Labs, USA	INTERCONNECTS	Chair:
12:00 — 12:30	In-Network Computing: Introduction to Programmable Networks Sebastian Kalcher, Mellanox		
12:30 — 13:15	Data forever – myth or reality? Understanding the fundamental economics of data that enables you to store and access every piece of data forever Mark Lucas, Technologist Enterprise Engineering Western Digital		
13:15 — 14:15	LUNCH		
14:15 — 15:00	Studying brain structure and function through detailed computational modelling and topological analysis Michael Reiman, École Polytechnique Fédérale de Lausanne, Switzerland	COMPUTATION BIOLOGY	Chair: Joanna Sukowska
15:00 — 15:45	Computational Challenges in understanding the structure and function of the microbiome Tomasz Kosciółek, University of California San Diego, USA		
15:45 — 16:25	Sequence Similarity Search for Large-scale Metagenomic Data using Liquid Immersion Cooling Supercomputer Yoshinori Kimura, Infinite Curation, Japan		
16:25 — 16:45	BREAK		
16:45 — 17:00	The effect of crowded environment on the dynamics of the hepatitis C virus protease NS3/4a Natalia Ostrowska, Michael Feig and Joanna Trylska	COMPUTATIONAL BIOLOGY	Chair: Piotr Bala
17:00 — 17:15	Transport of vitamin B12 through the outer membrane protein BtuB in E.coli Tomasz Pieńko and Joanna Trylska, University of Warsaw, Poland		
17:15 — 17:30	SimRNA: a coarse-grained method for RNA folding simulations and 3D structure prediction Michał Boniecki et al., International Institute of Molecular and Cell Biology in Warsaw, Poland		
17:30 — 17:45	Supercomputing and causality analysis for a better understanding of the dynamic and functional properties of complex (bio)molecular systems Bogdan Lesyng, Mossakowski Medical Research Centre Polish Academy of Sciences, Poland		

CONFERENCE DAY 3

■ Keynote Speaker

■ Invited Speaker

■ Sponsors

8:00 — 9:00	Registration				Chair: Jacek Kitowski
9:00 — 10:10	Exascale Supercomputing Paul Messina, Argonne National Laboratory, USA				
10:10 — 10:55	Why Topology is Necessary at Exascale – And Why it's Not Easy Hamish Carr, University of Leeds, UK				
10:55 — 11:15	BREAK				
11:15 — 12:00	Extreme Scale Graphs Torsten Hoefer, ETH Zürich, Switzerland				
12:00 — 12:30	Data-driven design process for deep learning HW Maksymilian Bubula and Grzegorz Opoka, Intel				Chair: Łukasz Szustak
12:30 — 12:45	Software Product Lines and Exascale Stencils Muniyappa Manjunathaiah, University of Hertfordshire			COMPUTATIONAL STENCILS	
12:45 — 13:00	Performance Limits Study of Stencil Codes on Modern GPGPUs Ilya Pershin and Vadim Levchenko				
13:00 — 14:00	LUNCH				
14:00 — 15:05	Infrastructural and virtualization aspects of HPC solutions. Neptune: Or how to optimize energy consumption and performance for HPC Solutions Rick Koopman, Lenovo			IND.	Chair: Jerzy Proflicz
15:05 — 15:20	Evaluation of the DASH implementation of the HPCG Benchmark Daniel Rubio Bonilla and Jose Gracia, High Performance Computing Center, Stuttgart			BENCHMARKS	
15:20 — 15:35	HPC processors benchmarking assessment for Global System Science applications Damian Kaliszan <i>et al.</i> , Poznan Supercomputing and Networking Center, Poland				
15:35 — 15:55	BREAK				
15:55 — 16:40	What's the next step of accelerated supercomputing? Taisuke Boku, University of Tsukuba, Japan				Chair:
16:40 — 16:55	Implementation of a RISC-V-Conform Fused Multiply-Add Floating-Point Unit Felix Kaiser <i>et al.</i> , Heidelberg University, Germany				
16:55 — 17:25	Get ready for the data deluge: where are we Jean-Thomas Acquaviva, DDN				
17:25 — 17:35	Closing words Marek Michalewicz, Interdisciplinary Centre for Mathematical and Computer Modeling (ICM), University of Warsaw, Poland			CLOSING	

