



ECRC Communication Meeting

29 September 2021



جامعة الملك عبد الله
للعلوم والتقنية
King Abdullah University of
Science and Technology

Extreme Computing
Research Center

Plan of today's meeting

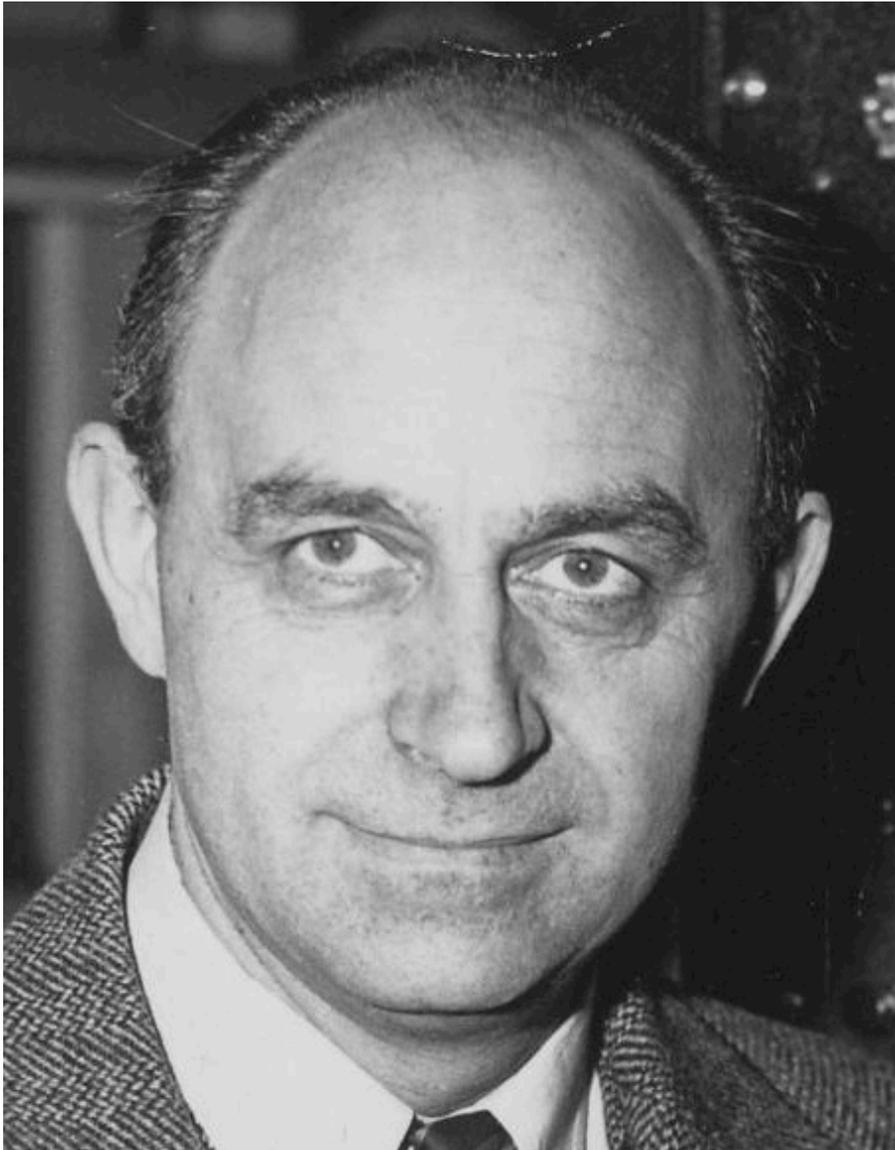
- People
 - new members, visitors, interns, departing
- Announcements
 - Recent, near-term & semester activity overview
- Communication meeting calendar and content
- Blitz presentations of ECRC sponsored projects
- Roundtable

Happy Birthday to one of the GPU Boys!



Fermi (2010)

Enrico Fermi



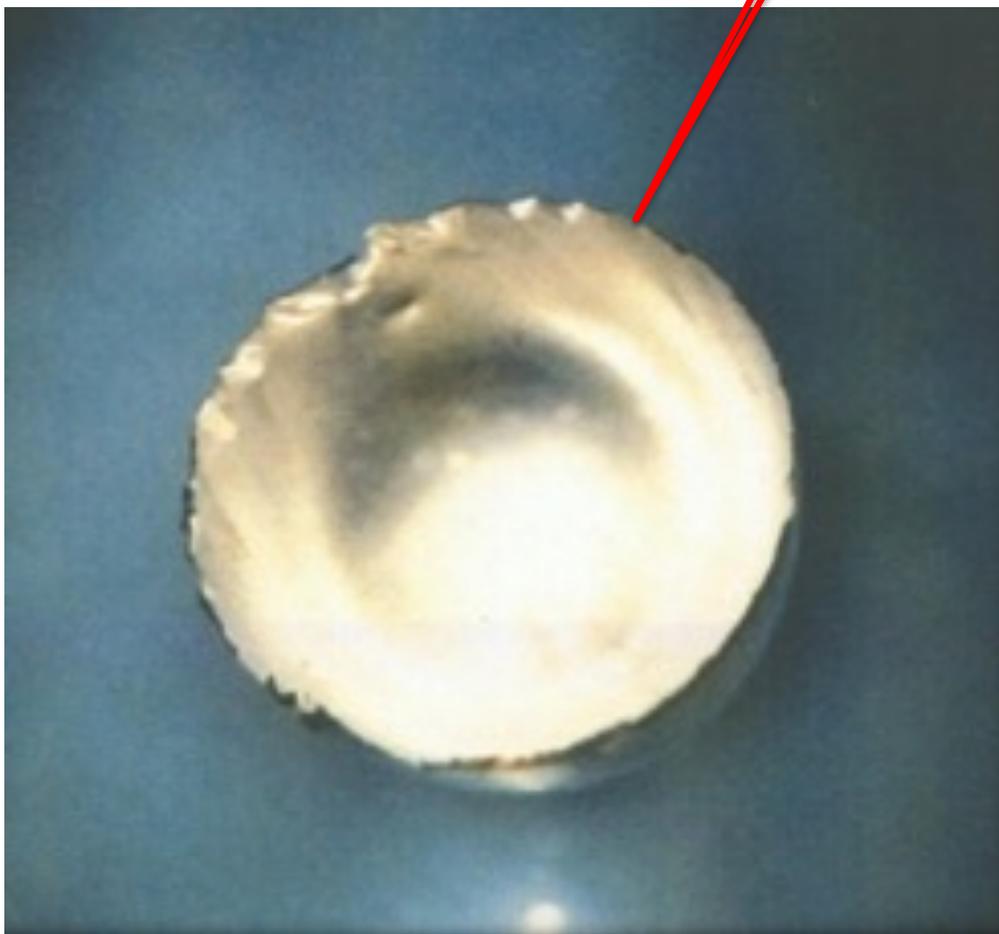
- Born **29 September 1901** Rome
 - *Laurea*, Scuola Normale Superiore, 1922
 - Nobel Prize for transuranic elements, 1938
 - Università di Roma, *Sapienza*, 1926-1939
 - Columbia University, 1939-1944
 - University of Chicago, 1944-1954
- Fermi was perhaps the last physicist who was equally respected in experiment and a theory. Though he died in 1954, he was an early computing pioneer.

Fermium, **Fm₁₀₀**, is the last stable element. “Fermions” include quarks and leptons (of which there are 24 kinds) and anything that can be built from them, such as protons and neutrons. So most of the matter in the universe, including you, is named after Fermi 😊.

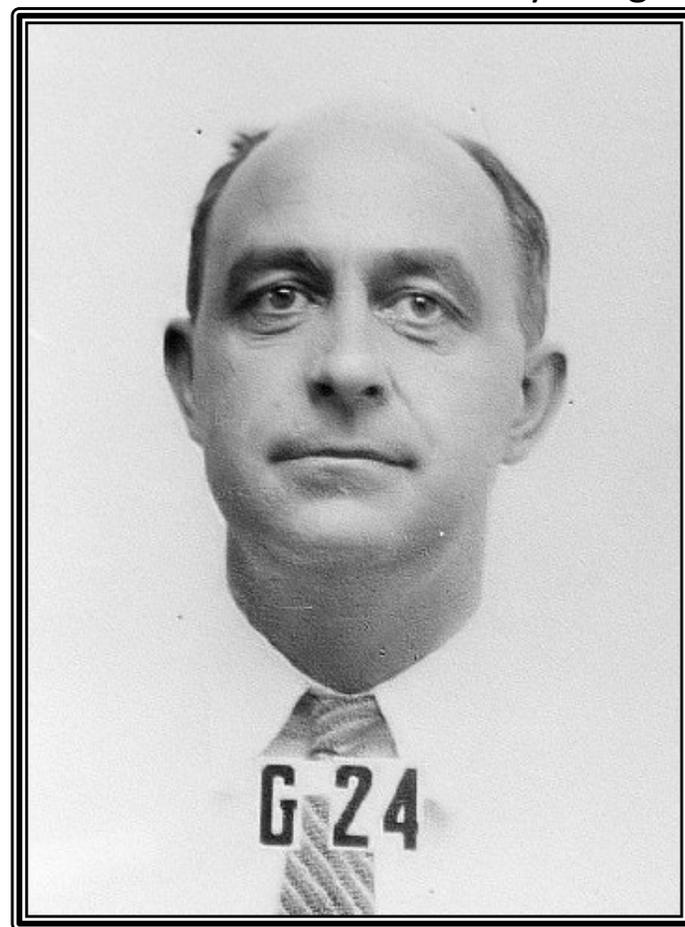
So is the DOE’s Fermi Laboratory complex in Illinois.

