Mario Acosta

Barcelona Supercomputing Center

mario.acosta@bsc.es

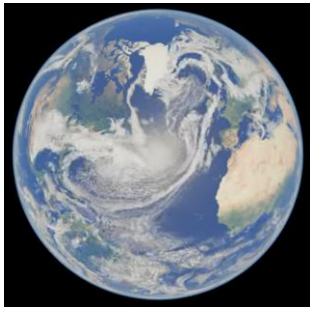

- **Head** of the Computational Earth Sciences group at BSC with more than 80 people.
- **Coordinator** of the Horizon Center of Excellence ESiWACE towards exascale in weather and climate simulations

Fun:

- Hiking in nature
- Football and other sports
- Surviving my little daughters

Background and research

- High performance computing applied to Earth Science applications
- Efficiently address the increasing complexity of Earth System Models with higher global resolutions and very long experiments for climate projections.
- My group focus on co-design, computational efficiency, model development, workflow manager integration, data efficiency and management, IA applications for climate, etc.

Big questions for workshop

- Have a representative weather and climate benchmark for postexascale initiatives
- Leverage accelerators such as GPUs and new technologies like the RISC-V processor or ML/IA units, including climate benchmarking in the co-design cycle.
- Development of new tools and architectures taking into account the scientific domain needs

Gabriel ANTONIU

gabriel.antoniu@inria.fr

Co-leader of the Digital Continuum Session

Gabriel Antoniu is a Research Director at Inria, Rennes, where he leads the KerData research team.
His recent research interests include scalable storage, I/O and in situ visualization, data processing architectures favoring the convergence of HPC, Big Data analytics and AI across the Edge-Cloud-HPC continuum.
He currently serves as an Executive Director of the JLESC international lab on HPC for Inria.
He has served as a PI for several international projects in these areas in partnerships including Microsoft Research, IBM, ATOS/Eviden, Argonne National Lab, the University of Illinois at Urbana Champaign, University of Tsukuba, Universidad Politécnica de Madrid, Barcelona Supercomputing Center.
He served as Program Chair for the IEEE Cluster conference in 2014 and 2017 and regularly serves as a PC member of major conferences in the area of HPC, Cloud computing and Big Data analytics (SC, HPDC, CCGRID, Cluster, Big Data, etc.).
He has acted as advisor for 20+ PhD theses and has co-authored over 150 international publications in the aforementioned areas.

Rosa M Badia - Barcelona Supercomputing Center

rosa.m.badia@bsc.es

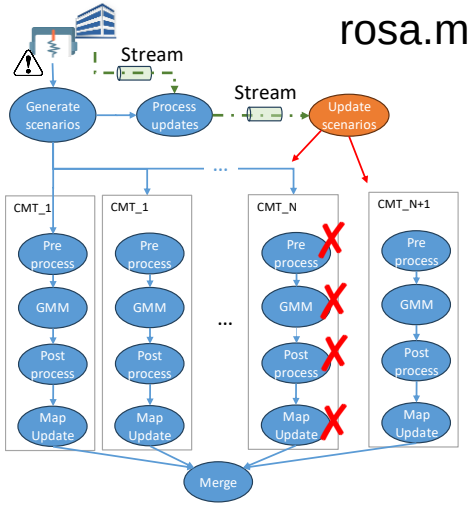

Rosa leads Workflows and Distributed Computing group, developing COMPSs programming model and runtime
HPC software res. area director (April 2025)

FUN:
- Swing dance
- Yoga

Background and Research
- Task-based programming models
- Convergence of HPC, AI and Big data
- Hybrid quantum-classic applications
- Swarm oriented platforms for edge-to-cloud
- Methodologies for developing digital twins

Big Questions for Workshop
- Can we improve programming and the management of the continuum?
- How LLMs can be used in the development of digital twins? Can we use agentic AI?
- Which will be the killer application for hybrid HPC-Quantum computing?

Pete Beckman

Northwestern University
pete.beckman@anl.gov

SAGE
Cyberinfrastructure for AI at the Edge
sagecontinuum.org

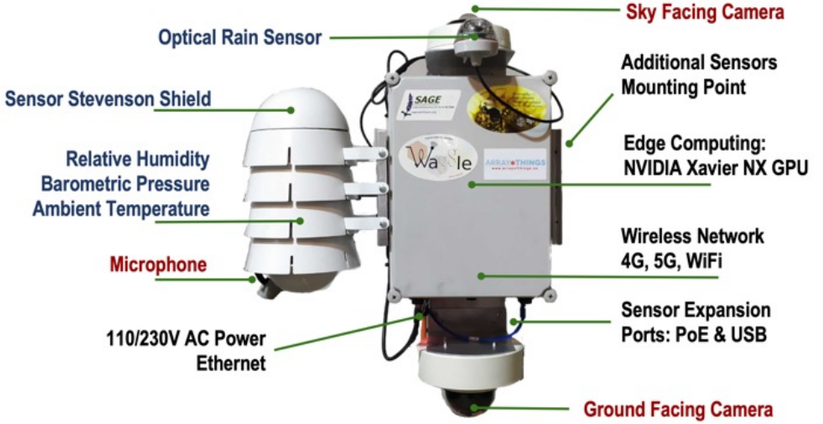

Pete leads the Sage project which has built a national AI cyberinfrastructure for the edge. Formerly, he spent 25 years in the DOE, and lead the HPC center, developed system software

Fun:

- Adventure kayaking, biking, canoeing, hiking
- Annual volunteer trips to west Africa
- Woodworking, electronics, photography, remote destination travel, book binding, ham radio, off-grid 4x4 vanlife

Background & Research

- Sage Grande is deploying 300 new AI-enabled nodes around the US
- Sage is exploring real-time hazards, drought, environmental resilience
- We are deploying LLMs to the edge for data analysis

Big Questions for Workshop

- How will Generative AI and HPC and Edge be linked into a computing continuum?
- What are the challenges around benchmarking, trustworthy AI, and LLMs for autonomous, distributed systems?
- How can HPC and Edge be linked internationally, and how can platforms be shared?

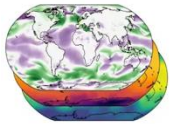



Jean-Yves Berthou, Inria

jean-yves.berthou@inria.fr

Jean-Yves is Deputy CEO of Inria, co-leader of the French Software Exascale program.

Fun: trail running, hiking, skying in the Alpes, sailing, cooking

Background and Research

Computer scientist in HPC and related stuff

Big Questions for Workshop

How to make applications efficient, portable and sustainable at the post-Exascale era?

What will be the new emerging usages?

What is Digital Continuum at the post-Exascale era?

What are the right AI methods/algo for a given problem so that resources consumption are minimized?

How to go towards validated/robust and trustworthy AI models/services?

Julien Bigot – CEA/MdlS

julien.bigot@cea.fr

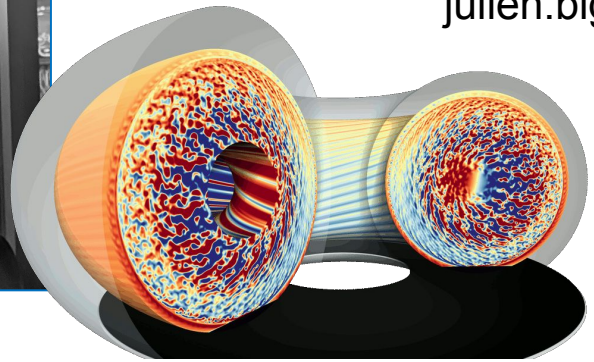

Julien is a CEA research scientist at *Maison de la Simulation* where he leads the *Science of Computing* team

Fun

I don't do any sport but I love skiing & outdoor climbing. I work on HPC but have fun embedding microcontrollers in model rockets. Those who know me well know my only strength is enjoying discussions with friends around beers

Background & research

- Separation of concern between application domain & optimizations
- Programming models for HPC
- In situ analytics workflows
- I co-lead NumPEx/ExaDoST (HPC data handling)
- I represent CEA in the HPSF
- I lead the CExA project (Kokkos at CEA)

Big Questions for Workshop

- How can we achieve a sustainable HPC software ecosystem
- What are the right interface between users & administrators for software deployment

INFN – Italy (National Institute for Nuclear Physics)

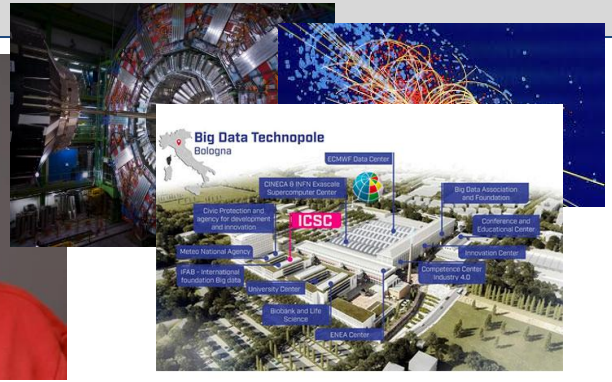

Director of Research at INFN; Supervisor for INFN of the DA.MA/Technopole; CMS/CERN former computing coordinator; CERN WLCG computing deputy leader

Fun:

- Everything having to do with technology (current pet project: porting sw to RISC-V)
- Hiking (compatibly with age!)
- Reading

Background and Research

- *High Energy Physics (HEP) researchers*
- *In (CERN and INFN) scientific computing since ~25 y*
- *Initiator of attempts to execute Experimental High Energy Physics workflows on HPCs*
- *2 PRACE project access + several smaller projects*
- *In several EU boards and committees (EOSC, EGI, Science Europe, EIC, e-IRG)*

Big Questions for the Workshop

- *How to integrate HPCs in a distributed data-intensive processing infrastructure?*
- *How to equalize (useful!) access to different HPCs?*

Taisuke Boku

Center for Computational Sciences, University of Tsukuba
taisuke@ccs.tsukuba.ac.jp

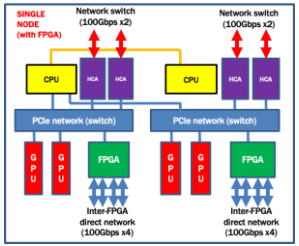

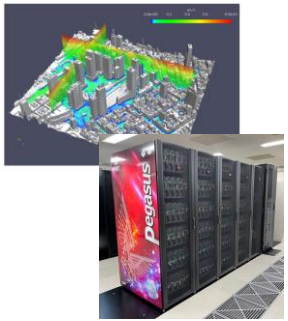

- Graduated from Keio University, Japan
- Working at Center for Computational Sciences, University of Tsukuba > 30 years, including six years as the director
- Member of most of MEXT HPCI related committees

FUN!!!

- Loving skiing and snow sports in mountains
- Speaking Korean as well as Japanese and English with frequent travels to South Korea
- Big fan of both Star Trek and Star Wars!!

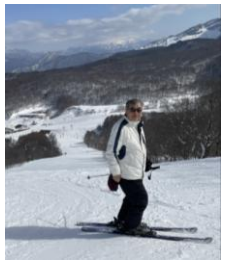

- **Background and Research**
  - Large scale parallel computing system architecture, system software and application codesigning on HPC applications
  - Heterogeneous computing with accelerators including many core, GPU and FPGA, on the system architecture and compilation
  - Leading most of U. Tsukuba's supercomputer design, deployment and operation
  - Joining Flagship system design WG of Japan on K and Fugaku
- **Big Questions for Workshop**
  - How to establish performance portable user-friendly programming environment for accelerators
  - What is the short-, mid- and long-term roadmap for AI for Science

Paola A. Buitrago

Pittsburgh Supercomputing Center
Carnegie Mellon University, University of Pittsburgh
pbuitrag@andrew.cmu; paola@psc.edu

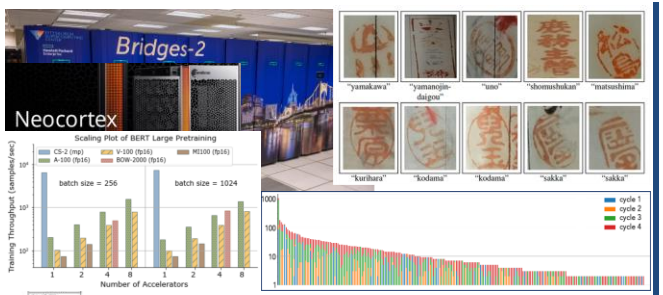

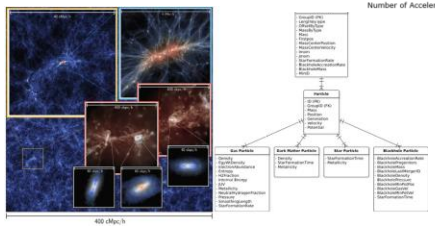

Founding Director of the AI and Big Data group at the PSC. Interested in bringing HPC and AI to wider audiences, domains, and applications.