WORKSHOPS

DAY 4

TUTORIAL 1: HIGH PERFORMANCE I/O AND IN SITU DATA PROCESSING USING THE ADIOS FRAMEWORK

Presenters:

Scott A. Klasky, group leader for Scientific Data in the Computer Science and Mathematics Division, Oak Ridge National Laboratory
Norbert Podhorszki, team leader in the Scientific Data Group in the Computer Science and Mathematics Division, Oak Ridge National Laboratory

Abstract:

As concurrency and complexity continue to increase on high-end machines, I/O performance is rapidly becoming a fundamental challenge to achieving exascale computing. Wider adoption of higher-level I/O abstractions will be critically important to address this challenge. Modern I/O libraries provide data models, portable APIs, storage abstractions, and self-describing data containers. They achieve high performance and scalability, allow data to be managed more effectively throughout the data life-cycle, and enable reproducible science.

Part I of this tutorial will provide an overview of parallel I/O systems and summarize the key techniques for obtaining high performance I/O on high-performance computing (HPC) resources at scale. Part II introduces the [ADIOS library](#), delving through their usage models and examples, showing how to achieve high performance scalable I/O. Part III explains data compression. Part IV covers techniques for creating in situ analytics and teaches how to generate visualization services. Over one half of this tutorial will be hands-on sessions, where we provide access to the software and go through live examples using a [Virtualbox VM](#). We will have users work with real examples including [the heat equation](#), a [brusselator example](#), and [the Gray Scott Model of Reaction Diffusion](#).

About Presenters:

Scott A. Klasky is a distinguished scientist and the group leader for Scientific Data in the Computer Science and Mathematics Division at the Oak Ridge National Laboratory. He holds an appointment at the University of Tennessee, and Georgia Institute of Technology. He obtained his Ph.D. in Physics from the University of Texas at Austin (1994), specializing in general relativity. Dr. Klasky is a world expert in scientific computing and scientific data management, co-authoring over 300 papers.

Norbert Podhorszki is a senior scientist and team leader in the Scientific Data Group in the Computer Science and Mathematics Division at the Oak Ridge National Laboratory. He is one of the key developers of ADIOS that won an R&D 100 award in 2013. His main research interest is in creating I/O and staging solutions for in situ processing of data on leadership class computing systems. He received his Ph.D. in Information Technology from the Eötvös Loránd University of Budapest. He has worked in the field of logic programming, performance monitoring and analysis of message-passing programs, application monitoring in Grid environments and application development in Desktop Grids, scientific workflow technologies in supercomputing projects, and high performance I/O. He is author and co-author of more than 100 scientific papers and book chapters. He has given many tutorials on ADIOS, including several times at the Supercomputing and the International Supercomputing conferences.

WORKSHOPS

DAY 4

TUTORIAL 2: DISTRIBUTED COMPUTING AND SPATIAL DATA ANALYSIS IN JULIA

Presenters:

Bogumił Kamiński, PhD, Associate Professor, SGH Warsaw School of Economics

Przemysław Szufel, PhD, Assistant Professor, SGH Warsaw School of Economics

Abstract:

The goal of this course is to provide hands-on experience for computational scientists to leverage the power of distributed computing and simulations in the Julia language. You are invited to bring a laptop and install the software that will be used during the workshop (see the installation instructions below).

During the workshop we will present practical case study requiring massive computations. Both thread- and process- based parallelization in Julia will be discussed. We will also show how to create various types of distributed computational clusters.