Fermium #100

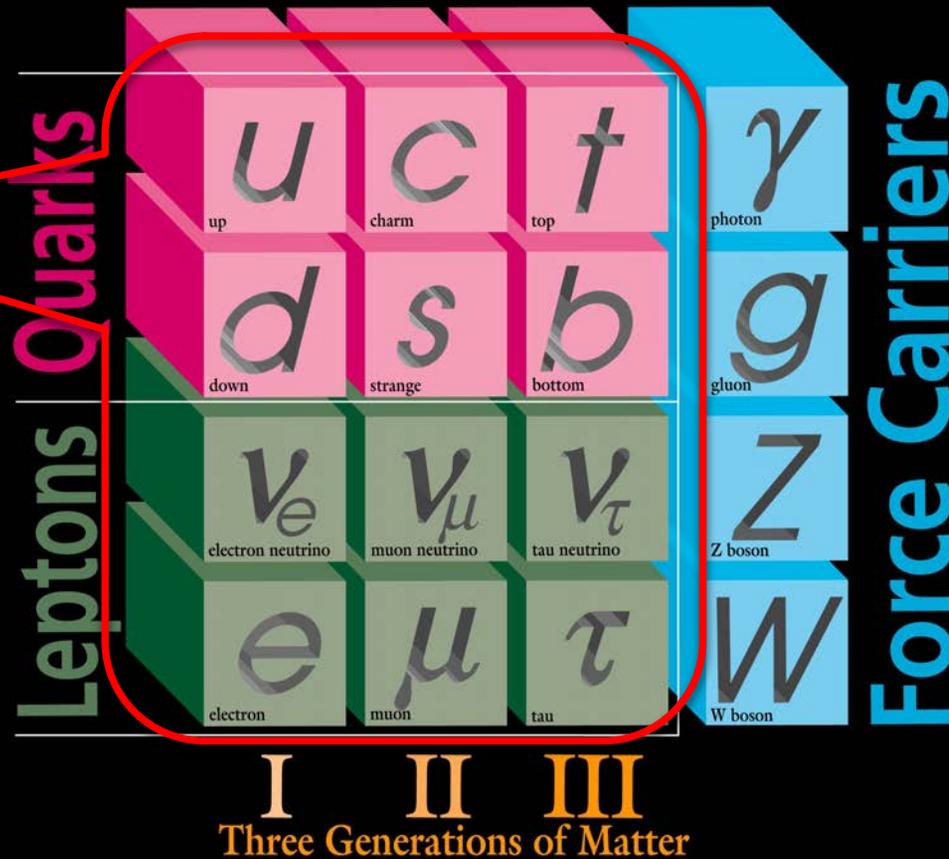
H																	He														
Li	Be											B	C	N	O	F	Ne														
Na	Mg											Al	Si	P	S	Cl	Ar														
K	Ca											Sc	Ti	V	Cr	Mn	Fe	Co	Ni	Cu	Zn	Ga	Ge	As	Se	Br	Kr				
Rb	Sr											Y	Zr	Nb	Mo	Tc	Ru	Rh	Pd	Ag	Cd	In	Sn	Sb	Te	I	Xe				
Cs	Ba	La	Ce	Pr	Nd	Pm	Sm	Eu	Gd	Tb	Dy	Ho	Er	Tm	Yb	Lu	Hf	Ta	W	Re	Os	Ir	Pt	Au	Hg	Tl	Pb	Bi	Po	At	Rn
Fr	Ra	Ac	Th	Pa	U	Np	Pu	Am	Cm	Bk	Cf	Es	Fm	Md	No	Lr	Rf	Db	Sg	Bh	Hs	Mt	Ds	Rg	Cn	113	114	115	116	117	118



Fermi's Los Alamos security badge



ELEMENTARY PARTICLES



There are 12 Fermions:
6 quarks and 6 leptons.
Electrons (e^-) are leptonic Fermions, so you are actually full of Fermions.

Fermions have half-integer spin and obey the exclusion principle, as opposed to Bosons (named for Satyendra Bose), which have integer spin (such as the Higgs Boson discovered in 2012 at CERN).

Fermi Laboratory, Illinois



Fermi Laboratory, Illinois



Fermi's letter to the University of Pisa

11 August 1954

Professor Enrico Avanzi, Rector of the University of Pisa

Dear Professor,

On the occasion of my stay at the Varenna School, Professors Conversi and Salvini mentioned that the University of Pisa might have at its disposal a large amount of money to use for the development and progress of Italian research. On being questioned about the various possibilities to use such funds. I thought that the idea of building an electronic computer in Pisa was by far the best.

An electronic computer would constitute a research instrument from which all science and research activities would profit, in such a manner that is currently inestimable.

*I know that the Institute for Applications of Computation, headed by Professor Picone, is in the process of buying a machine of the same kind. However, I do not think that this event would reduce the need for such a machine for a centre of study like the University of Pisa. **Experience shows that the possibility of executing complex calculations with great speed and accuracy quickly creates such an enormous demand for these services which would soon be beyond the capacity of just one machine. Furthermore, there are the advantages for the students and researchers who will be trained in the use of these new tools for computation.***

Yours sincerely,
Enrico Fermi

Quotations by Fermi

“Before I came here I was confused about this subject. Having listened to your lecture I am still confused. But on a higher level.”

Quotations by Fermi

“There are two possible outcomes: if the result confirms the hypothesis, then you’ve made a measurement. If the result is contrary to the hypothesis, then you’ve made a discovery.”

Quotations by Fermi

“If I could remember the names of all these particles, I'd be a botanist.”

Quotations by Fermi

When asked what characteristics Nobel Prize winning physicists had in common:

*“I can't think of a single one.
Not even intelligence.”*

Fermi, the mentor

Fermi's students include:

- T. D. Lee
- C. N Yang
- Mildred Dresselhaus
- Marvin Goldberger

Eight of Fermi's students or close mentees went on to win Nobel Prizes.

But we celebrate in the ECRC his prescient ideas about scientific computing.

Promotion 😊

Sumathi has taken over the corner office as ECRC Executive Secretary. Sumathi is currently on a short vacation leave in India.

Linzhi Gu has been recruited as our new Administrator and Workshop Coordinator.



Other introductions?

By self?

By sponsor?

ECRC @ HPEC'21

StressBench: A Configurable Full System Network and I/O Benchmark Framework [Best Paper Award]

Dean G Chester (University of Warwick); Taylor Groves (NERSC); Simon Hammond (Sandia National Laboratories); Tim Law (AWE); Steven Wright (University of York); Richard Smedley-Stevenson (AWE); **Suhaib A. Fahmy (KAUST)**; Gihan Mudalige (University of Warwick); Stephen Jarvis (University of Birmingham)

Implications of Reduced Communication Precision in a Collocated Discontinuous Galerkin Finite Element Framework

Marcin Rogowski; Lisandro Dalcin; Matteo Parsani; David Keyes (KAUST)

ECRC @ HPEC'21

StressBench: A Configurable Full System Network and I/O Benchmark Framework [Best Paper Award]

Dean G Chester (University of Warwick); Taylor Groves (NERSC); Simon Hammond (Sandia National Laboratories); Tim Law (AWE); Steven Wright (University of York); Richard Smedley-Stevenson (AWE); **Suhaib A. Fahmy (KAUST)**; Gihan Mudalige (University of Warwick); Stephen Jarvis (University of Birmingham)

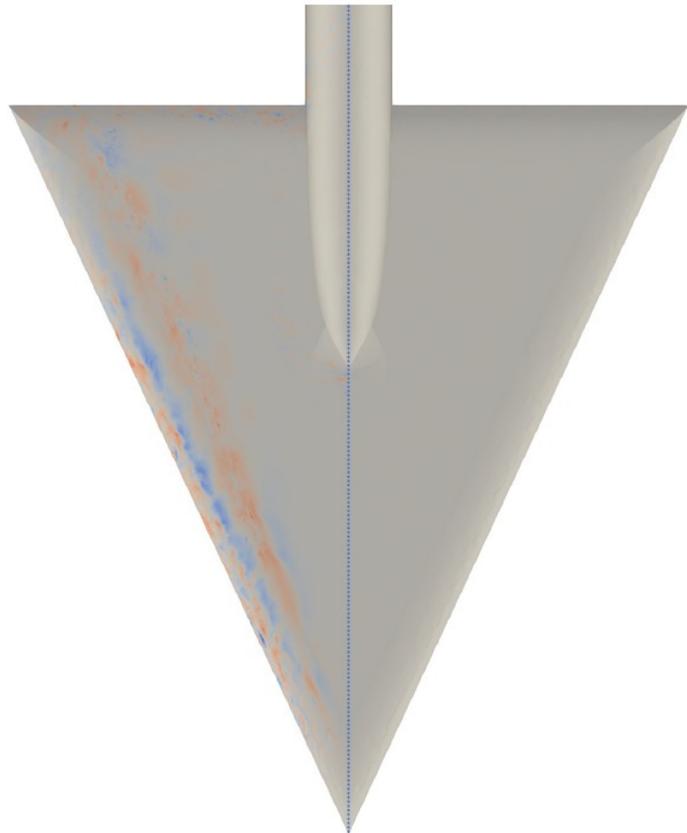
Implications of Reduced Communication Precision in a Collocated Discontinuous Galerkin Finite Element Framework

Marcin Rogowski; Lisandro Dalcin; Matteo Parsani; David Keyes (KAUST)

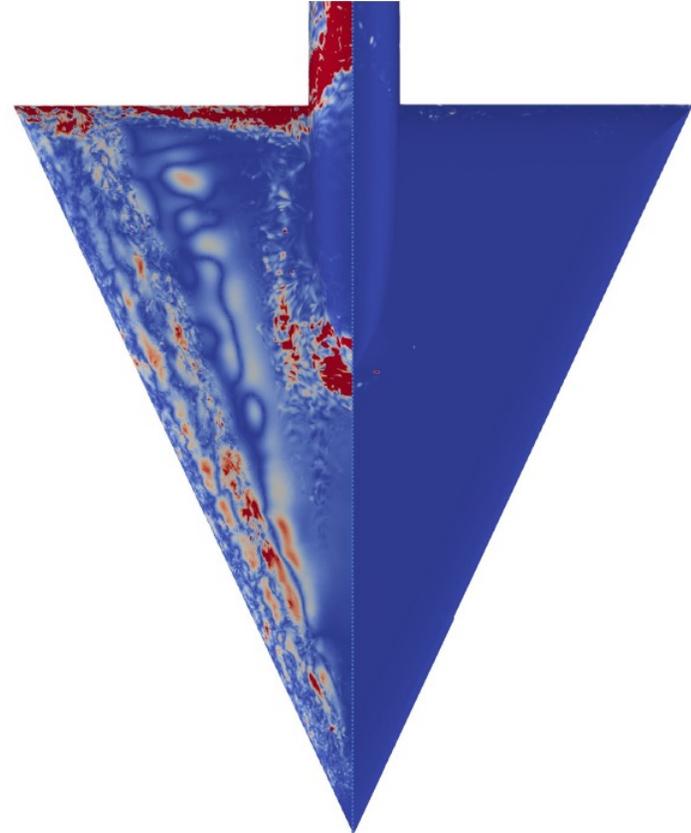
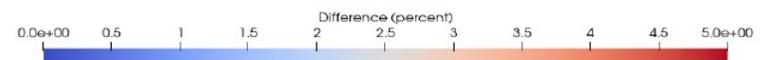
StressBench

embargoed

Low-precision communication



FP64 \longleftrightarrow FP32 FP64 \longleftrightarrow FP64
+FP32 halo



FP64 \longleftrightarrow FP32 FP64 \longleftrightarrow FP64
+FP32 halo

SANDS Lab – Marco Canini

SIGCOMM '21

Efficient Sparse Collective Communication and its application to Accelerate Distributed Deep Learning

Jiawei Fei*
NUDT
KAUST

Chen-Yu Ho*
KAUST

Atal Narayan Sahu
KAUST

Marco Canini
KAUST

Amedeo Sapia
Intel

[KAUST Discovery highlight](#) | [Video presentation](#)

All-Reduce and other collectives are well known to HPC folks. Recently discovered by AI/ML to run data parallel training.

[OmniReduce](#) natively supports sparse inputs and adopts [SwitchML](#)'s streaming aggregation. It's currently the fastest All-Reduce method we know of. Implemented using RDMA. Open sourced and adopted in production by [Meituan](#) (aka Uber Eat in China) where Jiawei is now for an internship.

SOSP '21

LineFS: Efficient SmartNIC Offload of a Distributed File System with Pipeline Parallelism

Jongyul Kim
KAIST

Insu Jang*
University of Michigan

Waleed Reda
KTH Royal Institute of Technology
Université catholique de Louvain

Jaesung Im
KAIST

Marco Canini
KAUST

Dejan Kostić
KTH Royal Institute of Technology

Youngjin Kwon
KAIST

Simon Peter
The University of Texas at Austin

Emmett Witchel
The University of Texas at Austin
Katana Graph

Simon gave a [CS seminar in Spring'21 on Assise](#), the predecessor to LineFS.