Fun:
I love spending time in the outdoors. I enjoy hiking (mountains, waterfalls) and cycling. I have recently started exploring pickle ball with friends! I love teaching and volunteer teaching kids in middle and high school.

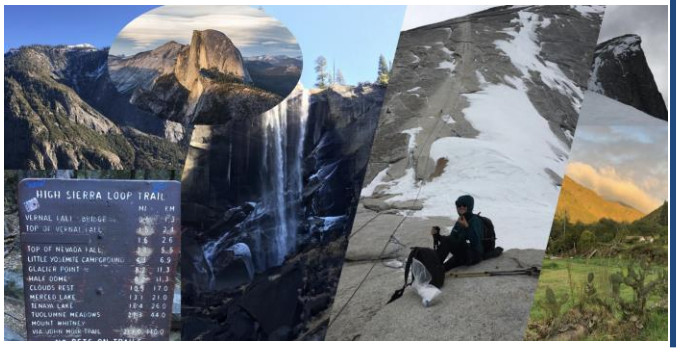

Background and Research

- (Generative) AI for Science. Foundation models for time series and for neuroscience.
- Research at the intersection of AI and HPC (AI for HPC and HPC for AI).
- Joined PSC after working in the Scientific Computing Division at Fermilab.
- Chemical Engineering, Systems and Computing Engineering, Machine Learning, and Electrical and Computing Engineering.
- Data intensive and AI powered scientific workflows.
- Innovative hardware technologies for AI training and inference.

Big Questions for Workshop

- Reproducibility in HPC+AI.
- How to establish stronger national and international collaborations that allow us to share best practices and advance transformative projects together?
- Workforce development for the needs of the HPC+AI community.

Paola A. Buitrago

Pittsburgh Supercomputing Center
Carnegie Mellon University, University of Pittsburgh
pbuitrag@andrew.cmu; paola@psc.edu

Alfredo Buttari

CNRS-IRIT

alfredo.buttari@irit.fr

$$Ax=b$$

Alfredo is a CNRS research director at the IRIT laboratory of Toulouse (France) in the APO (parallel algorithms and optimization) team

Fun

- Sports (running, cycling)
- Baking bread and pizza
- Gardening

Background and Research

My research work lies mostly in the domain of parallel numerical linear algebra:

- Mathematical libraries: dense and sparse
- Approximate computing: mixed-precision, randomization, low-rank approximations
- Parallel programming models: task-based parallelism, PGAS

Big Questions for Workshop

- Will programming become a thing of the past?
- What are the upcoming disruptive technologies in HPC?

Franck Cappello, Argonne National Laboratory

cappello@anl.gov

Proposed Methodology				
Techniques	MCQ Benchmarks	Open Response Benchmarks	Lab Style Experiments	Field Style Experiments
Main Goal	Testing knowledge **breadth, basic reasoning**	Testing knowledge **depth, planning, reasoning**	**Realistic** testing	**Realistic trend** analysis and weakness diagnosis
Problem Type	**Predetermined**, Fixed Q&As with known solutions	**Predetermined**, Fixed Free-Response Problems with known solutions	**Individual Human** Defined Problems with **unknown** solutions	**Many Human** Defined Problems with **(un)known** solutions
Verification	**Automatic** response verification	**Automatic or Human** response verification	**Humans detailed** response analysis	Scalable **automatic** summary of **human response**
Examples	Astro, Climate, AI4S (multi-domain), **Existing Benchmarks**	SciCode, ALDbench	**see "lab style experiments"**	**see "field style experiments"**
Cross Cutting Aspects	← Trust and Safety (ChemRisk), Uncertainty Quantification, Scalable Software Infrastructure (STAR) →			

Franck leads the scientific data management and reduction group at Argonne.
Before Argonne, Franck spent 10 years at CNRS and 10 years at Inria leading teams in high performance and distributed computing.

Sailing

Late art (novice)

Background and Research

- Scientific data reduction and exploring the notion of scientific information
- Scalable low overhead fault-tolerance/resilience for HPC and AI
- Evaluation methodology for LLMs as research assistants

Big questions for the workshop

- What scientific information should we inject into AI reasoning models?
- How to develop rigorous, scalable, and actionable evaluation of AI scientific capabilities?
- How generative AI (next generation of LLMs) will change research practices (what should we start training for/teaching)?

Daniele Cesarini – CINECA

d.cesarini@cineca.it

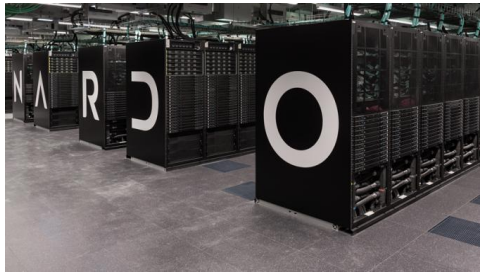

Daniele is a Project Manager & HPC Technology Specialist at the HPC department of CINECA where his works is focused on the definition of next-generation HPC architectures to define the strategic roadmap of CINECA's computing infrastructures.

Fun:

- Tennis, Crossfit, traveling
- Food, wine, gin tonic, margarita and many more

Background and Research

- His range of expertise includes the development of SW-HW co-design strategies to support large-scale workloads on AI/HPC supercomputers.
- He leads the energy efficient HPC activities of CINECA where his work is focused to improve the power efficiency of the CINECA's data centres.

Big questions for workshop

- Can AI models be redesigned to minimize data movement during inference?
- What is the trade-off between model accuracy and energy consumption on HPC platforms?

Vipin Chaudhary

Case Western Reserve University, vipin@case.edu

CASE WESTERN RESERVE UNIVERSITY

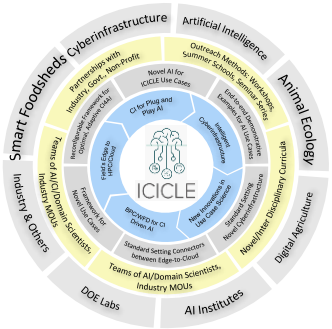

I am Kevin J. Kranzusch Professor and Chair of Computer and Data Sciences. Have been at NSF and quite a few startups.

Fun:

- Played badminton competitively
- Sports fanatic (any sport) and will do anything to keep fit now with injuries
- Read lot of STEM stuff

Background and Research

- Improving trustworthiness & performance of AI models in continuum (including spatial)
- Anomaly detection in large systems
- Reasoning and formal methods for LLMs
- Applications of AI in medicine
- Optimizing quantum circuits

Big Questions for Workshop

- Can we create benchmarks for reasoning about trust in LLMs for large systems?
- What are the challenges for plug-and-play AI in the computing continuum?
- What are challenges for MCP (model context protocol) to be used for HPC?
- What are the most viable applications that will benefit from hybrid quantum systems?

Susumu Date

D3 Center at University of Osaka

Bio

Susumu Date: a Professor of D3 Center at University ofOsaka. He is in charge of administration, management and system design of supercomputing systems at University of Osaka as well as R&D.

Funs

- Skiing
- Walking
- Driving
- Beer

Background and Research

- work for a supercomputing center in Japan and so have a responsibility to provide high performance computing resources for researchers in Japan.
- interested in high performance computing infra. and data infrastructure for (university) research
- also interested in missing (elementary) technologies to prepare for next-generation science research

Big Questions for Workshop

- What is supercomputing infrastructure like for convergence of HPC & AI & IoT?
- What is SC Infra. Like for OpenScience?

Jack Dongarra
U of Tennessee, Oak Ridge Nat Lab, & U of Manchester

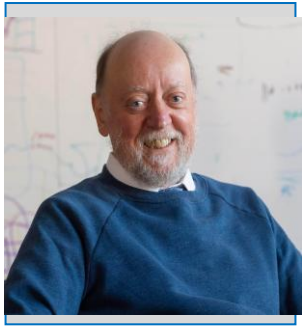

THE UNIVERSITY OF TENNESSEE KNOXVILLE

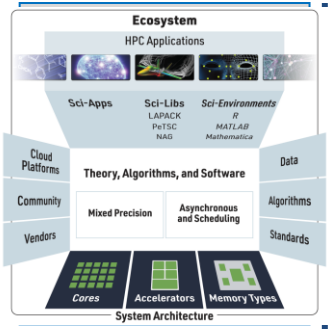

Jack is involved in the design and implementation of numerical algorithms for high performance computing.

FUN:
- Enjoys photography
- Spending time with eight grandchildren

Background and Research
- Specializes in numerical algorithms in linear algebra, parallel computing, the use of advanced computer architectures, programming methodology, and tools for parallel computers.

Big questions for workshop
- Co-design Benchmarks of hybrid workflows
- Software Integration
- Methodologies and tools for reproducibility in HPC

Anshu Dubey

Argonne National Laboratory
adubey@anl.gov

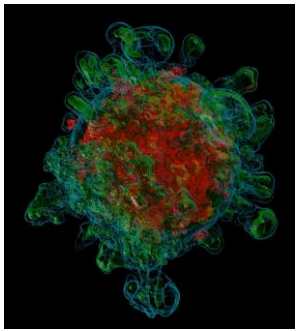

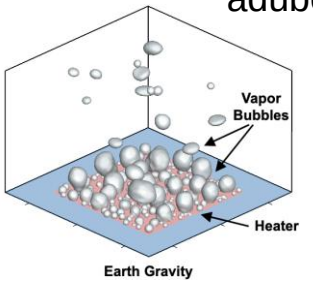

- Lead development of Flash-X, a Multiphysics, multidomain software system for solving PDEs
- Interested in productivity and quality improvement in scientific software development

FUN
Tinkering – cooking / electronics / fabrics
Reading and listening to audio books
Travelling to historic places – I read about them and then I want to see them
Engaging in interesting discussions on various topics

- Using GenAI for code translation from Fortran to C++ for a production science codes that uses Monte-carlo method.
- Also using GenAI for generating code for a new parallel algorithm I am developing for using N-body massive PIC in Flash-X – experimenting with test-driven development
- Big Questions for Workshop : I am effectively writing code in English – how should teaching programming change in view of that ?

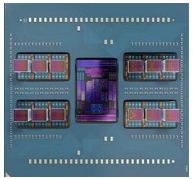
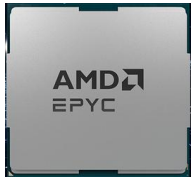

Yasuko Eckert (AMD)

Yasuko.Eckert@amd.com

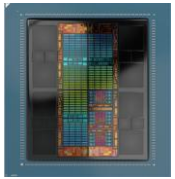

AMD

Yasuko is a Fellow at AMD Research and Advanced Development. She is based in Seattle, Washington.

FUN!!!