The Julia language use case presented during the workshop will be a multi-agent, large-scale, vehicle routing simulation model that can be used analyze commuter's behavior in large cities. The numerical simulations will be run on actual road transportation system data taken from the Open Street Map project.

The course assumes intermediate prior knowledge of scientific programming in at least one language (e.g. Python, Gnu R, Java or Julia). In the first part of the tutorial we will provide an introduction to Julia programming. No significant prior knowledge of distributed computing or experience with large-scale simulation models is required.

Important Notes:

1. Please download Julia from <https://julialang.org/downloads/> and follow the installation instructions presented at <https://julialang.org/downloads/platform.html>
2. Install Julia packages that will be used throughout the course. In the command line (console) type `julia` and paste the following Julia code:

```
using Pkg
Pkg.add("DataFrames")
Pkg.add("PyCall")
Pkg.add("Conda")
Pkg.add("PyPlot")
Pkg.add("Plots")
Pkg.add("CSV")
Pkg.add("LightGraphs")
Pkg.add("IJulia") # optional for participants wanting to use Jupyter Notebooks
```

3. A convenient programming environment for the Julia language is **Atom** plugin named Juno. In order to install Atom with Juno please follow the steps below:

1. Download and install Atom (available at <https://atom.io/>).
2. Start Atom and press Ctrl + , (Ctrl key + comma key) to open the *Atom settings* screen.
3. Select the *Install* tab.
4. In the *Search packages* field, type `uber-juno` and press *Enter*.
5. You will see the `uber-juno` package developed by JunoLab—click *Install* to install the package.

4. Optionally, it is possible to use Julia within Jupyter notebook.

In order to try Julia inside a Jupyter notebook start the Julia console and run the two following commands:

```
using IJulia
notebook(dir=pwd()) # a new web browser tab should open with Jupyter Notebook
# showing current Julia working directory
```

About Presenters:

Bogumił Kamiński is the Head of Decision Analysis and Support Unit at SGH Warsaw School of Economics and Adjunct Professor at Data Science Laboratory, Ryerson University, Toronto. He is a member of the Management Committee of European Social Simulation Association (ESSA), and Vice President of Institute for Operations Research and Management Sciences (INFORMS) Polish Chapter. His field of expertise is operations research, with special focus on industrial applications of forecasting, optimization and simulation. He has been involved in development of core Julia language and its packages related to data science workflow. He is one of the top answerers for julia-lang tag on StackOverflow.

Przemysław Szufel is an Assistant Professor in Decision Support and Analysis Unit at Warsaw School of Economics. He is a member of the Management Committee of European Social Simulation Association (ESSA). His current research focuses on methods for execution of large-scale simulations for numerical experiments and optimization. He is an author or a co-author of several Open Source tools for high performance and numerical simulation (such as KissCluster, D-MASON, Isislab SOF, SilverDecisions, PyCX), and actively participates in their development and a co-author of various algorithms for distributed simulation-optimization models (such as AKG, AOCBA).

WORKSHOPS DAY 4



TUTORIAL 3: ADVANCED SCIENTIFIC VISUALIZATION WITH VISNOW PLATFORM

Brief agenda:

1. Introduction to Scientific Visualization and Visual Analysis.
2. Visualization systems and paradigms
3. Generic data structures
4. Introduction to VisNow
5. Hands-on Session #1 – 2D data visualization
6. Hands-on Session #2 – 3D data visualization
7. Hands-on Session #3 – Vector data visualization
8. Hands-on Session #4 – Unstructured data visualization.