Both are distributed FSES that co-locate persistent memory (aka NVM) with the application. The paper offloads a host of FS operations to a BlueField SmartNIC, which happens to be on the data path anyway. The main benefit is freeing up CPU cycles but also it beats Assise.

NeurIPS '21 (less than 3% of 9122 submissions)

Rethinking gradient sparsification as total error minimization

**Atal Narayan Sahu, Aritra Dutta, Ahmed M. Abdelmoniem,
Trambak Banerjee[†], Marco Canini, Panos Kalnis**
KAUST, [†]University of Kansas

Introduce a new perspective to think about gradient sparsification in the context of data parallel ML training in which aim to optimize for the optimality for the entire training instead of the per-iteration optimality.

Advocate for absolute compression (based on a threshold) that adapts the amount of data volume exchanged at each iteration so as to minimize the total error under a communication budget. Solid theory work backed up by extensive evaluation.

KAUST Research Open Week

(second last week of semester)

KAUST Research Open Week Draft Agenda

November 28 – December 2, 2021

Theme: **Sustainability: Science for the Future**



	Sunday Nov 28				Monday Nov 29				Tuesday Nov 30				Wednesday Dec 1				Thursday Dec 2							
Thrust	Water				Food & Health				Environment				Energy				Digital							
Open Doors: All week each KAUST Core Labs and Research Centers are welcoming delegates to come to through their doors for interactive and engaging tours.																								
09:00	Auditorium Bldg. 20 Welcome: President Tony Chan & Vice President for Research Donal Bradley External Keynote Lecture: TBD				Auditorium Bldg. 20 External Keynote Lecture TBD				Auditorium Bldg. 20 External Keynote Lecture TBD				Auditorium Bldg. 20 External Keynote Lecture TBD				Auditorium Bldg. 20 External Keynote Lecture TBD							
10:00	Bldg. 19 Scientific Demonstrations	Auditorium Bldg. 2&3 Talks	Outside Library Science on the Spine	Auditorium Bldg. 4&5 Talks	Bldg. 19 Scientific Demonstrations	Auditorium Bldg. 2&3 Talks	Outside Library Science on the Spine	Auditorium Bldg. 4&5 Talks	Bldg. 19 Scientific Demonstrations	Auditorium Bldg. 2&3 Talks	Outside Library Science on the Spine	Auditorium Bldg. 4&5 Talks	Bldg. 19 Scientific Demonstrations	Auditorium Bldg. 2&3 Talks	Outside Library Science on the Spine	Auditorium Bldg. 4&5 Talks	Bldg. 19 Scientific Demonstrations	Auditorium Bldg. 2&3 Talks	Outside Library Science on the Spine	Auditorium Bldg. 4&5 Talks				
11:00																								
12:00	Library Poster Session & Networking Lunch				Library Poster Session & Networking Lunch				Library Poster Session & Networking Lunch				Library Poster Session & Networking Lunch				Library Poster Session & Networking Lunch							
	Mosti, Bldg. 19 KAUST Smart Demo & Tour to Smart House				Mosti, Bldg. 19 KAUST Smart Demo & Tour to Smart House				Mosti, Bldg. 19 KAUST Smart Demo & Tour to Smart House				Mosti, Bldg. 19 KAUST Smart Demo & Tour to Smart House				Mosti, Bldg. 19 KAUST Smart Demo & Tour to Smart House							
13:00	Auditorium Bldg. 20 Internal Keynote Lecture: TBD				Auditorium Bldg. 20 Internal Keynote Lecture: TBD				Auditorium Bldg. 20 Keynote Lecture: Donal Bradley				Auditorium Bldg. 20 Internal Keynote Lecture: TBD				Auditorium Bldg. 20 Internal Keynote Lecture: TBD							
14:00	Bldg. 19 Scientific Demonstrations	Auditorium Bldg. 2&3 Talks	Outside Library Science on the Spine	Auditorium Bldg. 4&5 Talks	Bldg. 19 Scientific Demonstrations	Auditorium Bldg. 2&3 Talks	Outside Library Science on the Spine	Auditorium Bldg. 4&5 Talks	Bldg. 19 Scientific Demonstrations	Auditorium Bldg. 2&3 Talks	Outside Library Science on the Spine	Auditorium Bldg. 4&5 Talks	Auditorium Bldg. 20 Visit from McLaren: Q&A with Lando Norris				Bldg. 19 Scientific Demonstrations	Auditorium Bldg. 2&3 Talks	Outside Library Science on the Spine	Auditorium Bldg. 4&5 Talks				
15:00													Auditorium Bldg. 20 KAUST Sci-Café: KAUST X McLaren Collab: Fostering a Sustainable Future for F1											
16:00	Research Center Open Doors	Auditorium Bldg. 2&3 Talks	Core Lab Open Doors	Auditorium Bldg. 4&5 Talks	Research Center Open Doors	Auditorium Bldg. 2&3 Talks	Core Lab Open Doors	Auditorium Bldg. 4&5 Talks	Research Center Open Doors	Auditorium Bldg. 2&3 Talks	Core Lab Open Doors	Auditorium Bldg. 4&5 Talks	Research Center Open Doors	Auditorium Bldg. 2&3 Talks	Core Lab Open Doors	Auditorium Bldg. 4&5 Talks	Research Center Open Doors	Auditorium Bldg. 2&3 Talks	Core Lab Open Doors	Auditorium Bldg. 4&5 Talks				
17:00																								
Installation					Evening Activities: SME Interactive Session with KAUST Startups					Evening Activities: National Poster Session					Evening Activities: KAUST School F1 Racing					Evening Activities: Closing Ceremony				

Please note this agenda is subject to change.

KAUST Research Open Week

KROW is a campus-wide event: external keynote speakers, talks by KAUST researchers, posters, demonstrations, *Science on the Spine*, and visit from McLaren on Dec 1.

Open Doors: All Research Centers should take part in “Open Doors” for one day during the week. This should showcase your center to the wider research community and KAUST community, including the KAUST schools. The responsibility lies with each center to organize their Open Doors. **Suggestion:** parents of TKS students take the lead. Assistance available from Naadiya Carrim.

So far, Matteo Parsani and I are down to share the “popular” lecture period on Thursday, 2 December at 4pm. Others interested in presenting at a lay person’s level, let me know!

IEEE IPDPS Lyon May-Jun '22

<https://www.ipdps.org/>
(abstracts due 8 Oct;
papers due 15 Oct)

ECRC has so far up to six submissions planned



IPDPS
2022 Lyon, France
Lyon, France • May 30 - June 3, 2022
ipdps.org

36th IEEE International Parallel and Distributed Processing Symposium

CALL FOR PAPERS

The five-day IPDPS program includes three days of contributed papers, invited speakers, industry participation, and student programs, framed by two days of workshops with peer reviewed papers that complement and broaden the main program. In 2022, IPDPS will return to meeting in person and will do so in the European metropolis of Lyon. For more details, visit the website at ipdps.org.

Authors for the main conference are invited to submit manuscripts that present original unpublished research in all areas of parallel and distributed processing. Work focusing on emerging technologies and interdisciplinary work covering multiple IPDPS focus areas is especially welcome. Topics of interest include:

- Parallel and distributed computing theory and algorithms (**Algorithms**)
- Experiments and practice in parallel and distributed computing (**Experiments**)
- Programming models, compilers and runtimes for parallel applications and systems (**Programming Models & Compilers**)
- System software and middleware for parallel and distributed systems (**System Software**)
- **Architecture**
- **Multidisciplinary**

Abstracts due	October 1, 2021
Submissions due	October 8, 2021
Author response/rebuttal	November 23-24, 2021
First round decisions	December 3, 2021
Revised submissions due	January 5, 2022
Final notification	January 15, 2022

Sponsored by






GENERAL CO-CHAIRS
Anne Benoit (ENS Lyon, France)
Laurent Lefèvre (Inria & ENS Lyon, France)

WORKSHOPS CHAIR AND VICE-CHAIR
Ananth Kalyanaram (Washington State University, USA)
Suren Byna (Lawrence Berkeley National Laboratory, USA)

PROGRAM CO-CHAIRS
Yves Robert (ENS Lyon, France & Univ. of Tennessee, USA)
Bora Uçar (CNRS, Laboratoire LIP, Lyon, France)

PROGRAM AREA CHAIRS & VICE-CHAIRS

- **Algorithms:**
Gagan Agrawal (Augusta University, USA)
Sherry Li (Lawrence Berkeley National Laboratory, USA)
- **Experiments:**
Kengo Nakajima (University of Tokyo/RIKEN R-CCS, Japan)
Keita Teranishi (Sandia National Laboratories, USA)
- **Programming Models & Compilers:**
Didem Unat (Koç University, Turkey)
Mohamed Wahib (AIST, Japan)
- **System Software:**
Thomas Herault (University of Tennessee, Knoxville, USA)
Ana Gaiaru (Oak Ridge National Lab, USA)
- **Architecture:**
Radu Teodorescu (Ohio State University, USA)
Mengjia Yan (MIT, USA)
- **Multidisciplinary:**
Alba Cristina Melo (University of Brasilia, Brazil)
Sunita Chandrasekaran (University of Delaware, USA)

IPDPS 2022 VENUE
Lyon, France is a major science and technology center where Joseph Marie Jacquard designed the first mechanical weaving machine - a precursor of the first computer! Lyon still bears witness to 2,000 years of history, enriched by the signatures of modern architects and the dynamics of a full menu of cultural attractions, parks and rivers, and green city life. Located in the South-East of France, Lyon is 2 hours by TGV from Paris and less than 2 hours from the Mediterranean coast and the airport connects to one hundred plus international destinations.

IEEE/ACM Supercomputing Conference

<https://sc21.supercomputing.org/>

The image shows the homepage of the IEEE/ACM Supercomputing Conference website. At the top, a banner reads "No Travel? No Problem." with a "REMOTE PARTICIPATION →" button. Below this is a dark navigation bar with links for "FIRST-TIMERS", "EXHIBITORS", "STUDENTS", "MENTORS", and "MEDIA". The main header features the SC21 logo and navigation links for "PROGRAM", "EXHIBITS", "SCINET", "ATTEND", and "SUBMIT", along with a "REGISTER →" button. The main content area is divided into several sections: "The International Conference for High Performance Computing, Networking, Storage, and Analysis" with location information for America's Center in St. Louis, MO; "Program" (November 14-19) and "Exhibits" (November 15-18) sections; a "science & beyond." section with a message from the SC21 General Chair; a "Health & Safety" section with a link to review safety protocols; "Learning Is Everything" section encouraging attendance at the HPC technical program; "Check It Out, Students" section promoting the Students@SC program; and a "Megawatts of Petaflops" section highlighting SC's high-capacity network. The SCinet logo is visible in the bottom right corner. The design uses a color palette of dark blue, light green, orange, and white, with various geometric shapes and icons.