- Swimming
- Ocean
- Hockey

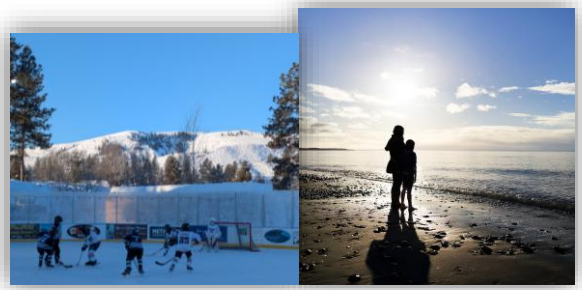

Background and Research

- PI of FugakuNEXT Feasibility Study
- Architecture pathfinding
- Exploring the future of supercomputing
- Contributed to DOE *Forward

Big Questions for Workshop

- How should future systems be designed to address sometimes conflicting needs of HPC and AI?
- How should we measure the success of post-Exascale systems?
- How can we foster/evolve public-private partnership to accelerate AI for Science?

Toshio Endo

Institute of Science Tokyo
endo@scrc.iir.isct.ac.jp

Toshio Endo is a professor at Supercomputing Research Center (SCRC), IIR, Science Tokyo.

FUN!!!

- I play with a daughter (8yo), a son (5yo)
- Video games including music games
- Karaoke?
- I played the E-bass and E-guitar long ago

Background and Research

- I'm working as the technical leader of TSUBAME4.0 supercomputer.
- High performance software on architecture with GPU accelerators and/or deep memory hierarchy
- Recently, job scheduling with very low latency for integrative jobs and digital twin

Big Questions for Workshop

- Parallel programming tools/framework, with which end-programmers are more unconscious with architectural detail
- Resource management for interactive / (quasi) real-time workload

Nicola Ferrier, Northwestern University & Argonne Nat. Lab

nicola-ferrier@northwestern.edu

NORTHWESTERN UNIVERSITY

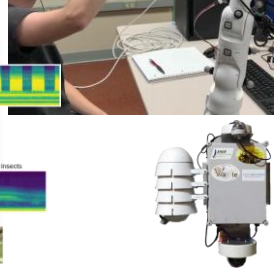

Argonne NATIONAL LABORATORY

Nicola is a research fellow at the Northwestern Argonne Institute for Science and Engineering & a Senior Computer Scientist at Argonne

FUN: cycling, swimming, hiking, cooking (Indian, Thai, Malaysian, Vegan), reading, gardening, bunnies

Background & Research

- Computer vision, robot eye-hand systems (was doing AI before it was cool!)
- Sage project - leads AI team
- AI across the Computing Continuum (edge to HPC – applies to digital twins, automated labs, and cool science)

Big Questions for Workshop

- With LLMs, should we still devote effort to “old school AI” – what is the balance?
- Can AI be trained to verify/validate output of AI or ensure we can trust results?
- How do we address “I like programming and problem solving – what will I do after AI takes over?”

HOKKAIDO UNIVERSITY

Takeshi Fukaya

Hokkaido University

fukaya@iic.hokudai.ac.jp

Hokkaido University Information Initiative Center

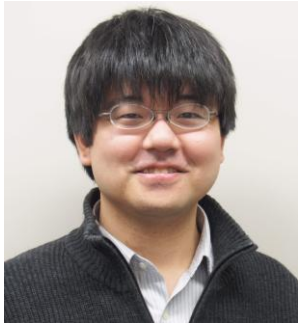

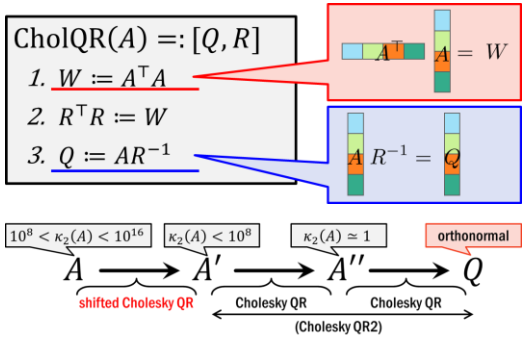

I am an Associate Professor at the Information Initiative Center, Hokkaido University, where the operation of the supercomputer system is one of my main missions.

FUN:

- Watching soccer games (J.League & European Leagues)
- Playing Japanese chess (将棋)

Background and Research

My research focuses on HPC in Numerical Linear Algebra (NLA):

- Comm.-avoiding/reducing algorithms
- Mixed-precision algorithms
- Algorithms for accelerators (e.g., GPU)

Big Questions for Workshop

- What kinds of technologies will be required in NLA in the future?
- How can traditional NLA methods be combined with recent AI technologies?
- What are the future directions for NLA research?

TEXAS ADVANCED COMPUTING CENTER

Niall Gaffney

ngaffney@tacc.utexas.edu

Artificial Intelligence Institute for Cosmic Origins
Funded by the U.S. National Science Foundation and the Simons Foundation

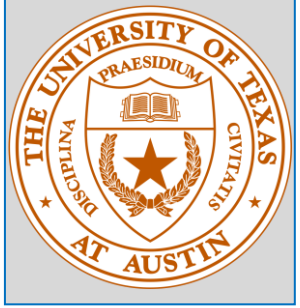

Director of Data Intensive Computing at TACC and Computing Director for the Cosmic AI Institute at UT Austin

FUN!!!
Travel, Cooking, Gardening, Formula 1 (go Oscar), and, when I am in places like Colorado, Road Cycling

Background and Research

- Observational Astronomer and Data Systems Architect for the Hobby Eberly Telescope and the Hubble Space Telescope archives
- Managing and utilizing complex datasets
- Sharing results, techniques, and workflows
- Copilots and surrogate modeling for research
- Services to simplify AI/ML adoption across a broader audience

Big Questions for Workshop

- What can HPC teach AI developers to lower costs and improve code portability in ML/AI?
- Can AI/ML help improve HPC codebases, or is that just a trap for future coders?
- How do we, on an HPC budget, keep up with AI/ML's “just throw money at it” approach?

Todd Gamblin

Lawrence Livermore National Laboratory
tgamblin@llnl.gov

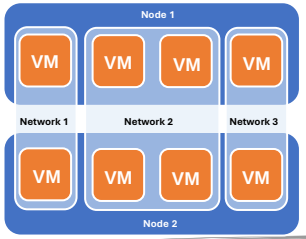

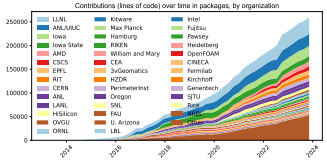

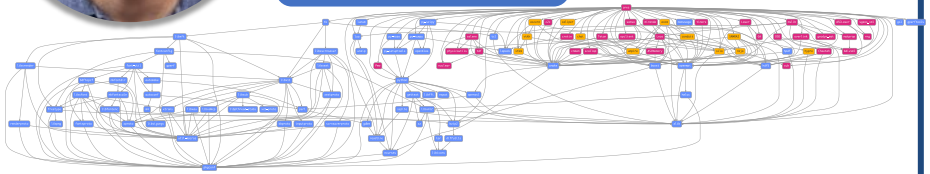

Todd is a Distinguished Member of Technical Staff at LLNL, and HPC strategist in Livermore Computing's Advaced Technology Office, and General Chair of the Governing Board of the High Performance Software Foundation (HPSF).

Fun:

- Family vacations w/kids
- Crossfit
- Go, Wordle, speed cubing (slowly)
- Open source, writing some code

Ik Kil Cenote

- Background and Research
  - Started out researching parallel performance tools
  - Started and lead the Spack Project, widely used HPC package manager / dev + build tool
  - Co-founded and currently lead HPSF: foundation for HPC open source
  - Interested in HPC/Cloud convergence, modernizing HPC data centers, and bringing cloud-like IaaS to HPC
- Big questions for the workshop
  - How can the performance-portable HPC software ecosystem going to go mainstream?
  - How are people going to deploy hybrid HPC/AI workflows in HPC centers?
  - Why can't HPC users have root and deploy their own services?

Lucas Gasparino

BSC-CNS (CASE)

lucas.gasparino@bsc.es

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

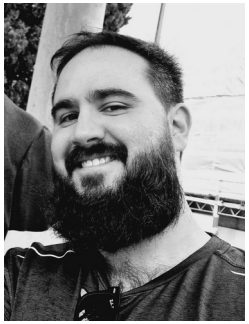

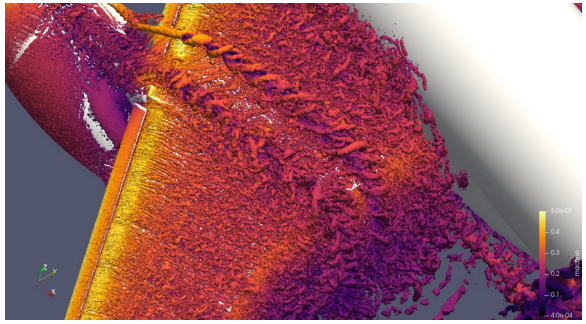

I'm a researcher in the LS-CFD group of Oriol Lehmkuhl at BSC. One of the main devs of the SOD2D code, a compressible solver for LES/DNS developed with current exascale architectures in mind.

FUN!!!

- Guitar and piano
- Programming
- Photography
- Cooking
- PARTY!!!

Background and Research

- Researcher at Oriol's LS-CFD group
- Dev/maintainer of SOD2D code
- Focus on efficient GPU-based algorithms

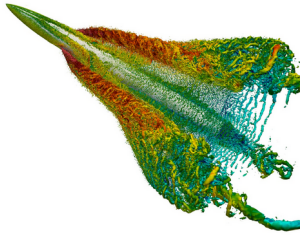

Big Questions for Workshop

- Role of AI in classical HPC workflows
- Challenges of developing vendor-agnostic scientific codes without sacrificing performance
- Role of RISC-V architectures in HPC

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

- Sergi Girona

Barcelona Supercomputing Center

sergi.girona@bsc.es

BSC 20 YEARS

Sergi Girona, PhD in Computer Engineering from the UPC, is the Operations Director of the Barcelona Supercomputing Center - National Supercomputing Center (BSC-CNS), since its creation in 2004.

Sailing in the Mediterranean, solo of with family/friend's crew.

Playing padel with friends, and the after padel

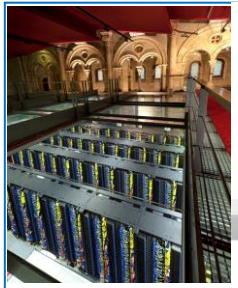

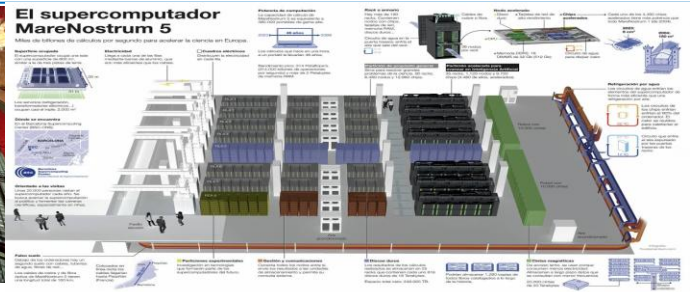

Background and Research

I start some year ago porting Unix V7, from PDP11 to M68010 dual processor, followed by a PhD on performance tools for parallel systems, and lately installing all MareNostrum, to achieve the best performance, in optimal infrastructure conditions, and maximizing the visibility of the systems.

Big Questions for Workshop

- What will be the impact to scientific international cooperation from the domestic political decisions?
- How to integrate AI innovation access for industry, not willing to queue, with normal supercomputing batch modes?

Maria Girone, CERN

maria.Girone@cern.ch

https://en.wikipedia.org/wiki/Maria_Girone
https://www.linkedin.com/in/maria-girone-4b375020/

Maria Girone is the Head of CERN openlab and leads the CERN IT Collaborations & Partnerships team. She is also a Principal Applied Scientist in CERN's IT department.

