Welcome to IXPUG Middle East 2018

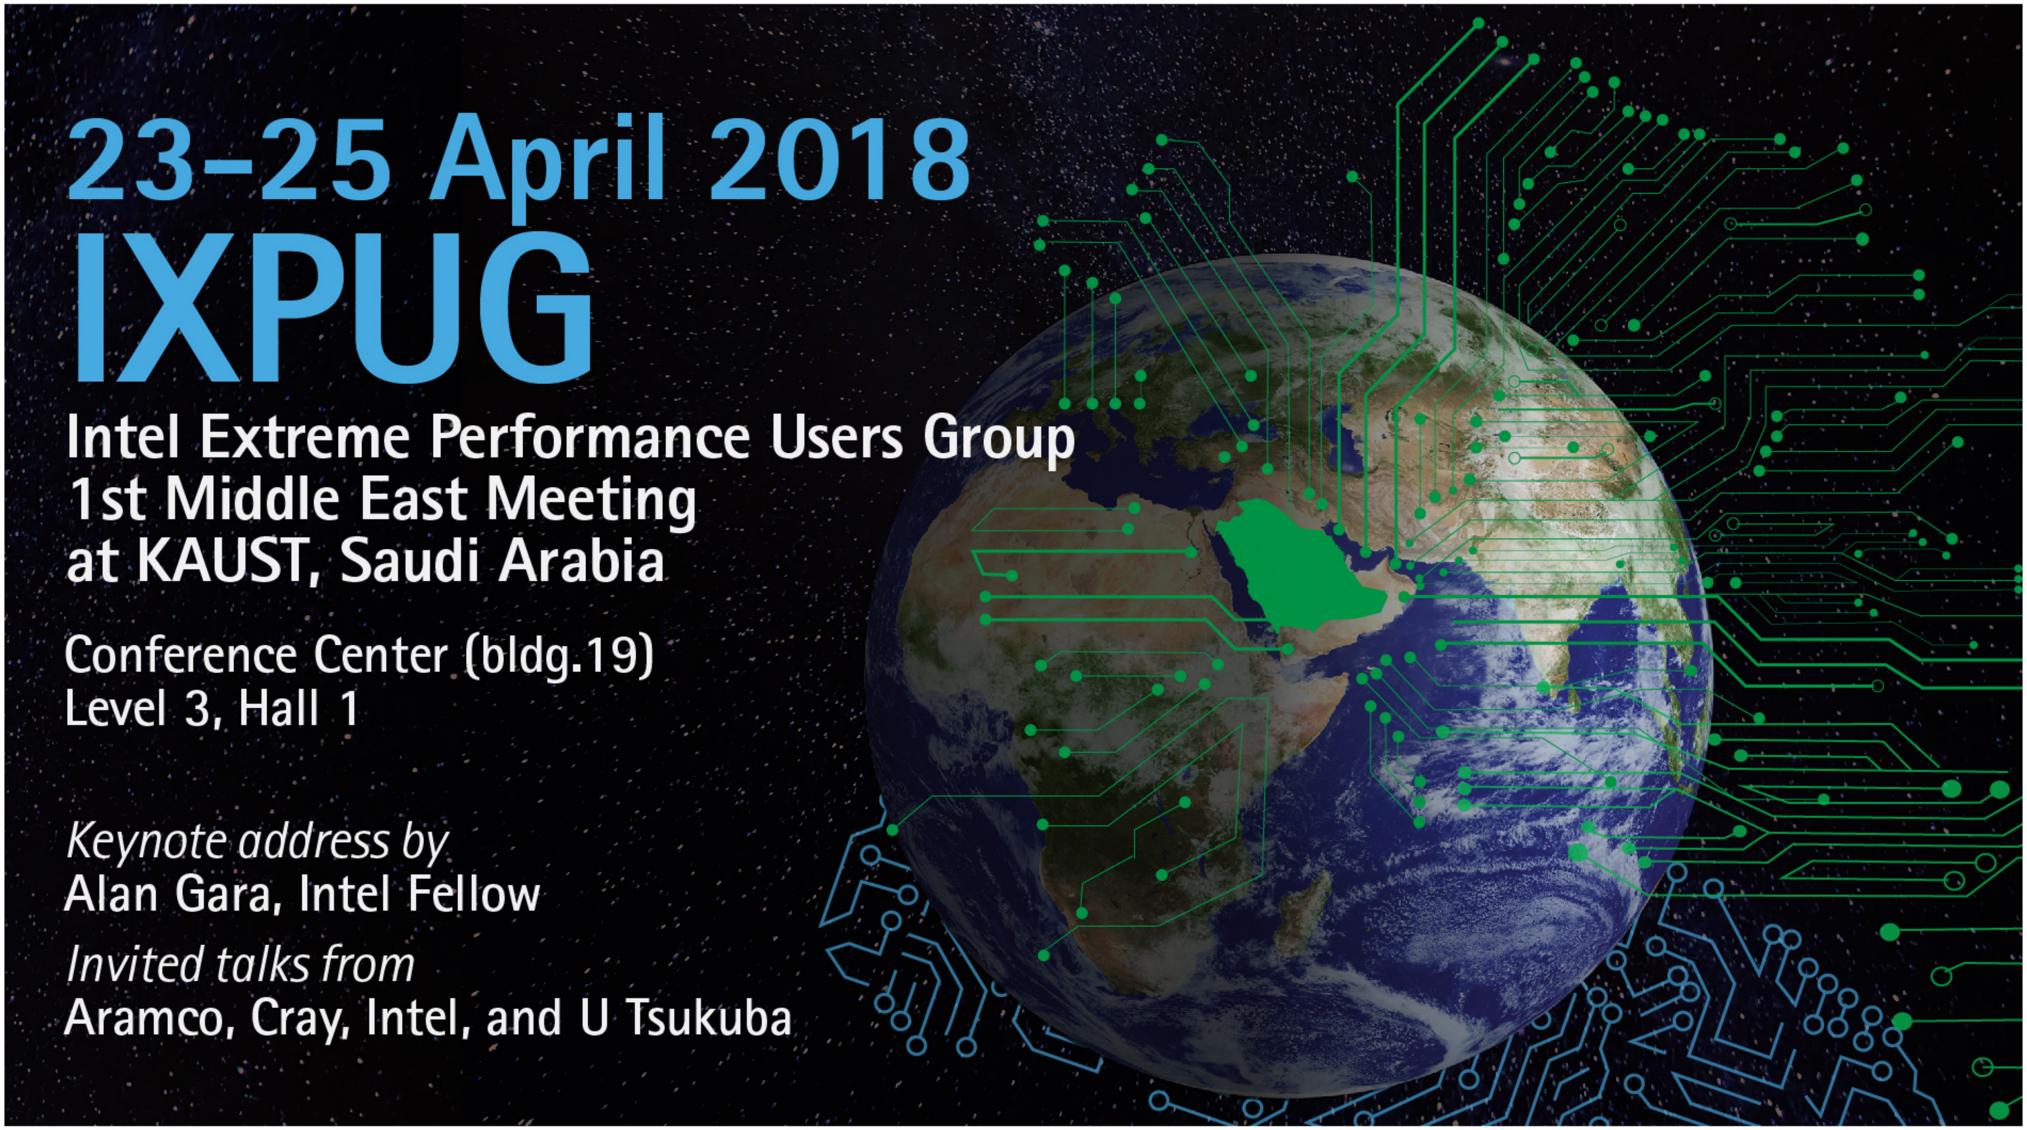

Haiku for IXPUG meetings

Held a·round the world,

IX·PUG ral·lies re·search·ers:

Man·y·core is here!

Where are we … on campus?