No Travel? No Problem. [REMOTE PARTICIPATION →](#)

FIRST-TIMERS EXHIBITORS STUDENTS MENTORS MEDIA

SC21 PROGRAM EXHIBITS SCINET ATTEND SUBMIT [REGISTER →](#)

The International Conference for High Performance Computing, Networking, Storage, and Analysis

America's Center
St. Louis, MO →

Program
November 14–19

Exhibits
November 15–18

2021

science & beyond.

A **message** from the SC21 General Chair →

Health & Safety

Your safety is our priority.
Review our **safety protocols.** →

Learning Is Everything

Put your brain into overdrive.
Attend the boldest, broadest **HPC technical program** in the multiverse. →

Check It Out, Students

You're generation next. Get with the **Students@SC program.** →

Megawatts of Petaflops

SC's **high-capacity network.**
Immensely powerful, insanely fast, totally collaborative. →

SCinet

SC'21 keynote



- Vice President and Chief Evangelist for Google
- 2014
- Queen Elizabeth Prize for Engineering, 2013
- Presidential Medal of Freedom, 2005
- National Medal of Technology, 1997
- Coinvented TCP/IP while Assistant Prof at Stanford
- UCLA PhD
- Stanford BS

Vint Cerf

"Computing and the Humanities"

A handwritten signature of Vint Cerf in black ink.

"Father of the Internet" Dr. Vint Cerf will join us as our keynote speaker for SC21. He'll share his perspective on how advanced computing may have a groundbreaking effect on how we can better appreciate and understand the study of languages and literatures, the arts, history, and philosophy.

SC'21 elements

- Top 500 unveiling
 - keynote
 - invited talks
 - selected papers
 - tutorials
 - workshops
 - posters
 - panels
 - emerging technology showcase
 - scientific visualization showcase
 - birds-of-a-feather sessions
 - early career events
- ECRC and KSL
well represented
by panels and
paper this year

IEEE/ACM Supercomputing Conference

<https://sc21.supercomputing.org/>

Early on-site registration closes on 15 October

Late is \$200 additional

Remote on-site registration closes 15 October

Late is \$100 additional

SC'21

- Began in 1988; today largest international conference for scientific supercomputing
- Typically >300 exhibitors from >100 different countries
- KAUST will be exhibiting for 13th time (began in 2008, no exhibit at SC'20)
- Attendee trends:
 - SC'16: 11,100
 - SC'17: 12,753 (+14.9%)
 - SC'18: 13,071 (+2.5%)
 - SC'19: 13,950 (+6.7%)
 - SC'20: 7,440 (virtual)

Participants in a normal year:

- ECRC
- KSL/KVL
- IT/KAUST Smart
- HR
- Admissions
- CEMSE
- PSE
- SNA (KGSP)
- Aramco
- GBC

SC'19 (last physical meeting)



ECRC short course, 11-13 October

Theory and Application of Neural Networks

Professor Jinchao Xu (PennState) will give a self-contained introduction to the theory of the neural network function class and its application to image classification and numerical solution of partial differential equations (tentatively scheduled for the second week of October).

- * Definition of the neural network function class as a generalization of classical finite element functions
- * Deep ReLU neural networks versus the classic piecewise linear finite element functions
- * Classical approximation theory of neural network functions
- * New optimal approximate theory of stable neural network functions
- * Classic machine learning methods: logistic regression and support vector machine
- * Deep learning: convolutional neural networks (CNN) for image classification
- * MgNet: a special CNN obtained from a minor modification of the classical multigrid method
- * Application and error analysis of neural network for numerical solutions of partial differential equations (PDEs)
- * Numerical quadrature and Rademacher complexity analysis
- * Old and new training algorithms for machine learning and numerical PDEs



Jinchao Xu

Professor of Mathematics, Penn State U

Verified email at psu.edu - [Homepage](#)

[Numerical analysis](#) [multigrid](#) [domain](#)
[numerical methods for parti...](#)

TITLE

[Iterative methods by space decomposition and subspace](#)

J Xu
SIAM review 34 (4), 581-613

[Parallel multilevel preconditioners](#)

JH Bramble, JE Pasciak, J Xu
Mathematics of computation 55 (191), 1-22

[Two-Grid Discretization Techniques for Linear and Nonlin](#)

J Xu
SIAM journal on numerical analysis 33 (5), 1759-1777

[A novel two-grid method for semilinear elliptic equations](#)

J Xu
SIAM Journal on Scientific Computing 15 (1), 231-237

[Convergence estimates for multigrid algorithms without re](#)

JH Bramble, JE Pasciak, JP Wang, J Xu
Mathematics of Computation 57 (195), 23-45

TITLE	CITED BY	YEAR
Approximation Properties of Deep ReLU CNNs J He, L Li, J Xu arXiv preprint arXiv:2109.00190		2021
Helicity-conservative finite element discretization for incompressible MHD systems K Hu, YJ Lee, J Xu Journal of Computational Physics 436, 110284	3	2021
Improved Convergence Rates for the Orthogonal Greedy Algorithm JW Siegel, J Xu arXiv preprint arXiv:2106.15000	1	2021
Characterization of the variation spaces corresponding to shallow neural networks JW Siegel, J Xu arXiv preprint arXiv:2106.15002	3	2021
Sharp lower bounds on the approximation rate of shallow neural networks JW Siegel, J Xu arXiv preprint arXiv:2106.14997	1	2021
Multiscale and monolithic arbitrary Lagrangian–Eulerian finite element method for a hemodynamic fluid–structure interaction problem involving aneurysms W Hao, P Sun, J Xu, L Zhang Journal of Computational Physics 433, 110181		2021
ReLU Deep Neural Networks from the Hierarchical Basis Perspective J He, L Li, J Xu arXiv preprint arXiv:2105.04156		2021
A Priori Analysis of Stable Neural Network Solutions to Numerical PDEs Q Hong, JW Siegel, J Xu arXiv preprint arXiv:2104.02903	3	2021
Rademacher Complexity and Numerical Quadrature Analysis of Stable Neural Networks with Applications to Numerical PDEs Q Hong, JW Siegel, J Xu arXiv e-prints, arXiv: 2104.02903	2	2021
Robust BPX preconditioner for the integral fractional Laplacian on bounded domains JP Borthagaray, RH Nochetto, S Wu, J Xu arXiv preprint arXiv:2103.12891		2021
Optimal Approximation Rates and Metric Entropy of ReLU^k and Cosine Networks JW Siegel, J Xu arXiv preprint arXiv:2101.12365	2	2021
An Extended Galerkin analysis in finite element exterior calculus Q Hong, Y Li, J Xu arXiv preprint arXiv:2101.09735		2021

Upcoming external review

All KAUST Centers and Programs are being reviewed

Center reviews have started under the VPR. Program reviews are beginning soon under the Provost.

- 5 committee members chosen by VPR
- 3 consecutive days, 1-5pm, if virtual
- Professionally produced virtual lab tour
- 3-page templated description for each project
- 4 hours of technical research talks
- Separate meetings with students, post-docs, and research scientists
- Discuss the added value of a Center
- Discuss the translational accomplishments of the Center

Order:

- KCC
- CCRC
- KSC
- WDRC
- AMPM
- VCC
- ECRC
- ...

TTO Briefing, 12 October

KAUST's Technology Transfer Office will take 20 minutes on 12 October to make a short presentation and lead a discussion on how they help in the research journey with:

- Patents
- Intellectual Property
- Commercialization of the research
- Developing and utilizing industry contacts

ECRC Culture

For discussion on
12 October, with this
10-page document
circulated in advance

Extreme Computing Research Center: our culture¹

This document describes the culture we seek for the Extreme Computing Research Center (ECRC) at KAUST. We have not achieved it in full, but we articulate it as a shared pursuit. Describing the “culture” of a research group should include discussion of: What are our key goals and characteristics? What sets us apart and makes us worthy of support? How can we achieve a good work atmosphere? We provide some general principles in the KAUST context to make this tangible.

1. The KAUST Mission
2. The KAUST Model
3. ECRC Research Culture and Best Practices

1. The KAUST Mission

The first mission of KAUST is to revive the ideals of the *Bayt al Hikmah* (“House of Wisdom”) of the ninth century pan-Islamic world in the context of the contemporary world of research. The *Bayt al Hikmah* was the leading repository and generator of scientific discovery and technological innovation of its period, worldwide. It was especially important to the development of algebra. The word “algorithm” comes from Al Khwarizmi, who worked there, and algorithms are key deliverables for those of us who work in the ECRC today, who thus lay claim to an inspiring heritage! We strive to be globally influential, through research that is ranked near the top of the international scale. Without quality, nothing else matters. Volume without quality certainly does not matter; it accrues only to KAUST’s “denominator.” The ECRC must strive so that its every activity or investment also contributes to the “numerator” in a significant way.

The second mission of KAUST is to contribute to a transformation of Saudi Arabia’s economy from one heavily dependent on natural resources to one driven by human resources and powered by innovation: a “service economy” or a “knowledge economy.” Our goal is to translate *from discovery to delivery*. This second mission embraces many aspects, from supplying a creative and skilled scientific and engineering workforce for national academia, industry, and agencies; to fostering the creation of new companies; to developing and licensing technology and software. KAUST research need not be limited to projects for which there is a clear and immediate pathway to economic exploitation, but the preponderance of the investment should contribute to environmentally and economically sustainable advance – for the Kingdom, the region, and the world. The ECRC contributes to this mission primarily by providing core capabilities. The quality of our research is measured by comparison within its discipline; the relevance of our research is measured, in large part, by how it enables the advance of *other* disciplines that

¹ This document substantially follows the “Simula Culture” document at <https://www.simula.no/about/simula-culture>. The imitation is without apology since the Simula document is so well thought out, by such good friends.

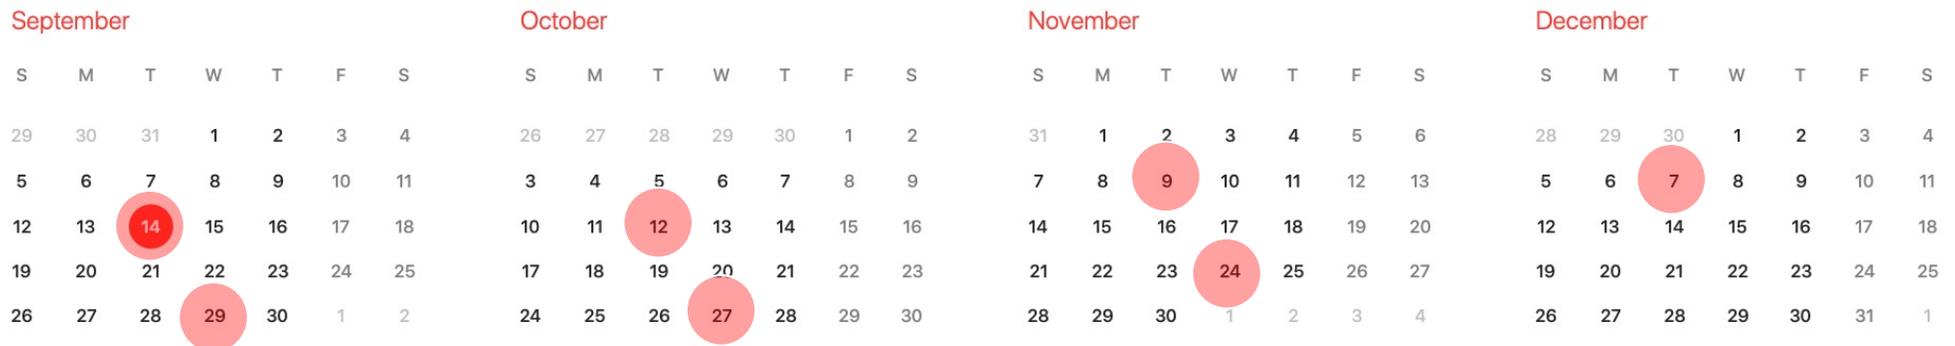


Other announcements?
Other conference notices?
Other training coming up?