FUN!!!
cooking, hiking, skiing, pilates and dogs

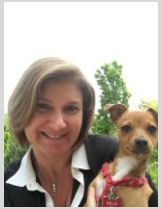

- Holds a PhD in particle physics and has extensive expertise in computing for HEP, having worked in scientific computing since 2002.
- Recognized leader in driving the adoption of HPC within the HEP community.
- Played a key role in the WLCG, founding and leading the operations coordination team responsible for core services and commissioning new tools.
- From 2014 to 2015, she was the software and computing coordinator for CMS, overseeing operations at about seventy computing centers worldwide.

Jason HAGA

The National Institute of Advanced Industrial Science and Technology (AIST)
jh.haga@aist.go.jp

Jason is a Chief Senior Research Scientist in the Continuum Computing Architecture Research Group of the Intelligent Platforms Research Institute.

Fun

- Exploring Japanese culture and Japanese sake
- Traveling, walking, exercising

Background and Research

- Immersive visualization and analytics to derive insight from data using interactive environments
- Data intensive applications for continuum computing
- Cyberinfrastructure for AI applications and development

Big Questions for Workshop

- How can visualization be used effectively with datacenter metrics to enable decision making by users?
- How can HPC centers innovate to address rising costs of energy consumption?

Toshihiro Hanawa, ITC, Univ. of Tokyo

hanawa@cc.u-tokyo.ac.jp

JCAHPC

Toshihiro Hanawa is a professor of the Supercomputing Research Division, Information Technology Center, U. Tokyo.

FUN!!!
- Trumpet
- Concert
- Cat

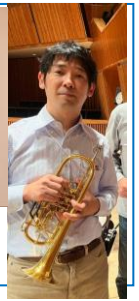

- Background and Research
  - GPU Computing
  - Interconnect for Accelerators
  - GPU Direct Storage
  - Energy Efficient Operation
- Big Questions for Workshop
  - How to overcome the barrier of GPU vendors?
  - Not only the programming model but also the total solution including communication and IO
  - What is energy-efficient operation that satisfies both users and administrators?

Yuan He

RIKEN Center for Computational Science
isaacyhe@acm.org / yuan.he.uw@riken.jp

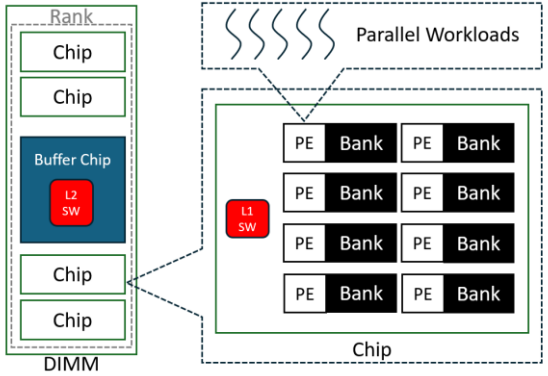

Yuan is a Research Scientist at the Next Generation High Performance Architecture (Kondo) Research Team of RIKEN Center for Computational Science

FUN!!!

- Jogging, playing table tennis and bowling
- Watching stand-up comedy and many sports, especially soccer, tennis and snooker
- Reading history and science fiction books

Background and Research

- Mainly focus on memory- and power-related topics to improve energy efficiency of computer systems
- Currently work on enabling the adoption of processing-in-memory and performance analyses of parallel workloads on emerging devices

Big Questions for Workshop

- Will future education turn into a training for writing better prompts to Generative AIs?
- When are we going to see Generative AI conducts scientific simulations from prompts?
- Are we carbon-based lives a precursor for silicon-based ones? (And if so, is this disobeying Darwin's theory of evolution?)

Toshiyuki Imamura

RIKEN Center for Computational Science
[Imamura.Toshiyuki@riken.jp]

理化学研究所
計算科学研究センター
RIKEN Center for Computational Science

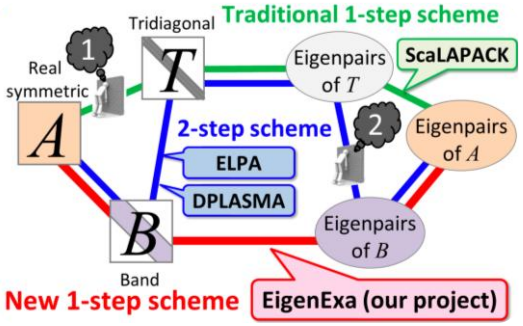

Toshiyuki leads the Large-Scale Parallel Numerical Computing Technology Team at RIKEN R-CCS, developing numerical libraries for K, Fugaku, and FugakuNEXT.

FUN!!!

- Sake: Japanese Rice-Wine, Beer in the world, as well as Japanese cuisine
- Ceering Japanese MLB players: Imanaga@CHC, and pitching trio Ohtani, Sasaki, Yamamoto@LAD

Background and Research

- Large-Scale Parallel Eigensolver: EigenExa and EigenG-Batched
- HPL-MxP: 2.0EFLOPS on Fugaku
- Mixed-precision for numerical linear algebra, algorithm, and software.
- Ozaki-scheme: significantly impacted the processor-design and AI.

Big Questions for Workshop

- How can we address the diversity of upcoming HWs, leveraging a combination of CPUs, GPUs, NPUs, TPUs, and QPUs?
- How will Generative AI enhance the future developments in Numerical Linear Algebra software?
- How about Quantum computing devices?

Takeshi Iwashita (Kyoto University, Japan)

[iwashita@i.Kyoto-u.ac.jp]

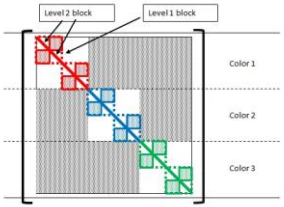

He received a B.E., an M.E., and a Ph.D. from Kyoto University in 1992, 1995, and 1998, respectively. In 2014-2023, he worked as a professor in Hokkaido University He currently works as a professor in Kyoto University.

I loved playing and watching soccer. I am a big fan of AS ROMA. I also cheer Hanshin Tigers, a Japanese professional baseball team. Also, I like visiting shrines and temples in Japan.

I conduct research in the area of high performance computing and numerical linear algebra. One of my research results is a parallel processing method for IC/ILU preconditioning and GS smoother (SISC'05, IPDPS'12). Recently, my colleagues and I investigated integer arithmetic based linear iterative solvers. We developed an algebraic multigrid preconditioned GMRES solver using iterative refinement framework in which integer arithmetic instructions are only used in an inner solver. The research result was accepted in ACM TOMS (open-access article).

[1] K. Suzuki, T. Fukaya, T. Iwashita, An Integer Arithmetic-Based AMG Preconditioned FGMRES Solver, ACM Transactions on Mathematical Software, Volume 51, 2025.

Shantenu Jha

Princeton Plasma Lab and Princeton University; Rutgers-New Brunswick
shantenujha@acm.org

Shantenu is a professor of computer engineering at Rutgers University-New Brunswick, the head of computational sciences at the DOE's Princeton Plasma Lab, and concurrently a research scholar in computer science at Princeton University. His research interests lie in AI for science, situated at the intersection of high-performance distributed computing and computational and data science. He is the lead of the RADICAL computing Lab.

Perfecting the art of being boring.
Pun for Fun.
Wander (wonder) the mountains of British Columbia.
Obsessive reading: from dismal to abysmal science.

Background and Research

- Computing for science
- Surely You're Joking, Mr Feynman? There might be plenty of room at the bottom, but there is a lot more room on the side!

Big Questions for Workshop

- How do you get 10 orders of magnitude (OOM) increase in scientific performance in less than 10 years?
- When you have 10+ OOM, what does it mean to be an intelligent algorithm?
- What happens when algorithms disobey their authors? (RADICAL computing ideas will not be enough. Nothing less than pure intellectual anarchy will do!)

Kimmo Koski, CSC

Kimmo.Koski@csc.fi

Kimmo is the CEO at CSC – IT Center for Science, Finland (since pretty many years…). Previous positions at CSC in 90's, CERN and Nokia.

FUN:

- Wife and three daughters. Had to get a male dog few years ago for equality reasons.
- A lot of veteran sports (Soccer, karate, weight lifting, gym, …) make me the best customer of the CSC medical center

Background and Research

- CSC is the national "HPC and almost everything else" center in Finland, working intensively also in Europe and globally. We host national and European systems such as LUMI andfuture AI Factory
- I have made my dissertation in 1996, topic Metacomputing technology, ever since less research and more management… good thing at least for quality of science ☺

Big Questions for Workshop

- Global collaboration – how to boost it ?
- AI future and how to get most out of it (and what not to get)
- Developing the HPC, AI, Quantum etc. ecosystem

Michael KRAJECKI,

michael.krajecki@polytechnique.edu

Agence ministérielle pour l'intelligence artificielle de défense

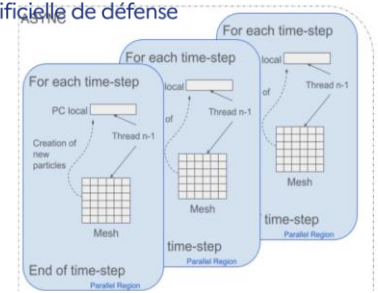

- Michael is the director of the research departement of the French AI for Defense Ministerial Agency
- He is professor at the university of Reims in France
- He is also co-director for CNRS of the national program NumPEx (French exascale program)

Background and Research

- Michael has established and led the ROMEO computing center, https://top500.org/site/50497/
- Research Focus : High-Performance Computing, Parallel Algorithms, Computing on Accelerators and Many Core Architectures, AI

Big Questions for Workshop

- How to integrate AI models into scientific simulations to improve the accuracy and efficiency of the results?
- How to design HPC infrastructures that effectively support both AI workloads and scientific computing?

Erwin Laure

Max Planck Computing and Data Facility (MPCDF)
erwin.laure@mpcdf.mpg.de

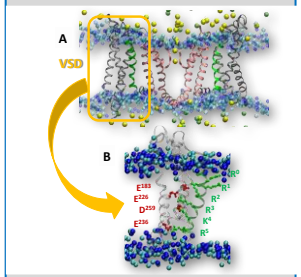

Erwin is the director of the Max Planck Computing and Data facility in Germany and Professor in HPC at TU Munich. He has over 20 years of experiences in HPC and is particularly engaged in EuroHPC where he was a member of the infrastructure advisory board and founder and member of several Centres of Excellence

Theater
Opera
Classical Concerts

- Building and operating HPC infrastructures for over 15 years
- Engaged in interdisciplinary HPC applications projects (Centres of Excellence)

- Big Questions for Workshop
- What role will AI really play in scientific workloads
- How to deal with reduced precision hardware

FIU FLORIDA INTERNATIONAL UNIVERSITY

Jason Liu

liux@fiu.edu

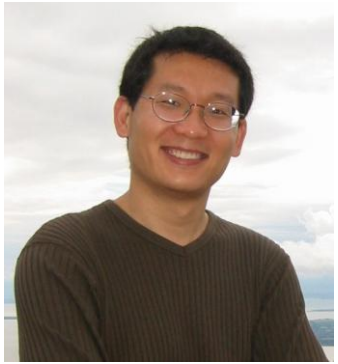

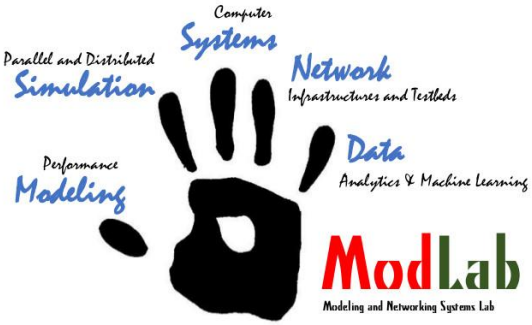

Jason Liu is an Endowed Computer Science Professor and Director at FIU's Knight Foundation School of Computing and Information Sciences (KFSCIS)

FUN!!!