Presenters:

Bartosz Borucki, Interdisciplinary Centre for Mathematical and Computational Modelling University of Warsaw, Poland
Krzysztof Nowiński, Interdisciplinary Centre for Mathematical and Computational Modelling University of Warsaw, Poland

Abstract:

Visual analysis is one of the most powerful tools for data exploration and interpretation. It takes advantage of visualization techniques and allows scientists to work with their research data in interactive and intuitive way. In today's HPC environment and Big Data era, data analysis techniques, together with visualization, gain on importance. However, the amounts of data and the sizes of single datasets impose the need for adequate software tools. In this tutorial we will address this problem by providing participants with strong tool for data processing, visualization and visual analysis – VisNow, an open source generic platform based on data flow paradigm. The goal of this tutorial is to introduce the audience to the concept of visual analysis, show basic ideas of scientific visualization and to go step-by-step through several case studies in hands-on sessions based on our platform. Problems of visualization of common HPC data structures, including 2-D and 3-D, scalar and vector, regular and unstructured data will be covered and adequate elements of the software described to give participants the basics of VisNow usage.

TUTORIAL 4: HOW TO BRING SECURE AND SCALABLE HPC TO YOUR ENTERPRISE WITH LENOVO AND SUSE SOLUTIONS

Presenters:

Artur Duszczak, Senior Technical Sales Representative, Lenovo
Marcin Madey, General Manager, SUSE Polska

Abstract:

High-performance computing (HPC) is no longer the sole domain of well-funded research institutions. Enterprises across industries such as finance, manufacturing and healthcare are employing HPC today to gain a competitive edge. Capturing these HPC benefits for your enterprise is now easier than ever. By teaming up with Lenovo and SUSE, not only do you get a powerful and efficient HPC solution with strong support, you also partner with organizations renowned for their HPC leadership. Attend our session and learn more how together, Lenovo and SUSE offer

TUTORIAL 5: IBM Q – PROGRAMMING A QUANTUM COMPUTER

Presenters:

Tomasz Stopa, PhD, IBM-Q Ambassador, IBM Kraków Software Laboratory

Abstract:

The session will be divided into two parts.

First part (approximately 1h) will cover a short introduction to quantum computing and overview of IBM's quantum computer architecture. During this session both physical implementation as well as software environment of IBM-Q will be presented with examples.

The second part will be a hands-on workshop using real IBM quantum computer in a cloud. Both interactive composer as well as QISKit SDK examples will be used to implement simple quantum algorithms.

The course is designed as introduction into the topic of quantum computation and as such does not require any specific prior technology nor programming knowledge. However, basic knowledge of quantum physics and python programming will help to fully benefit from the session.

Important Notes:

In the second part of the workshop, participants should have Python 3.7 programming environment with qiskit installed on their computers. The easiest way to do that is to download and install Anaconda distribution: <https://www.anaconda.com/distribution/#download-section>, followed by qiskit installation using 'pip install qiskit'.

About presenter:

Tomasz Stopa works in IBM Software Laboratory in Kraków, Poland. He obtained his Ph.D. in Physics from AGH University of Science and Technology in Kraków. With background in theoretical solid state physics he moved to software development 12 years ago. Currently works as development manager in software asset management space. Tomasz is also IBM-Q Ambassador and IBM Master Inventor with several patents and publications.

ACKNOWLEDGEMENT

Honorary Patronage



MINISTRY
OF ENTREPRENEURSHIP
AND TECHNOLOGY



MINISTRY
OF INVESTMENT
AND ECONOMIC DEVELOPMENT



Ministry
of Digital Affairs



UNIVERSITY
OF WARSAW

Patronage



Ministry of Science
and Higher Education

Republic of Poland

Media Partners

COMPUTERWORLD
FROM IDG

Primeur
Magazine

ENTERPRISETECH
INSIDE ADVANCED SCALE CHALLENGES

HPC | wire

datanami
DATA SCIENCE • AI • ADVANCED ANALYTICS

CIO
FROM IDG

PLab.pl
DIGITAL COMMUNITY

SPONSORS

Platinum Sponsors



comtegra



CRAY

Gold Sponsors

Western Digital

Lenovo



Silver Sponsors




**Hewlett Packard
Enterprise**

Bronze Sponsors

inspur