Calendar and content

ECRC Communication Meeting planned for every other week:

- Sep 14 (Tue, lunch)
- Sep 29 (Wed)
- Oct 12 (Tue, lunch)
- Oct 27 (Wed)
- Nov 9 (Tue, lunch)
- Nov 24 (Wed)
- Dec 7 (Tue, lunch)



The second Tuesday of every month will be a luncheon meeting 12:15-1:45pm and the fourth Wednesday of every month we will meet 1:15-2:45pm. This schedule avoids conflicts with CS faculty meetings on the fourth Tuesday of each month and it avoids teaching conflicts with all Center faculty affiliates.

Calendar and content

Is this combination of people, announcements, and technical overview useful?

- We don't need another seminar series!
- What is missing that ECRC can supply?
- Send ideas to david.keyes@kaust.edu.sa

Center-based funding

- Center baseline funds
 - support people on the organizational chart (administrators and research scientists)
- Center competitive funds (CCF) also called Center Applied Research Funds (CARF)
 - support projects proposed every April for Jul 1-Jun 30 fiscal year execution
- **Center partnership funds (CPF)**
 - **For 2021-2022, these support translational collaborations with in-Kingdom ministries and industries – deadline 1 November**

ECRC project funding updates

CCF projects

- ExaGeoStat (Marc/Sameh)
- Hierarchical matrix technology at scale (George)
- High performance CFD (Matteo P)
- Optimization in machine learning (Peter)
- Exploring emerging architectures (Suhaib)
- Data compression and prediction (Panos)
- Hybrid discretizations for wave propagation (David/Longfei)
- Improving RTM-based petroleum exploration (David/Long)

Translational project

- Hardware-Software Solutions to the Challenges of Next-Gen Wireless Communication Networks (Adel)

AI initiative project

- Hierarchical Matrix Approximations of Hessians for Machine Learning Applications (George)

ExaGeoStat – Sameh Abdullah

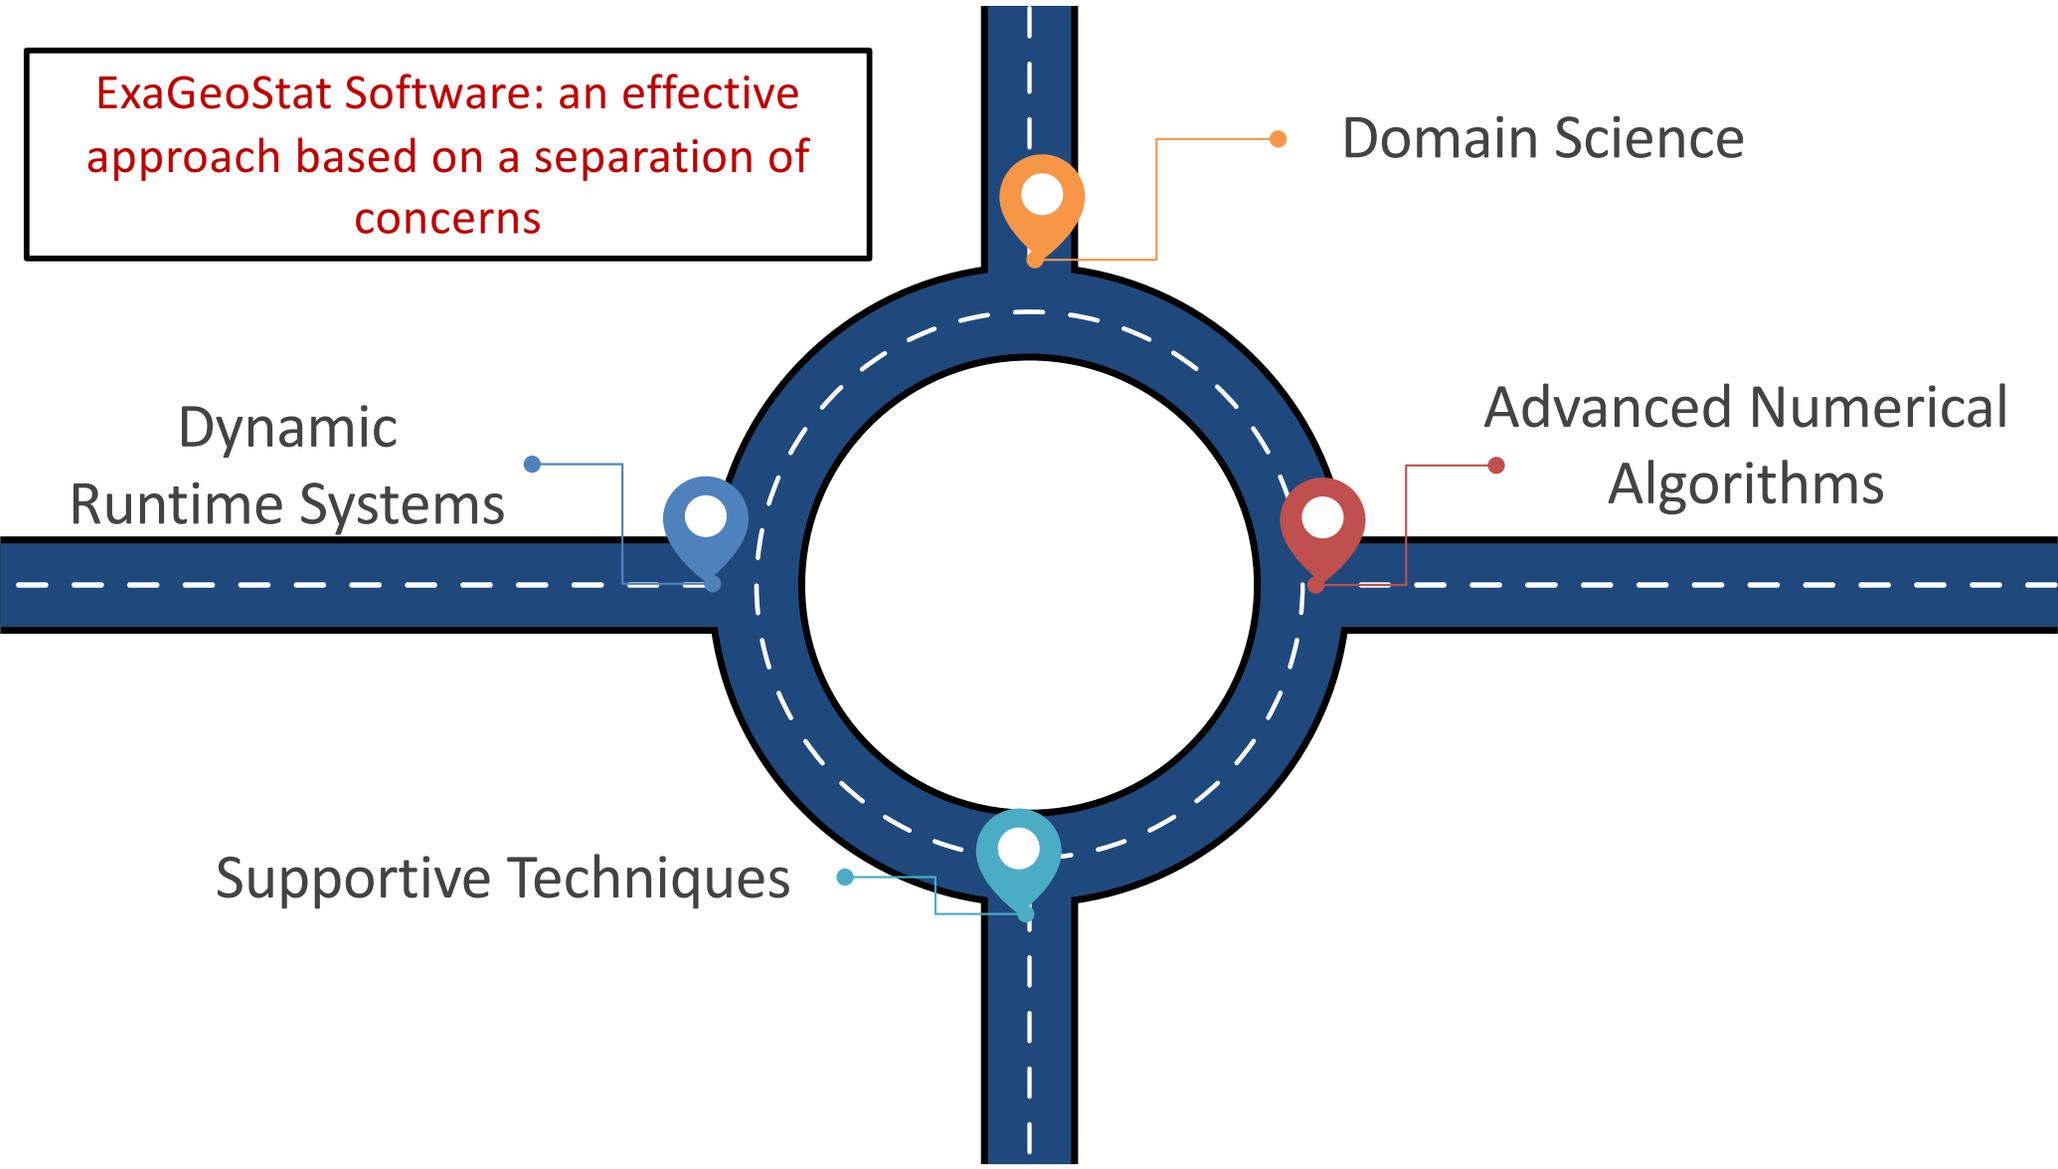
ExaGeoStat Software: an effective approach based on a separation of concerns

Domain Science

Dynamic Runtime Systems

Advanced Numerical Algorithms

Supportive Techniques



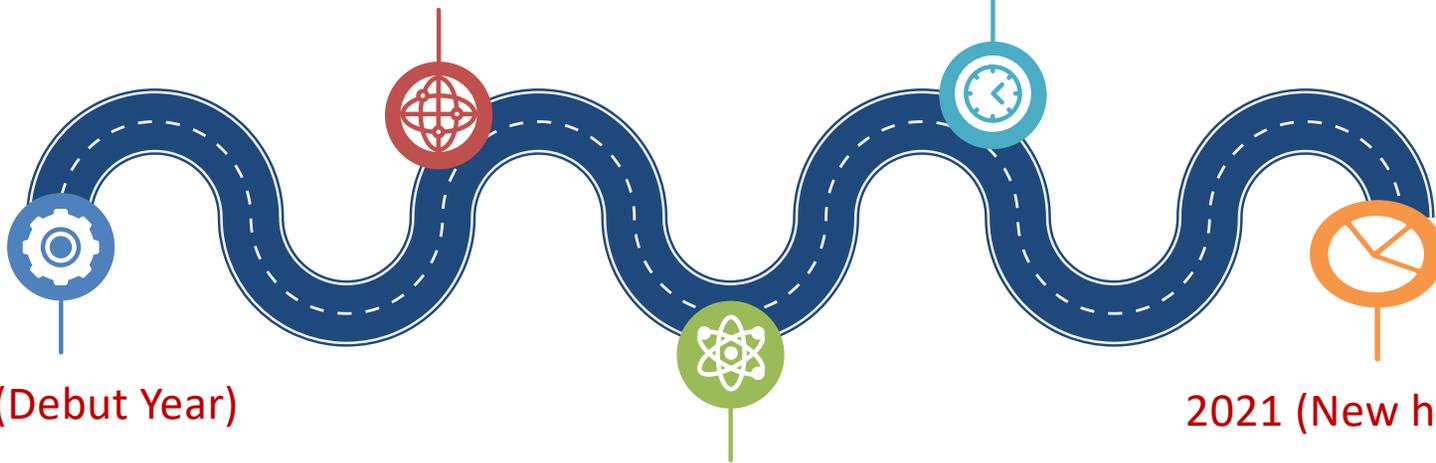
ExaGeoStat Development Timeline

2018 (TLR Year)

ExaGeoStat V1.0.0/ ExaGeoStat V1.0.0 were released with TLR support through KAUST HiCMA library.