- Tennis
- Hiking in nature
- Traveling and culture sightseeing

Background and Research

- Parallel discrete-event simulation
- Performance modeling and simulation of computer systems and networks
- HPC interconnection networks
- Memory, storage systems, I/O, caching
- ML applications
- Cyberinfrastructures (computing and storage)

Big Questions for Workshop

- Adaptive and reconfigurable HPC to support diverse science applications
- Integrate campus computing infrastructure running ML applications
- Efficient data workflow for large-scale ML science applications

Yohei MIKI (University of Tokyo)

[ymiki@cc.u-tokyo.ac.jp]

東京大学情報基盤センター
INFORMATION TECHNOLOGY CENTER, THE UNIVERSITY OF TOKYO

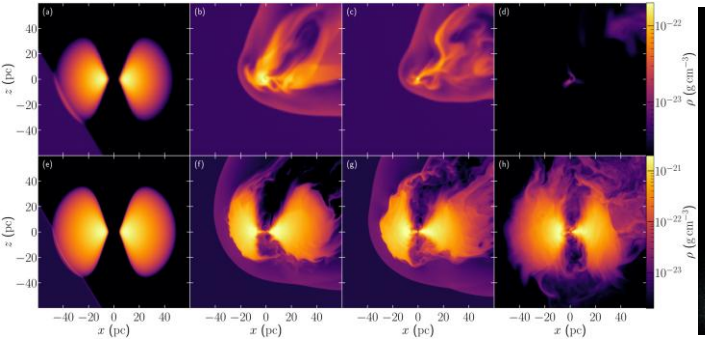

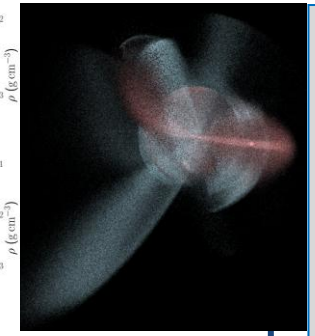

I am a theoretical/computational astrophysicist and a developer of software useful for astrophysical studies in Information Technology Center, University of Tokyo.

FUN!!!
Watching football games
Reading detective stories

Background and Research

- Formation and evolution of galaxies
- Galaxy merger
- Coevolution of massive black holes and galaxies
- Wandering massive black holes
- Missing satellite problem (aka subhalo problem)
- Galactic archaeology (aka near-field cosmology)
- *N*-body simulation
- GOTHIC: Gravitational Oct-Tree code accelerated by HIerarchical time step Controlling
- MAGI: MAny-component Galaxy Initializer
- Hierarchical Hamiltonian splitting
- Vendor-neutral GPU computing

Big Questions for Workshop

- What is the best (or suitable) algorithm and development environment for (galactic) astrophysical simulations?

Bernd Mohr, Jülich Supercomputing Centre

b.mohr@fz-juelich.de

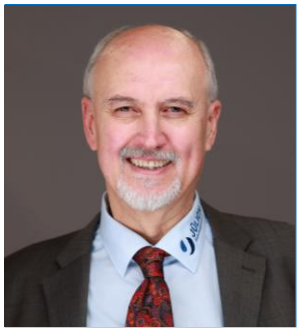

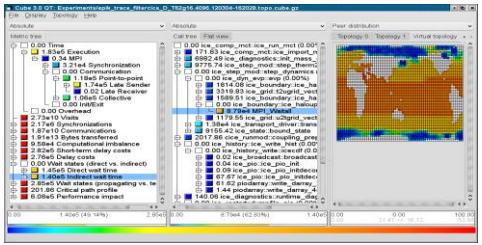

JÜLICH Forschungszentrum | JÜLICH SUPERCOMPUTING CENTRE

Division Head “Application Support”
Team Leader “Parallel Performance”
Super Computer Scientist

FUN!!!

Likes traveling and hiking
Cocktail named after him (“Mohrgarita”)
Playing “Real Racing 3”

Background and Research

- Tool design and development for performance instrumentation, measurement and analysis of HPC applications (SIMPLE, TAU, Vampir, KOJAK, Score-P, Scalasca, Cube)
- Parallel programming

Big Questions for Workshop

- How should we design and develop HPC HW + SW for post Exascale?
- Will the AI hype kill (or at least hinder/delay) successful HPC modeling and simulation work?

Thomas Moreau - Inria MIND

thomas.moreau@inria.fr

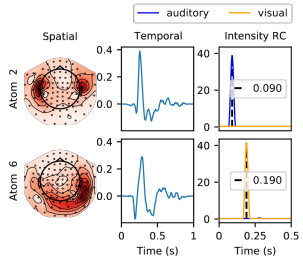

Researcher in Paris - AI for signals and benchmarks
Open-source developer and maintainer

FUN!!!

- Biking, Hiking, Jogging (up for a run?)
- Surfing, Skiing, Windsurfing, …
- Talking and remaking the world around beer or Sake

Background and Research

- **Core AI research:** representation learning, self-supervised learning, bilevel optimization, benchmarking methodology
- **Collabs:** Neurosciences, Anaesthesia, Physiological signals, Physics simulations
- **Open-source:** maintainer of benchopt, joblib, involved in scikit-learn ecosystem.

Big Questions for Workshop

- How to make benchmark so HPC and AI community can learn to collaborate ?
- How to foster concrete collaboration ?

Raymond Namyst
University of Bordeaux, Inria

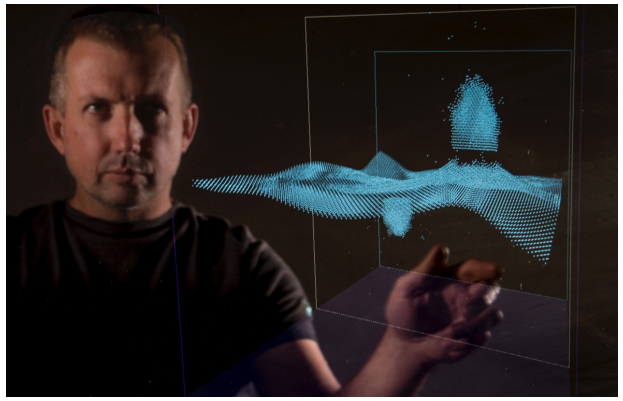

Director of HPC national French Program

Fun
- Road Cycling
- Photography of wild wild life
- Writing GPU shaders

Background and Research
- Runtime Systems for heterogeneous hardware
- Educational HPC environments

Big Questions for Workshop
- Hybridizing HPC with AI: real challenges and killer applications
- How should we train the next generation of HPC engineers?
- Is AI the most power inefficient way of computing?

Gautham Narayan

gsn@Illinois.edu
NSF-Simons SkAI Institute, University of Illinois, Urbana-Champaign

Funded by the U.S. National Science Foundation and the Simons Foundation

I'm a cosmologist studying the accelerating expansion of the Universe with supernovae & supercomputers

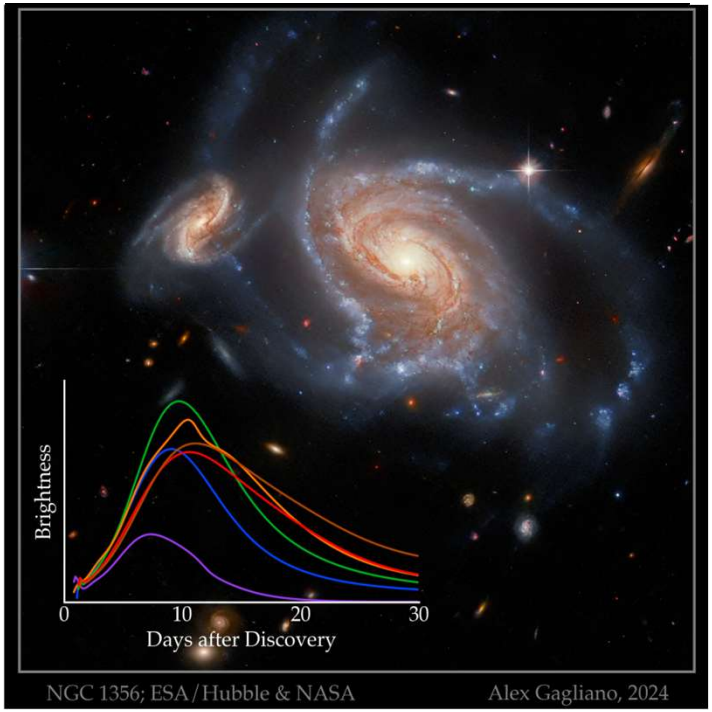

Fun: hiking, biking, kayaking, photography, travel, a good book or ten, telling your doggo that they are the best (cats OK too), unreasonable amounts of caffeine

Background and Research

- Deputy Director for NSF-Simons AI for the Sky Institute (SkAI)
- Analysis Coordinator and Spokesperson-elect for LSST Dark Energy Science Collaboration, Assoc. Prof @ UIUC Astronomy
- I use AI to discover/model/simulate stars that blow up to infer the matter-energy budget of the Universe (and other things)

Big Questions: How do we:

- Work with industry & govt. on Foundation AI models and policy to accelerate scientific discovery?
- Incorporate symmetries and statistics into massive interpretable AI models?
- Combine on-prem HPC, Cloud, and Edge to build a seamless compute environment that serves large international experiments?

Dhabaleswar K (DK) Panda

The Ohio State University

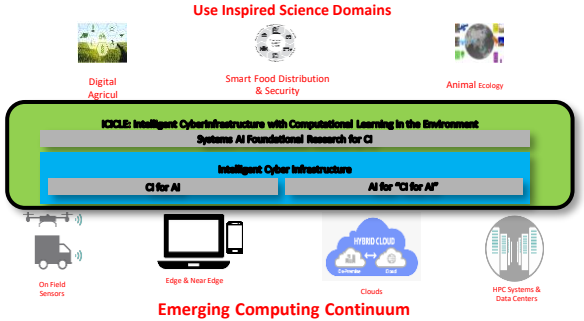

I am a Professor and University Distinguished Scholar at the Ohio State University. I also serve as the Director of the NSF-AI Institute, ICICLE (https://icicle.ai)

FUN!!!

- Cooking
- Reverse engineering food dishes
- Travel and exploring cultures, food, and history

Background and Research:

- High-Performance and Scalable Converged Software Stack
- MVAPICH MPI library (http://mvapich.cse.ohio-state.edu)
- High-Performance Deep/Machine Learning Libraries (http://hidl.cse.ohio-state.edu)
- Designing Intelligent Cyberinfrastructure and democratizing AI for domain scientists

Big Questions for Workshop:

- *Major innovations in the next decade: 1) HPC for AI, 2) AI for HPC, and 3) AI democratization*
- *Impact of the LLM revolution on designing and using next-generation HPC systems*

Manish Parashar

manish.parashar@utah.edu

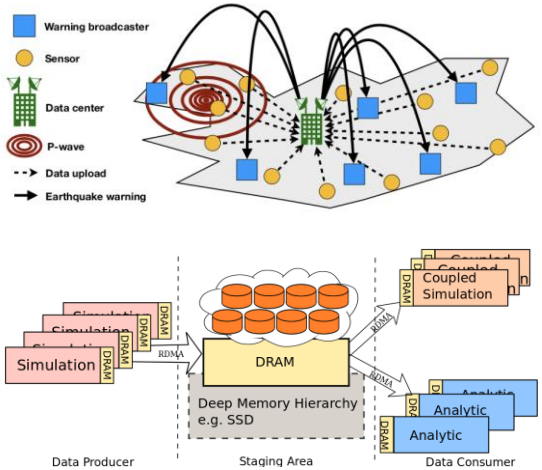

Manish Parashar is Director of the Scientific Computing and Imaging (SCI) Institute and Presidential Professor at the Kalhert School of Computing, University of Utah.

Running (run-walking), hiking, traveling, music (listening), traveling, cooking, …

Background and Research

- Urgent computing / Computing continuum
- Data management (National Data Platform, Science Data Exchange)
- Scalable computational & data-driven science (In-situ workflows, data staging)
- Translational computer science

Big Questions for Workshop

- How to democratize science across digital continuum?
- How to harness the digital continuum for urgent science (manage uncertainties, constraints)
- How to program data-driven, ai-enabled workflows across the digital continuum?
- How to scale the translational impact of computer science research?

Mark Parsons

EPCC, The University of Edinburgh
m.parsons@epcc.ed.ac.uk

Joined EPCC in 1994 after Particle Physics PhD at CERN. Director of EPCC since 2016 and Dean of Research Computing at the University of Edinburgh

Skiing, astrophotography (badly) and repairing/building 8-bit computers

Research interests

- Long term interest in system design and novel processor and interconnect technologies
- Over past decade coupled to large-scale data infrastructures for public sector research particularly with population data
- Most recently driving real Exascale applications – e.g. with Rolls Royce

Big Questions for Workshop

- Should we allow AI to dominate supercomputing?
- How do we take back the narrative? A new direction for our systems?

RIKEN
Center for
Computational Science

Theresa Pollinger, R-CCS

theresa.pollinger@riken.jp

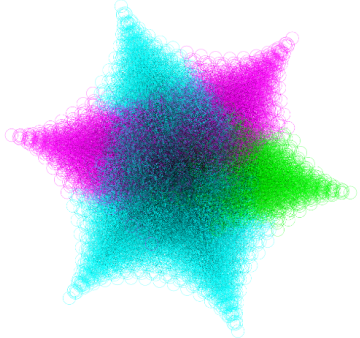

Postdoc at the intersection of numerics, HPC, AI, simulations, hardware & software

FUN!
Cycling, enjoying the beaches of Okinawa, reading, podcasts, learning about unpleasant facts and sharing them
(let me know in case you're interested!)

Background and Research:

- Engineering degree, boundless curiosity
- PhD focus: higher-dimensional simulations, numerics and HPC aspects
- Currently: intersection of HPC / AI (+more)

Big Questions for Workshop

- (How) Will AI help us outscale single human brains to generate new insight? What else can help?
- How do we build (work) environments that allow humans to be the most creative?
- How do we filter the really good stuff out of the AI hype? Is the current science workflow up to the challenge?
  If not, what is?

Andrea C. Riaño – Barcelona Supercomputing Center (BSC)

andrea.riano@bsc.es

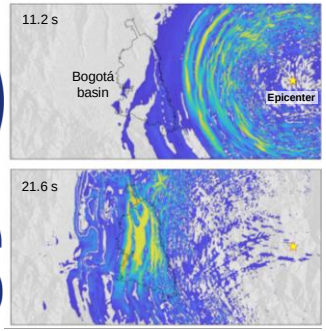

Researcher in the Wave Phenomena Group (BSC – CASE), specializing in earthquake engineering and computational seismology with a passion for large-scale earthquake simulations and seismic risk assessment.

FUN!!! I love binge-watching series, getting lost in paranormal activity podcasts, wandering aimlessly through cities, hiking, and soul-cycling.

Background and Research

- Established researcher focused on earthquake simulations.
- Applying CyberShake in the first use of physics-based simulations for seismic hazard assessment in Europe, as part of the EFERH community.
- Developing software to simplify regional integration into the CyberShake platform.

Big Questions for Workshop

- How can AI and HPC be combined to build real-time decision-support systems for disaster risk reduction?
- Are today's scientific tools and methods prepared to scale with emerging hardware, or do we need to rethink how we develop simulation code?
- Can AI help resolve uncertainty in complex simulations by learning from real data?

Kentaro Sano

Riken Center for Computational Sciences (R-CCS), kentaro.sano@riken.jp

PROCT

ESSPER
Elastic and Scalable System for High-Performance Re-configurable Computing

RIKEN R-CCS RIKEN TRIP

MOONSHOT RESEARCH & DEVELOPMENT PROGRAM

Kentaro is the leader of the processor research team, the advanced AI device development unit, and the next-generation HPC infrastructure system development unit.

FUN:

- Cycling with a road bike
- Snorkeling, Hot-spring
- Cooking, Brewing coffee
- Movies, Musicals, SLR

Background and Research

- Computer architectures for HPC and AI, especially reconfigurable one such as CGRA and specialized hardware with FPGAs. Implementing them in VerilogHDL, etc.
- Quantum error correction algorithm and hardware for fault-tolerant quantum computers. Developing FPGA cluster system for QEC.
- FugakuNEXT development
- Parallel visualization and vector-quantization in the past

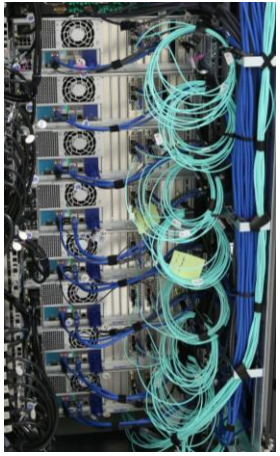

Big Questions for Workshop

- Any requirement for hardware research, in particular, FPGA and specialized architecture?
- Any interests in FTQC?

Kento Sato

RIKEN R-CCS / kento.sato@riken.jp

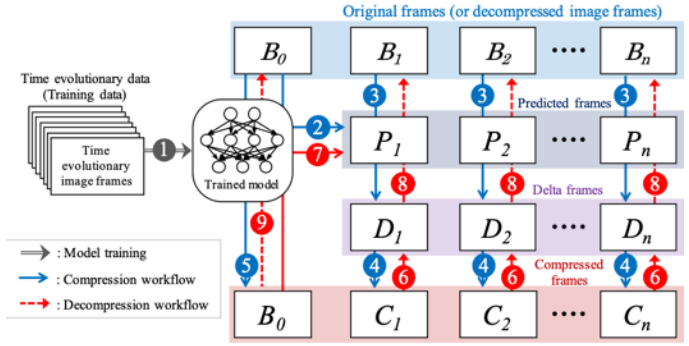

Figure: TEZip data compression

A team principal of High Performance Big Data Research Team in the Center for Computational Science at RIKEN

- **Fun:**
  - Futsal (Every Wednesday night)

- **Background & Research**
  - AI for HPC / HPC for AI
  - Application reproducibility (MPI/OpenMP)
  - Scalable fault tolerance (C/R)
  - I/O optimization / Data compression
  - Performance evaluation and analysis

- **Big Questions for Workshop**
  - What is the best solution to co-develop HPC software through international collaborations among universities, national labs and vendors ?
  - What is the best solution to build data infrastructure/platform to foster international collaborations ?

Martin Schulz

Technical University of Munich (TUM), CAPS Chair
schulzm@in.tum.de

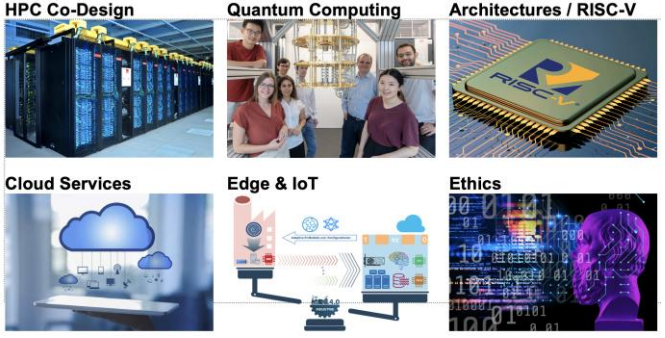

- Professor for Computer Architecture and Parallel Systems @ TUM
- Member of Board of Directors @ LRZ
- Before that 13 years @ LLNL

Fun:

- Beer brewing
- Home IoT
- Finishing the basement (slowly)

Background and Research

- Sustainable HPC
  - Energy efficiency in HW and SW
  - Dynamic resource management
  - Programming models, in part. MPI
- Integration of HPC and Quantum Computing
  - Munich Quantum Software Stack (MQSS)
- RISC-V processor extensions
  - AI, SmartNICs, Quantum control, Dataflow

Big Questions for Workshop

- How to efficiently program heterogenous systems with (very) different accelerators?
- How to combine ModSim+AI+QC?
- How to leverage more dynamic resources?

Takashi Shimokawabe

The University of Tokyo
shimokawabe@cc.u-tokyo.ac.jp

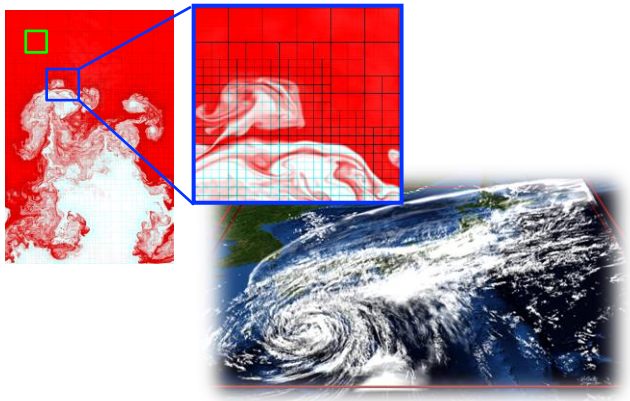

Takashi is an Associate Professor
in the Supercomputing Research Division,
Information Technology Center,
The University of Tokyo

FUN!!!

- Beer and Beer
- Travelling, mountain climbing, hiking, tennis

Background and Research

- Mesh-based applications on GPU-equipped supercomputers, such as fluid dynamics, weather simulations.
- Adaptive mesh refinement
- AI-based acceleration of simulations
- Performance portability programming

Big Questions for Workshop

- How can we accelerate conventional simulations (especially CFD) with the use of AI while keeping a high accuracy?
- How can we utilize LLMs for CFD?
- What languages and frameworks should we use for future application development?

Estela Suarez, SiPEARL, JSC, UniBonn

[estela.suarez@sipearl.com; e.suarez@fz-juelich.de]

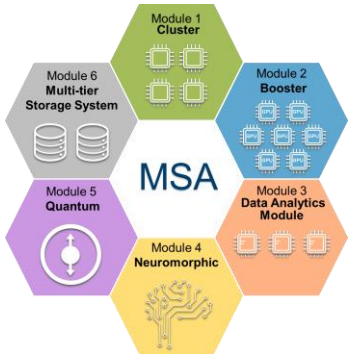

Bio

- *Senior Principal Solution Architect* at SiPEARL (sabbatical since Sept 2024)
- 15 Years at JSC: Joint Lead of division *Novel System Architecture Design*
- Associate Professor for HPC at UniBonn

Fun/Hobbies

- Running, cooking, painting, farm repairs

Background and Research

- HPC system architecture and codesign
- Modular Supercomputing Architecture
- Hardware prototyping, software environment, and applications
- Codesign: system and processor level
- System software for energy efficiency
- Processor development

Big Questions for Workshop

- Where is the largest potential to reduce the energy footprint in HPC? And does one *really* care?
- Which kind of processing units do we need in the future HPC systems? What comes after GPUs?

Toyo Suzumura

The University of Tokyo

suzumura@acm.org

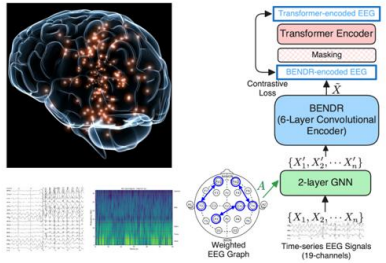

Graph-Enhanced EEG Foundation Model

47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society

Bio: Since 2021, I have served as a professor at the Information Technology Center of the University of Tokyo. Prior to that, I spent 17 years at IBM Research, working in Tokyo, Dublin, and at the T.J. Watson Research Center in New York

FUN: Swimming, Volleyball

Background and Research

- Graph Foundation Model
- Scalable User Behavior Modelling with Graph Neural Networks and Large Language Models
- EEG Foundation Model for Brain Science
- Large-Scale Data Science System (e.g. mdx)

Big Questions for Workshop

- How can we integrate large language model and brain foundation model for better human understanding ?
- What is the system architecture to combine data science platform and supercomputers ?