2020 (Covid19 Year)

- ExaGeoStat V1.1.0/ ExaGeoStat V1.0.1 were released.
- UTK PaRSEC runtime system integration.
- Three-level mixed-precision ExaGeoStat.



2017 (Debut Year)

ExaGeoStat V0.1.0/ExaGeoStatR V0.1.0 were released with large-scale spatial data generator, exact modeling, and prediction capabilities.

2019 (MP Year)

Two-level mixed-precision ExaGeoStat with high accuracy achievements.

2021 (New heights)

- KAUST Spatial Statistics Competition.
- KAUST-NVIDIA Workshop - Geospatial data science competition.
 - Space-Time modeling support.
- Parallel optimization with PSwarm algorithm.
- Non-Gaussian Modeling (Exact/TLR/HODLR).

ExaGeoStat Ongoing/Future Roadmap

TLR/MP Modeling

Develop a new approximation method that combines both TLR and MP approximation techniques.

Spatial Statistics Benchmarking Framework

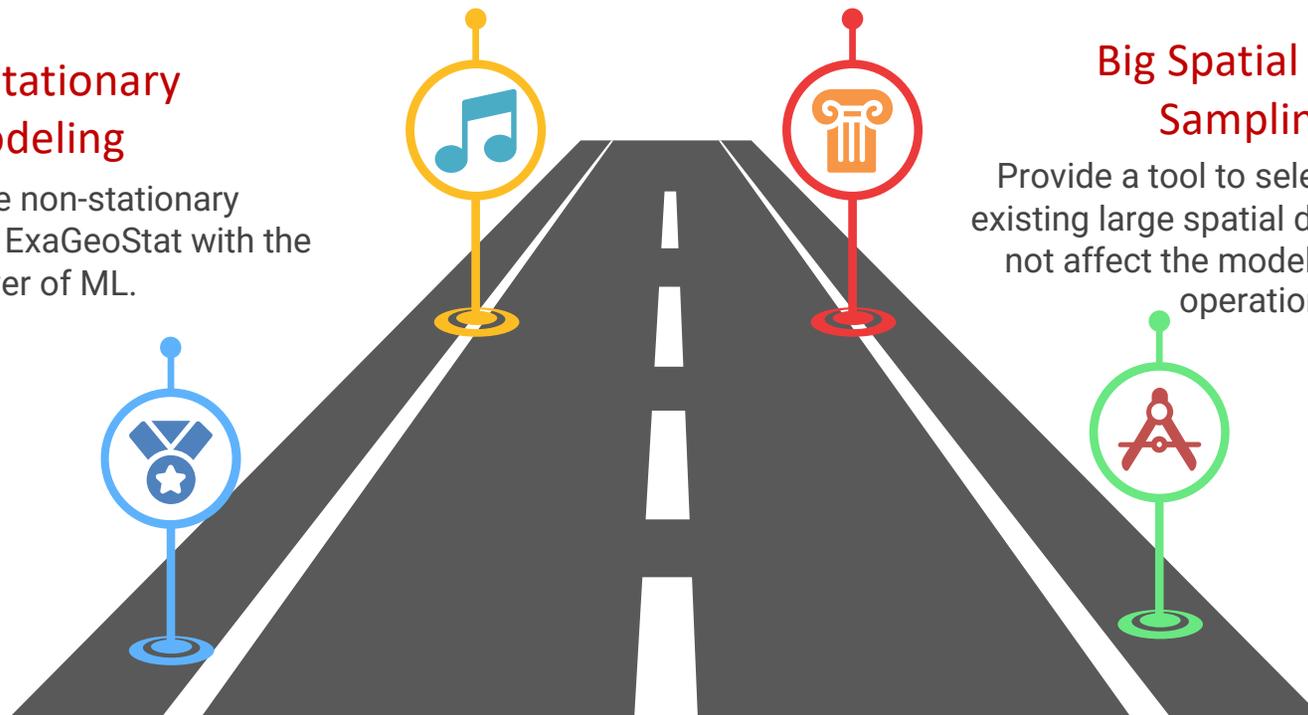
Provide a general framework to assess existing spatial statistics modeling/prediction algorithms with several metrics.

Non-Stationary Modeling

Large-scale non-stationary modeling using ExaGeoStat with the power of ML.

Big Spatial Data Sampling

Provide a tool to select a sample of existing large spatial dataset that does not affect the modeling/prediction operations.



ExaWave2 – Long Qu

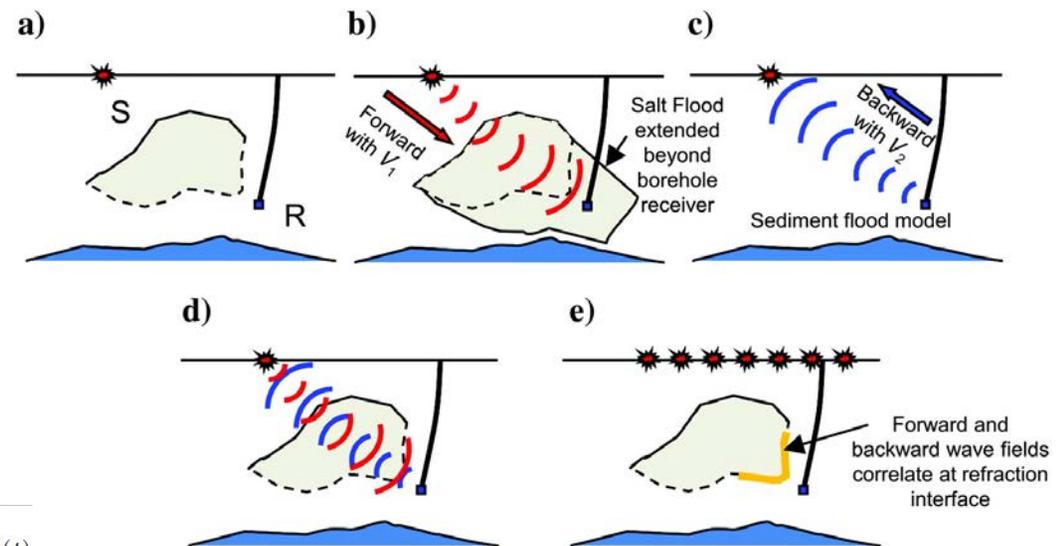
Overview on 3D Reverse Time Migration

- a state-of-the-art seismic migration technique used in the O&G industry for accurate imaging of subsurface with great structural and velocity complexities.

- based on 3D acoustic wave equation.

$$\frac{1}{V_P^2(\mathbf{x})} \frac{\partial^2 P(\mathbf{x}, t)}{\partial t^2} = \frac{\partial^2 P(\mathbf{x}, t)}{\partial x^2} + \frac{\partial^2 P(\mathbf{x}, t)}{\partial y^2} + \frac{\partial^2 P(\mathbf{x}, t)}{\partial z^2} + \delta(\mathbf{x} - \mathbf{x}_s) s(t)$$

$$\frac{P_{i,j,k}^{n+1} - 2P_{i,j,k}^n + P_{i,j,k}^{n-1}}{\Delta t^2} = V_{i,j,k}^2 [O_{ii}^8(P_{i,j,k}^n) + O_{jj}^8(P_{i,j,k}^n) + O_{kk}^8(P_{i,j,k}^n)],$$



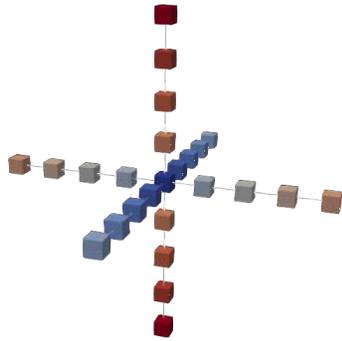
Credential : Reverse time migration of transmitted wavefields for salt boundary imaging, C. Willacy, M. Kryvohuz, Geophysics V84 Issue 2

Top500 list [2021/06] Oil and Gas industry

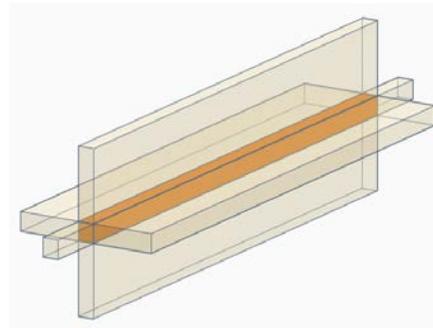
Rank		Rmax
9	<u>HPC5 - PowerEdge C4140, Xeon Gold 6252 24C 2.1GHz, NVIDIA Tesla V100, Mellanox HDR Infiniband, Dell EMC</u> <u>Eni S.p.A.</u>	35.4 Pflops
11	<u>Dammam-7 - Cray CS-Storm, Xeon Gold 6248 20C 2.5GHz, NVIDIA Tesla V100 SXM2, InfiniBand HDR 100, HPE</u> <u>Saudi Aramco</u>	22.4 Pflops
19	<u>Ghawar-1 - HPE Cray EX, AMD EPYC 7702 64C 2GHz, Slingshot-10, HPE</u> <u>Saudi Aramco</u>	19.25 Pflops
21	<u>PANGEA III - IBM Power System AC922, IBM POWER9 18C 3.45GHz, Dual-rail Mellanox EDR Infiniband, NVIDIA Volta GV100, IBM</u> <u>Total Exploration Production</u>	17.86 Pflops
.....	Unpublished machine from other energy majors : BP, Exxon Mobile, Chevron, etc.	

Stencil computation and Spatial Blocking

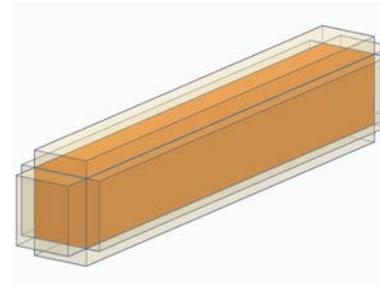
- Spatial blocking (X,Y,Z)
(classic method)



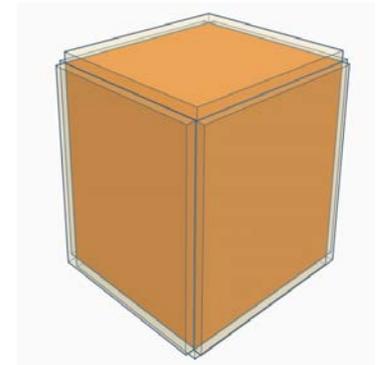
3D finite-difference stencil
with 8th order in space
with constant coefficients



1D



2D

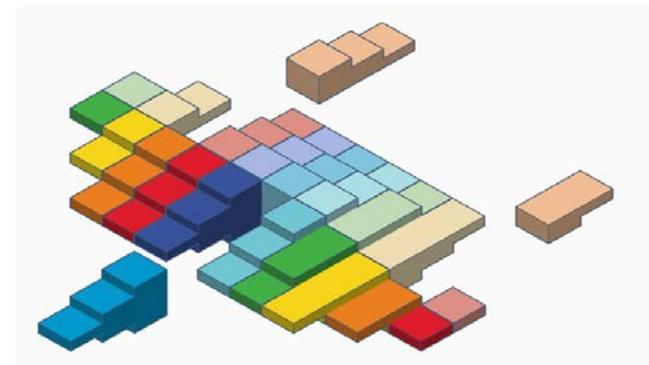


3D

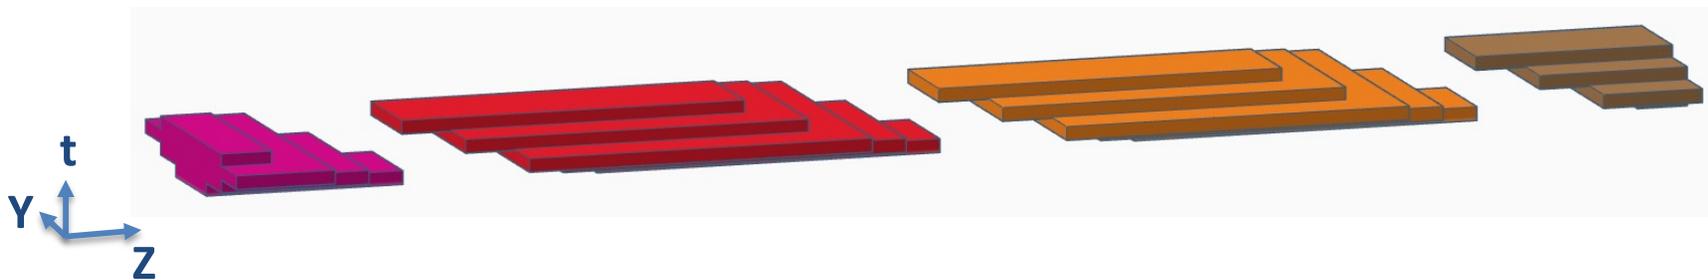
Illustration of Spatial Blocking applied to the stencil computation

Multi-Wavefront Diamond tiling Temporal Blocking

- Temporal blocking (X,Y,Z,t)
(new method)
- Combining diamond tiling for temporal blocking and wavefront parallelism techniques to enhance cache reuse
- GIRIH: <https://github.com/ecrc/girih>

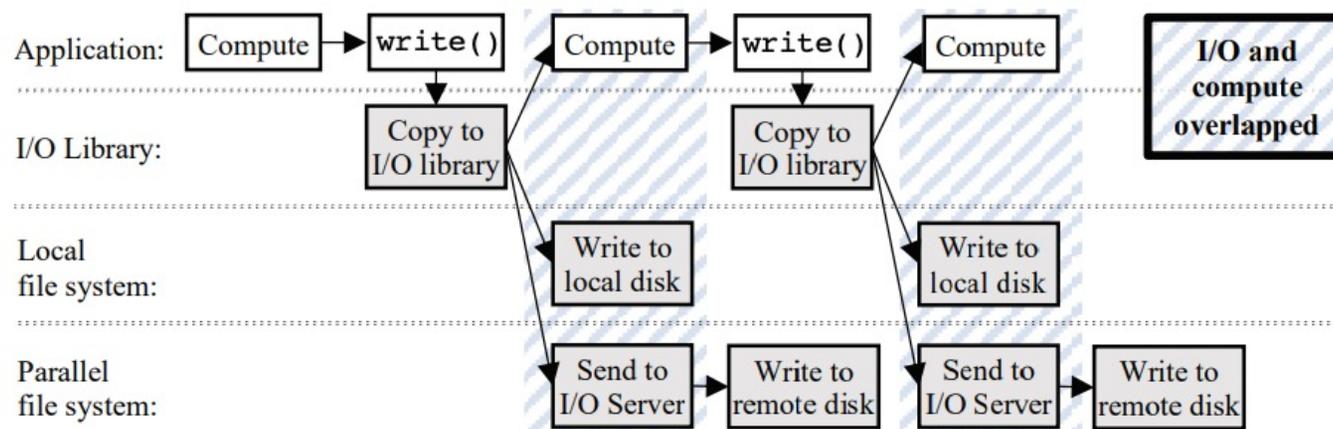


Well-balanced workload
among threads



Multi-Layered Buffer Storage (MLBS)

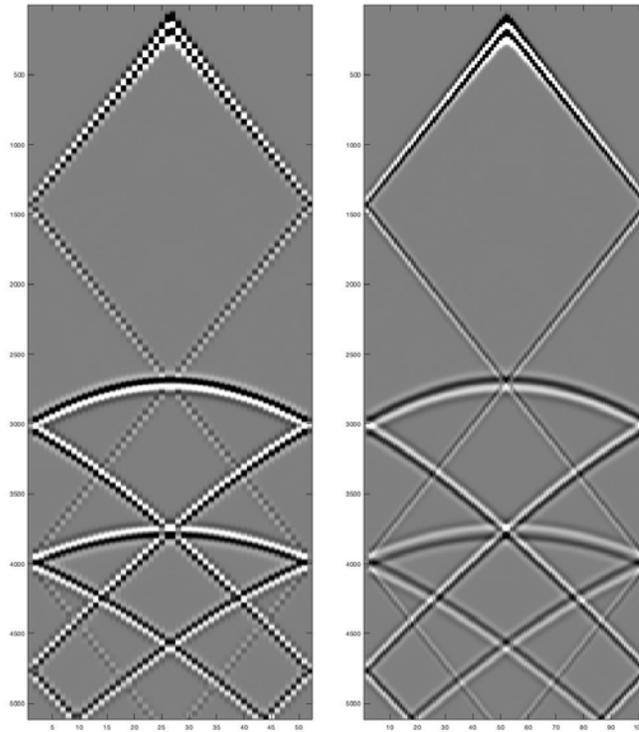
MLBS : **caching and prefetching data** in out-of-core scientific applications to perform expensive I/O operations asynchronously.



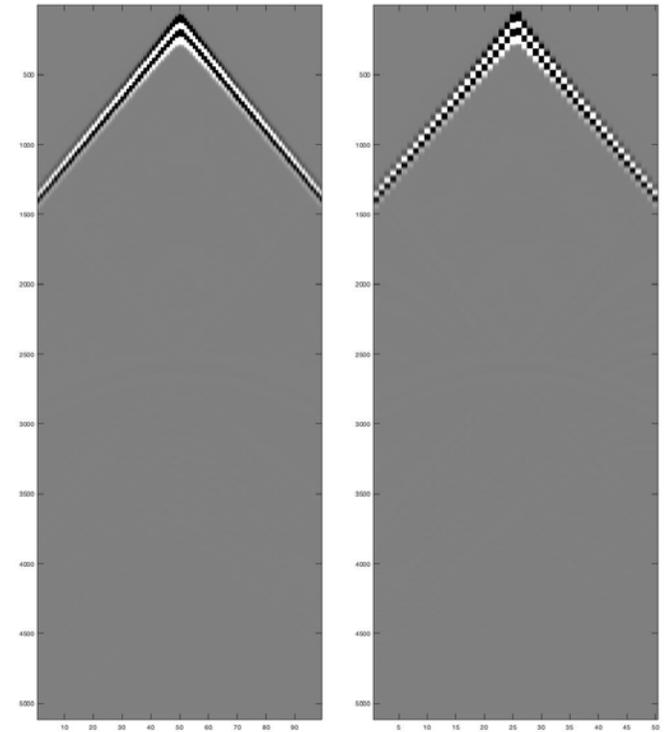
Reference : *MLBS: Transparent Data Caching in Hierarchical Storage for Out-of-Core HPC Applications*, T. Alturkestani, T. Tonellot, H. Ltaief, R. Abdelkhalak, V. Etienne & D. Keyes, *HiPC*, 2019

Modeling using $V_p = 4000$ m/s

- Grid size :
520 x 520 x 520
- Timesteps :
5120
- CPML:
10 pixels
- Receivers :
50 x 99

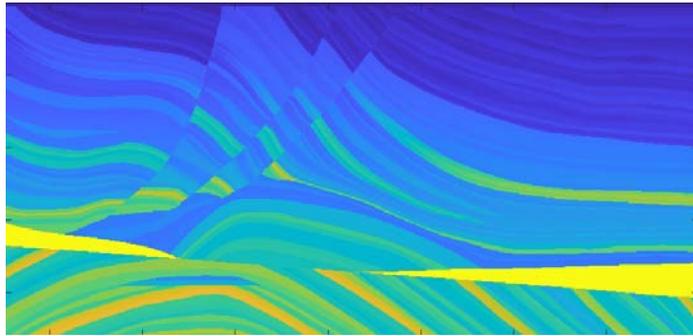


Modeling results without CPML
(Inline, Xline)

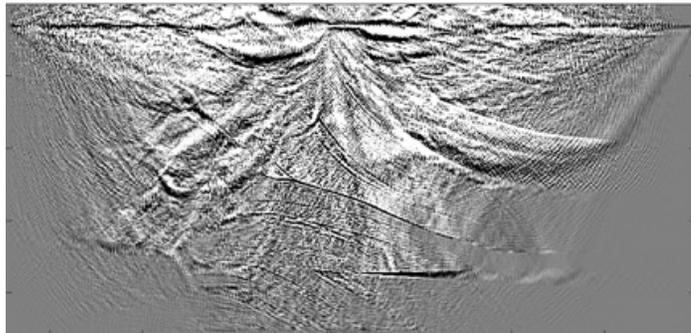


Modeling results with CPML
(Inline, Xline)