Ryousei Takano

National Institute of Advanced Industrial Science and Technology (AIST)
<takano-Ryousei@aist.go.jp>

Ryousei Takano is a principal research manager at AIST and currently leads a Quantum-AI hybrid computing infrastructure project (ABCI-Q).

FUN!!!
- Jogging
- Exploring craft beer spots and ramen shops

Background and Research
- System software (operating system, communication system, virtualization)
- High performance computing systems
- Quantum-HPC hybrid system

Big Questions for Workshop
- When and how will quantum computers be operated in HPC centers?
- What is the most important problem of Quantum-HPC hybridization?

Hiroyuki Takizawa

Cyberscience Center, Tohoku University
takizawa@tohoku.ac.jp

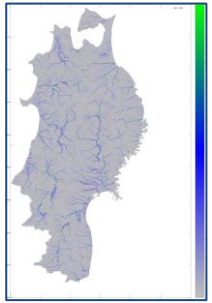

Hiroyuki (Hiro) is Professor and Deputy Director of Cyberscience Center, Tohoku University, leading the system operation of Supercomputer AOBA.

FUN!!!
Visiting historical places and aquarium.
Eating, drinking, singing? (not dancing)
…and I like being lazy and relaxing at home☺

- **Background and Research**
  - Programming models, languages, and AT for performance portability
  - Urgent and on-demand computing for disaster resilience (e.g. Tsunami)

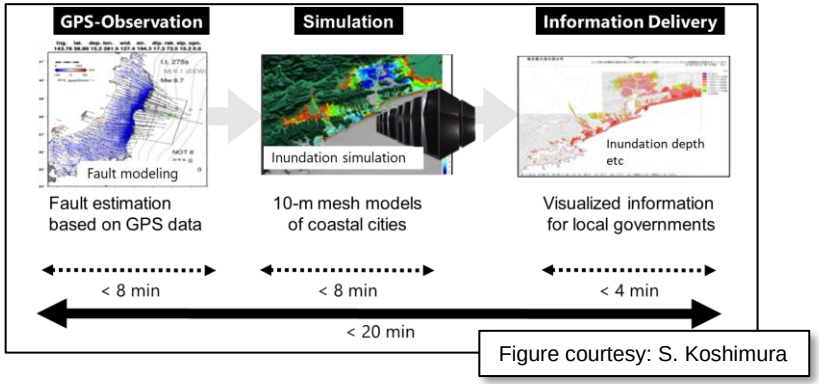

Figure courtesy: S. Koshimura

- **Big Questions for Workshop**
  - How can we make the most of HPC systems for diverse demands (HPC+AI)?

Yusuke Tanimura

National Institute of Advanced Industrial Science and Technology (AIST)
yusuke.tanimura@aist.go.jp

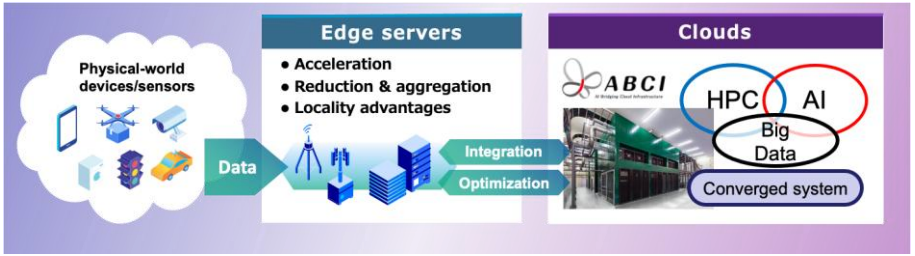

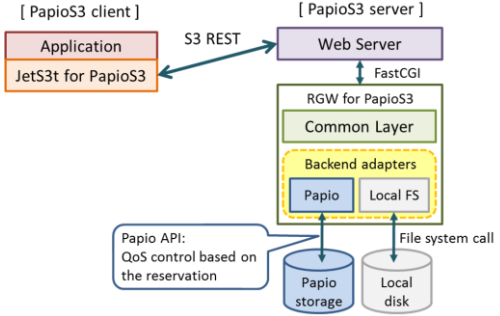

Yusuke is a chief senior researcher in the Intelligent Platform Research Institute, AIST.

FUN!!!

- Travel, drinking, food, ski, tennis, hiking, watching sports

Background and Research

- HPC storage and parallel data processing
- System management of AI-HPC converged insfrastructure like ABCI
- Data management in edge-to-cloud continuum

Big Questions for Workshop

- How the big HPC system should connect to data sources in the real world?
- How data should be managed in edge-to-cloud continuum?

筑波大学
University of Tsukuba

Osamu Tatebe

University of Tsukuba
tatebe@cs.tsukuba.ac.jp

University of Tsukuba
Center for Computational Sciences

Gfarm

HPCISS

CHFS

Osamu is a vice director and a professor in the Center for Computational Sciences at the University of Tsukuba. He is also a visiting scientist at RIKEN R-CCS.

FUN
- Playing the violin
- Hiking in low mountains

Mt. Houkyo (461m), Tsukuba on Jul 18, 2022

Univ Tokyo Orchestra played Shostakovich Symphony no.10 on Jan 19, 1989

Background and Research
- I have led the Gfarm distributed file system project for over twenty years. This system is used for 200PB national HPCI shared storage.
- Design and implementation of parallel file systems in HPC

Big Questions for Workshop
- What IT technology is driving AI for Science?
- What are the critical issues of AI for Science?
- What is truly meaningful AI for Science?

Tenindra Abeywickrama - RIKEN R-CCS

tenindra.abeywickrama@riken.jp

I'm a newly minted Research Scientist at RIKEN R-CCS. Previously I was a lead data scientist at Grab (the Uber of SE Asia)

Fun!

- Cycling (mainly Zwift lately), hiking, the gym
- Travel and miles chasing (next Shikoku, Chongqing, Ko Samui *fingers-crossed*)
- Escape rooms, board games, the local F&B scene, enthusiastic but offkey karaoke

Background:

- Spatial data management: indexing and querying techniques for POI, trajectory, road network datasets
- Graph heuristics for search, shortest path, combinatorial problems (TSP, VRP, etc)
- Deployed my research: dispatch in a major ride-hailing app (6 billion+ annual trips)

Big Questions for Workshop

- How do we close the gap between peak processing power and actual utilization in common graph operations?
- Which forgotten problems/algorithms should we revisit in anticipation of the post-exascale era?

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN

John Towns
NCSA/University of Illinois
jtowns@illinois.edu

NSF | ACCESS | LEADERSHIP-CLASS COMPUTING FACILITY

NAIRR Pilot National Artificial Intelligence Research Resource Pilot

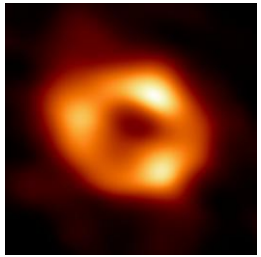

I am the Deputy Director at the National Center for Supercomputing Applications (NCSA). His history on HPC goes back 35+ years evolving from being a user of large systems, to building them, to building national collaborations providing access and support.

For fun: spending time with 5 grandchildren, brewing beer, gardening, reading, camping,

Background and Research

- Astrophysicist by training; building research infrastructure these days

Current interests:

- Practical applications of AI
- Prototyping practically useful, hybrid HPC-quantum computing environments

Big Questions for Workshop

- Will AIs fail, take over, or be practically useful (not necessarily mutually exclusive)?
- What will be the path to practical use and adoption of quantum computing?

Miwako Tsuji

[University of Tsukuba CCS, RIKEN R-CCS miwako.tsuji@riken.jp]

筑波大学 計算科学研究センター Center for Computational Sciences

RIKEN 理化学研究所 R-CCS 計算科学研究センター RIKEN Center for Computational Science

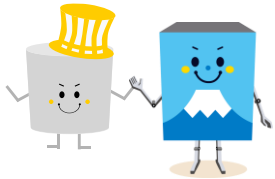

Unit Leader of Quantum HPC Hybrid Software Environment Unit, RIKEN R-CCS
Professor, Division of High Performance Computing Systems, CCS, University of Tsukuba

FUN!!!
camera, photo
pokemon-go

- Background and Research
  - Programming Model
  - Performance Model
  - Quantum HPC Hybrid
- Big Questions for Workshop
  - Let’s talk more about Quantum HPC

Brian Van Essen - Lawrence Livermore National Laboratory

Lawrence Livermore National Laboratory
vanessen1@llnl.gov

CASC Center for Applied Scientific Computing
NNSA National Nuclear Security Administration

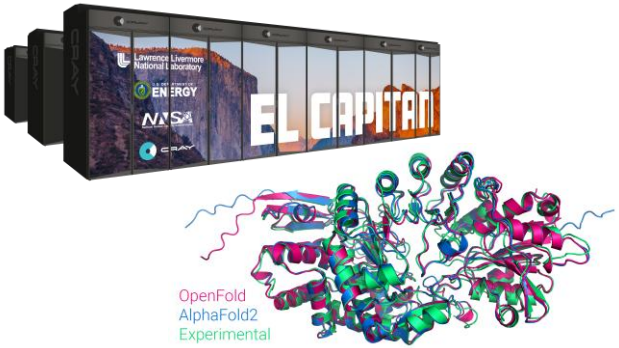

Brian is a Senior Computer Scientist & Architect in LC ATO & CASC. He leads research into leadership-class AI. Prior to his PhD he co-founded two reconfigurable computing startups (which failed).

FUN!!!

- Backpacking, kayaking, cross-country skiing
- Scoutmaster, Philmont (x2), Northern Tier
- Astrophotography, astronomy, ham radio
- Recently rescued a Husky

Background and Research

- AI at leadership-class HPC scale: developing tools and algorithms
- Foundation models for molecular discovery
- Cognitive Simulation: coupling HPC ModSim with AI augmented multi-physics applications
- Integrating AI accelerators into HPC systems and centers
- Mapping scientific algorithms to spatial dataflow accelerators (AI HW)

Big Questions for Workshop

- Broadening the AI HW/SW ecosystem
- Development of robust complex frontier AI workflows
- HPC+AI leadership-class center of the future

Emil **VATAI**, Riken **R-CCS**

High Performance AI systems reseach team, emil.vatai@{riken.jp,gmail.com}, **vatai.github.io**

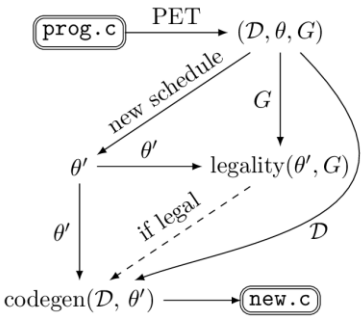

From SR RS, MSC+PhD in HU HU, CS+Math

Father 👶 Judo 🥋 Emacs 🐂 Maths
Gamer 🕹️ (when time allows)

Background:

- CS with *a lot of maths.*

Previous research:

- Comp. number theory (prime world record)
- Communication avoiding algorithms

Current Research

- **Tadashi: AI driven correct code transformation**
- RadicalPy: spin dynamics simulations
- Other projects: Sequence alignment, Carcinogenesys, QuantOm

Jeffrey Vetter

Oak Ridge National Laboratory
vetter@computer.org, +1-865-356-1649, https://vetter.github.io/

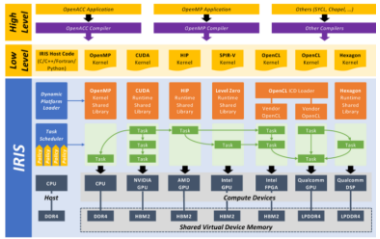

IRIS runtime architecture [6]

https://iris-programming.github.io/
https://github.com/ORNL/iris

Jeff leads the section for Advanced Computing Systems Research at ORNL: ~55 staff focusing on architectures, programming systems, distributed systems, research software engineering, etc.

FUN!!!: travel, maker, farming, etc.