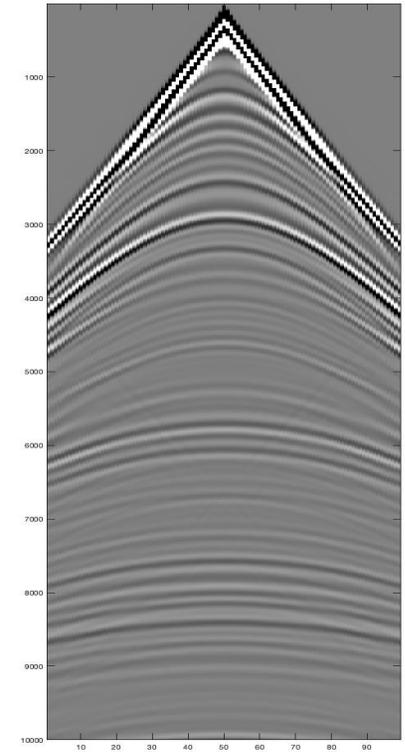
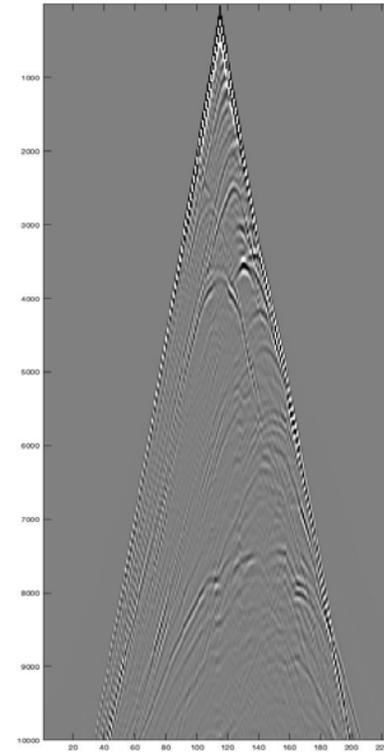
Modeling and Migration using 3D Marmousi dataset



Velocity Model



Migrated shot



Modeling results (Inline, Xline)

HPC-Connect – Adel Dabah

HPC-Connect:

Proposing an HPC Hardware-Software Solution to the
latency and accuracy issues in Next-Generation
Wireless Communication Networks.

Adel Dabah : adel.dabah.1@kaust.edu.sa

Prof. David Keyes (PI), Prof. Slim Alouini (co-PI), and Dr. Hatem Ltaief (co-PI)
KAUST Research Translation fund.

Technologies that make use of next-generation wireless communication networks.

Fields:

- Mobile communication
- Users with various terminals
- Smart buildings / smart cities
- IOTs
- UAVs
- Connected/Autonomous vehicles
- Augmented reality / Games
-

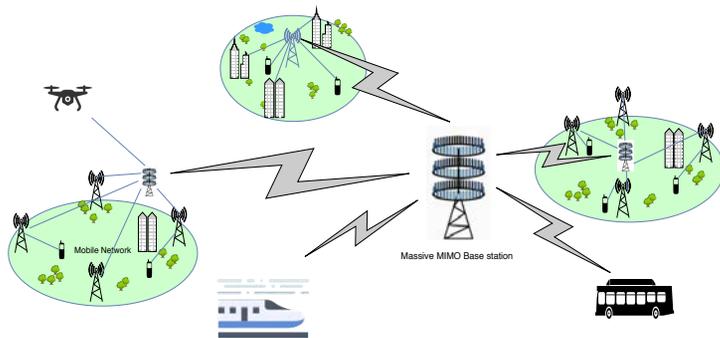
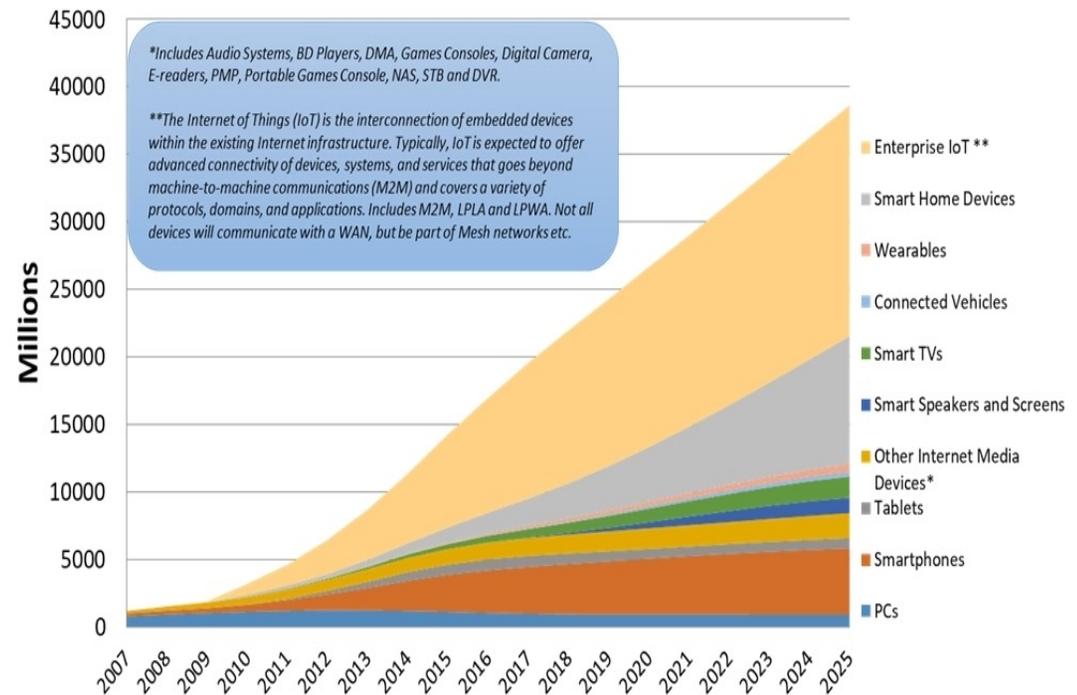


Figure 1: Massive MIMO for future wireless communication systems.



Source – Strategy Analytics research services, May 2019: IoT Strategies, Connected Home Devices, Connected Computing Devices, Wireless Smartphone Strategies, Wearable Device Ecosystem, Smart Home Strategies

Massive MIMO as a promising solution

- MIMO technology marked the transition from 3G to 4G wireless communication network.

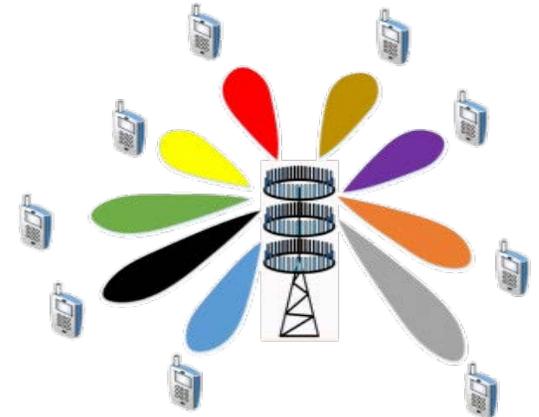
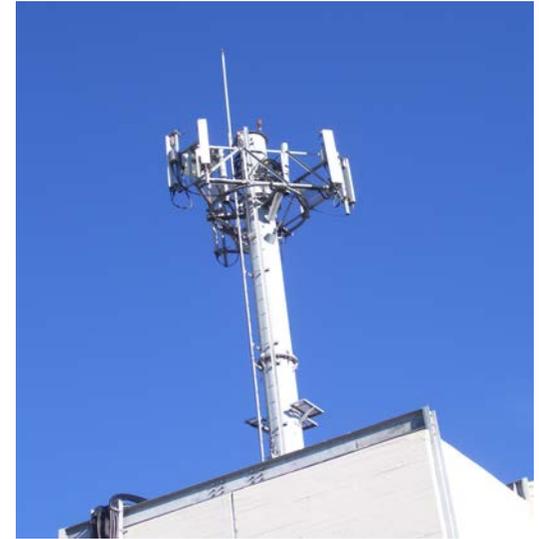
MIMO Benefits:

- Increasing the data-rate
 - Increasing the diversity gain
 - Increasing the area covered by the base station
- Massive MIMO: scale up the number of antennas up to several hundreds.
 - High data-rate
 - High accuracy
 - Good network reliability
 - Reduce interference by focusing on a specific user (beamforming)

Issues:

- High communication latency
- Accuracy of the detection
- The base station's power budget

Multiple-antenna base station



beamforming

When HPC meets Wireless:

HPC Hardware-Software Solutions for Massive MIMO systems.

Massive MIMO detection algorithms:

- Linear
- Nonlinear
 - Modeling the problem as search-tree
 - Complexity increases exponentially with number of antennas

Our Solution:

- GPUs: they deliver maximum performance per watt (flops/watt).
- New algorithms based on Trade-offs between accuracy and complexity.
 - Increase the arithmetic intensity by transforming the problem to be more compute bound.
 - Using Half-precision computation to reduce the latency.

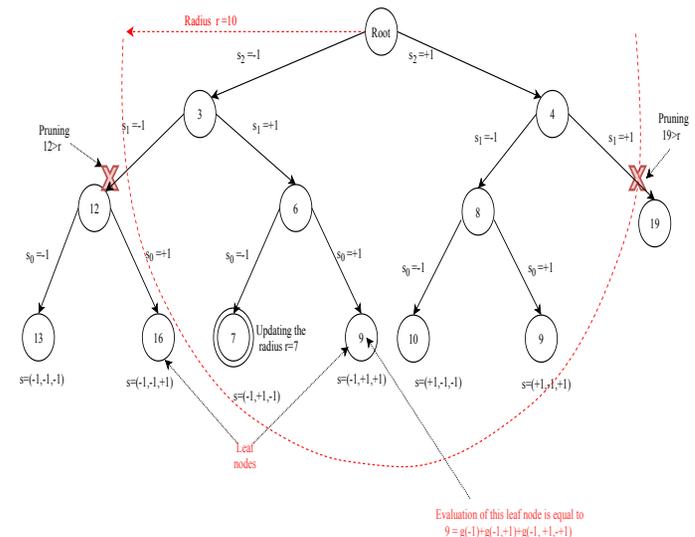


Figure 2: This figure represents an example of the SD search tree for a MIMO system with three transmit antennas. At each level, one symbol is fixed, starting with the last one. The partial evaluation of each node (sub-problem) is stored inside the circles. The pruning process uses the evaluation of the node and the sphere radius r to avoid unpromising branches.

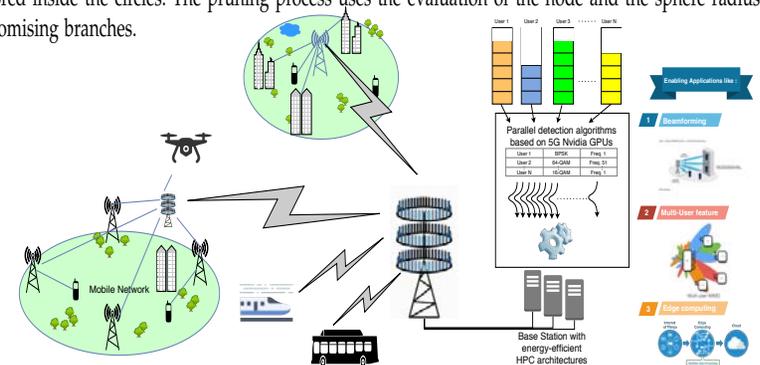


Figure 5: Global scheme of our framework which will be deployed in M-MIMO base station equipped with latest energy-efficient GPU hardware accelerators.

Extreme Compression – Turkiyyah

Majlis