- Background and Research
- Emerging architectures: GPUs, FPGAs, SoCs, Analog
- Programming systems:
  - OpenACC, OpenCL, Kokkos,
  - LLVM, IRIS Runtime, etc.
  - Julia, Rust, Python
- AI: ChatBLAS, ChatHPC, CodeLLAMA

- Big Questions for Workshop
  - How do we build software for extremely heterogeneous systems (think way beyond GPUs) over the next two decades?
  - What architectures will help us manage/mitigate the unprecedented power requirements over the next two decades?
  - Will AI help us manage the software ‘problem’ of porting, maintaining, accelerating applications?

Jean-Pierre Vilotte & CNRS (font size ≥20)

[jean-pierre.vilotte@cnrs-dir.fr]

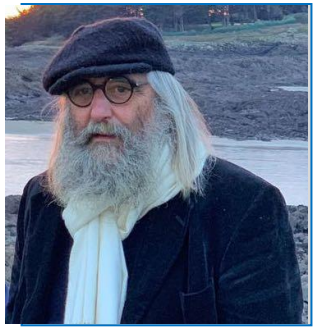

Dr. Jean-Pierre Vilotte (res.ID H-1552-2017), PhD, is Professor at the Institut de Physique du Globe de Paris (CNRS-UMR 7154), and Scientific Deputy at the French CNRS- INSU for HPC/HPDA/AI and open science. He is a member of the steering committee of the AI for Science and Science for AI (AISSAI) CNRS Centre. In 2015, he was awarded Fellow of the American Geophysical Union. He is actively involved in the French NumPEx program as co-PI of the Exa-DI project, in the international SKA project and in the French Data Terra E-Infrastructure for Earth System and Climate.

Fun !!!

Sailing & Jazz

- International expertise at the interface between computational geophysics, mathematics, statistical physics, AI and data sciences.
- How to develop international collaboration in exascale software co-design and co-development, shared benchmarking methodologies including complex AI-based HPC/HPDA workflows

Barr von Oehsen, Pittsburgh Supercomputing Center

barr@psc.edu

Barr is the Director of the Pittsburgh Supercomputing Center (PSC), a joint venture between the University of Pittsburgh (Pitt) and Carnegie Mellon University (CMU). He holds research professor appointments in electrical and computer engineering (Pitt, CMU) and biomedical informatics (Pitt).

FUN!!!

- Cycling and hiking
- Cooking
- Traveling
- Reading – many genres
- Working with the ISC 2025 Student Team

Background and Research

- Barr's NSF and NIH-funded research focuses on advanced computing, secure federated platforms, networking, edge computing, and workforce development.

Big Questions for Workshop

- Inference vs training, should we focus on both or one over the other?
- Edge computing, distributed data and AI resources, quantum - how do these fit into a national ecosystem?
- Cloud computing/industry – is this our future?
- What is the future of university on-prem resources?
- What new benchmarks should we develop/use (agree on) to better capture our hybrid computing reality?

Rich Vuduc – Georgia Tech

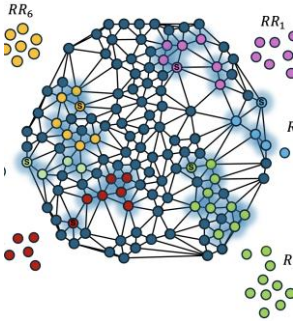

I am a professor at the School of Computational Science & Engineering and a co-director of CRNCH (“crunch”), a research center on post-Moore's Law computing.

Background & Research

- My lab, The HPC Garage, is performance analysis, performance engineering, and algorithm-architecture codesign, particularly for sparse linear algebra and graph computations.

Questions (big and small) **for the Workshop**

- *Foundations* — Is there a power/energy gap between biological and artificial/computational information processing? (sensorimotor + intelligence)
- *HPC+AI* — What (new?) computational workloads, proxy apps, or kernels should drive the design of future machines?

Mohamed Wahib, RIKEN-CCS

mohamed.attia@riken.jp

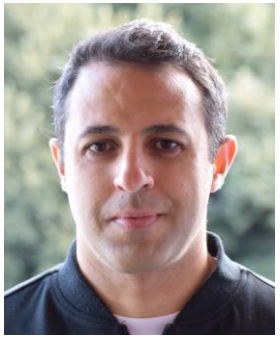

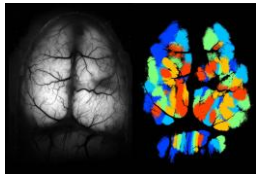

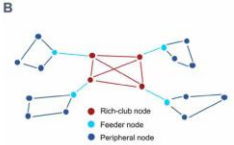

Mohamed Wahib is the team leader of the high-performance AI systems research team at RIKEN Center for Computational Science. His research interests are applying HPC in AI, and AI in HPC.

Fun:

• Crushing skulls!

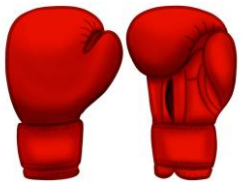

Background & Research

•Programming systems

•Large-scale AI: at different levels in HPC systems

•AI in science

•High-performance imaging applications

Big Questions for Workshop

•Can scientific data make use of LLMs?

•Can AI be seamlessly integrated with scientific sim?

•Correctness and uncertainty quantification?

Outcomes

•Identify major challenges for adopting AI in science

•Aligning AI4Science efforts between institutions

Michèle Weiland

EPCC, The University of Edinburgh
m.weiland@epcc.ed.ac.uk

Michèle is a Professor at EPCC and the Met Office Joint Chair for the University of Edinburgh

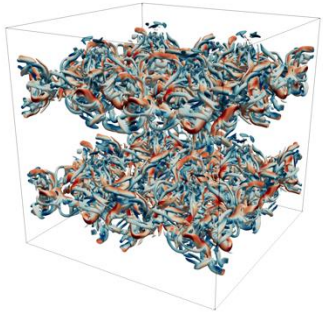

Fun: Gardening, camping, skiing, indoor rowing, cycling, gaming

Background

- Computer Science, ML & Music

Research focus

- Engineering (CFD & combustion) & weather/climate applications, meshing
- Sustainability & efficiency → system workload analysis
- Collaborative project with NCAR towards net zero in atmospheric sciences

Big Questions for Workshop

- AI is not magic, modelling & simulation is here to stay – how do we support a nuanced debate?

Stefan Wild

Lawrence Berkeley National Laboratory, wild@lbl.gov

BERKELEY LAB
Bringing Science Solutions to the World

AMCR

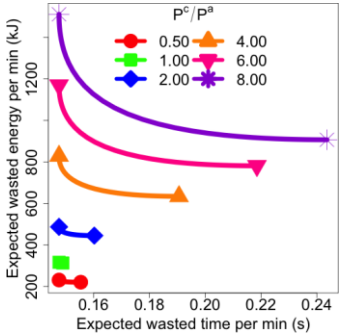

Stefan is director of the Applied Mathematics and Computational Research (AMCR) Division and is a Sr Scientist who solves challenging optimization and automated learning problems

skiing
scuba
alpine cheese loud music

Interests for this workshop

- AI-/data-driven acceleration of rigorous numerical schemes
- Optimization for next-generation efficiency
- Composable tools and outer loop methods
- Optimal design of experiments and automated science lab
- Pervasiveness of AI tools for productivity

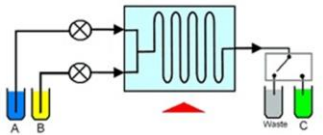

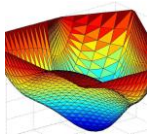

Rio Yokota

Institute of Science Tokyo / RIKEN R-CCS
rioyokota@rio.scrc.iir.isct.ac.jp

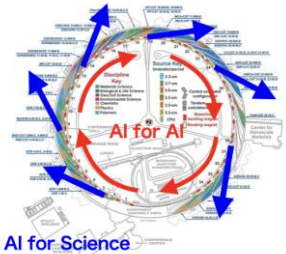

Rio is a Professor at the Supercomputing Research Center, Institute of Science Tokyo and Team Principle of the AI for Science Foundation Model Research Team at RIKEN R-CCS

Fun:

I used to do a lot of scuba diving when I was at KAUST, but not after moving to Japan.

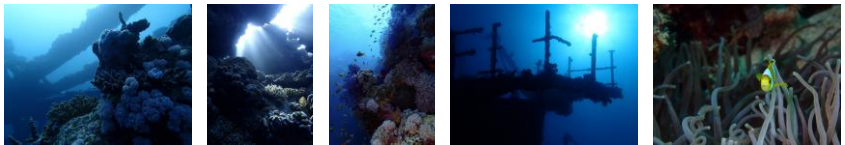

Background and Research

- I work at the intersection of HPC and AI
- Pre-training Japanese LLMs
- Synthetic data for vision models
- DGEMM Emulation on Tensor Cores

Big Questions for Workshop

- What are the evals for AI for Science?
- Pre-training vs Post-training vs In-context
- Do we really need to train on that data?
- What is the minimal reasoning model?
- How to collaborate with top AI companies?
- How can HPC apps utilize Tensor Cores?

Kazuki Yoshizoe

Kyushu University
yoshizoe@cc.kyushu-u.ac.jp

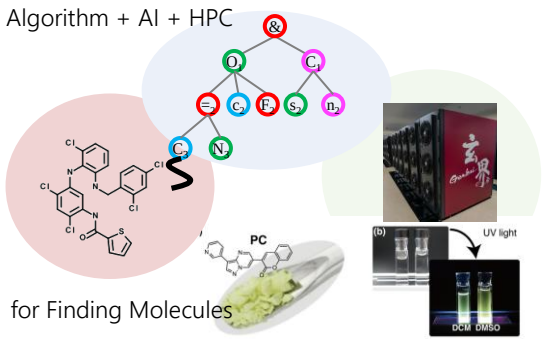

Kazuki leads the HPC department of Kyushu-U as a Professor in Research Institute of Information Technology (RIIT). He works in combining HPC with machine learning and algorithms.

Playing Go for 20+ years

Love German Black Beer

Recently Started: jogging, Watch Kabuki

Background and Research

- Parallel graph search algorithms
- Search + ML, mainly for game AI
- Currently combining parallel graph search, generative AI, and computational chemistry simulations to discover new molecules
- Works on HPC+AI at Kyushu University (formerly at RIKEN)

Big Questions for Workshop

- For HPC+AI workloads, how can we best provide flexible resources considering varying demands (CPU- GPU ratio, memory size, storage, etc.)?
- What kinds of systems are most suitable for “LLM + something” workloads?

Tomoya Yuki

RIKEN Center for Computational Science (R-CCS)
tomoya.yuki@riken.jp

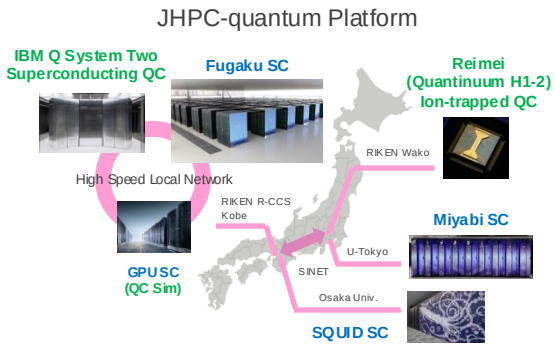

Tomoya is a Technical Scientist in the Quantum-HPC Hybrid Platform Operations Unit, Quantum-HPC Hybrid Platform Division, at RIKEN R-CCS

FUN:

- Fidget toys, Pocket tops
- Psychedelic Trance (Computer Music)
- Drawing, Oil Painting

Background and Research

- Working on building the Quantum-HPC Hybrid Platform, focusing on authentication, job scheduling, and frontend development
- Modeling microarchitectures, cache systems, and memory hierarchies (Previous work)

Big Questions for Workshop

- How can we simplify heterogeneous, complex systems into a unified platform while ensuring both high performance and usability?
- Can we harness the benefits of AI at the system layer - covering configuration, orchestration, and monitoring - by moving from automation toward autonomation?