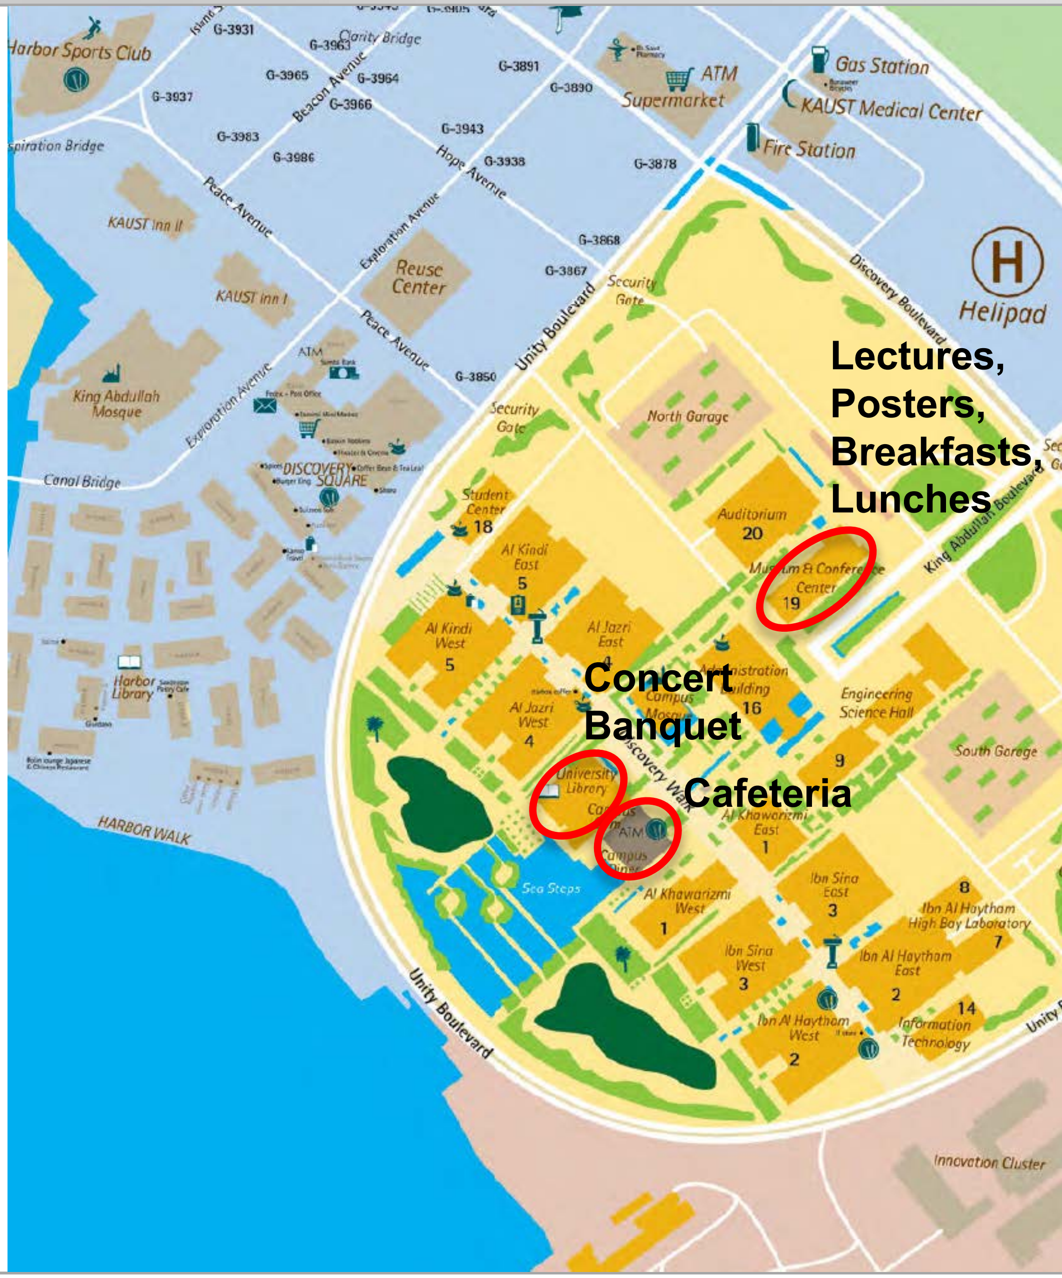

Where are we … in the conference center?

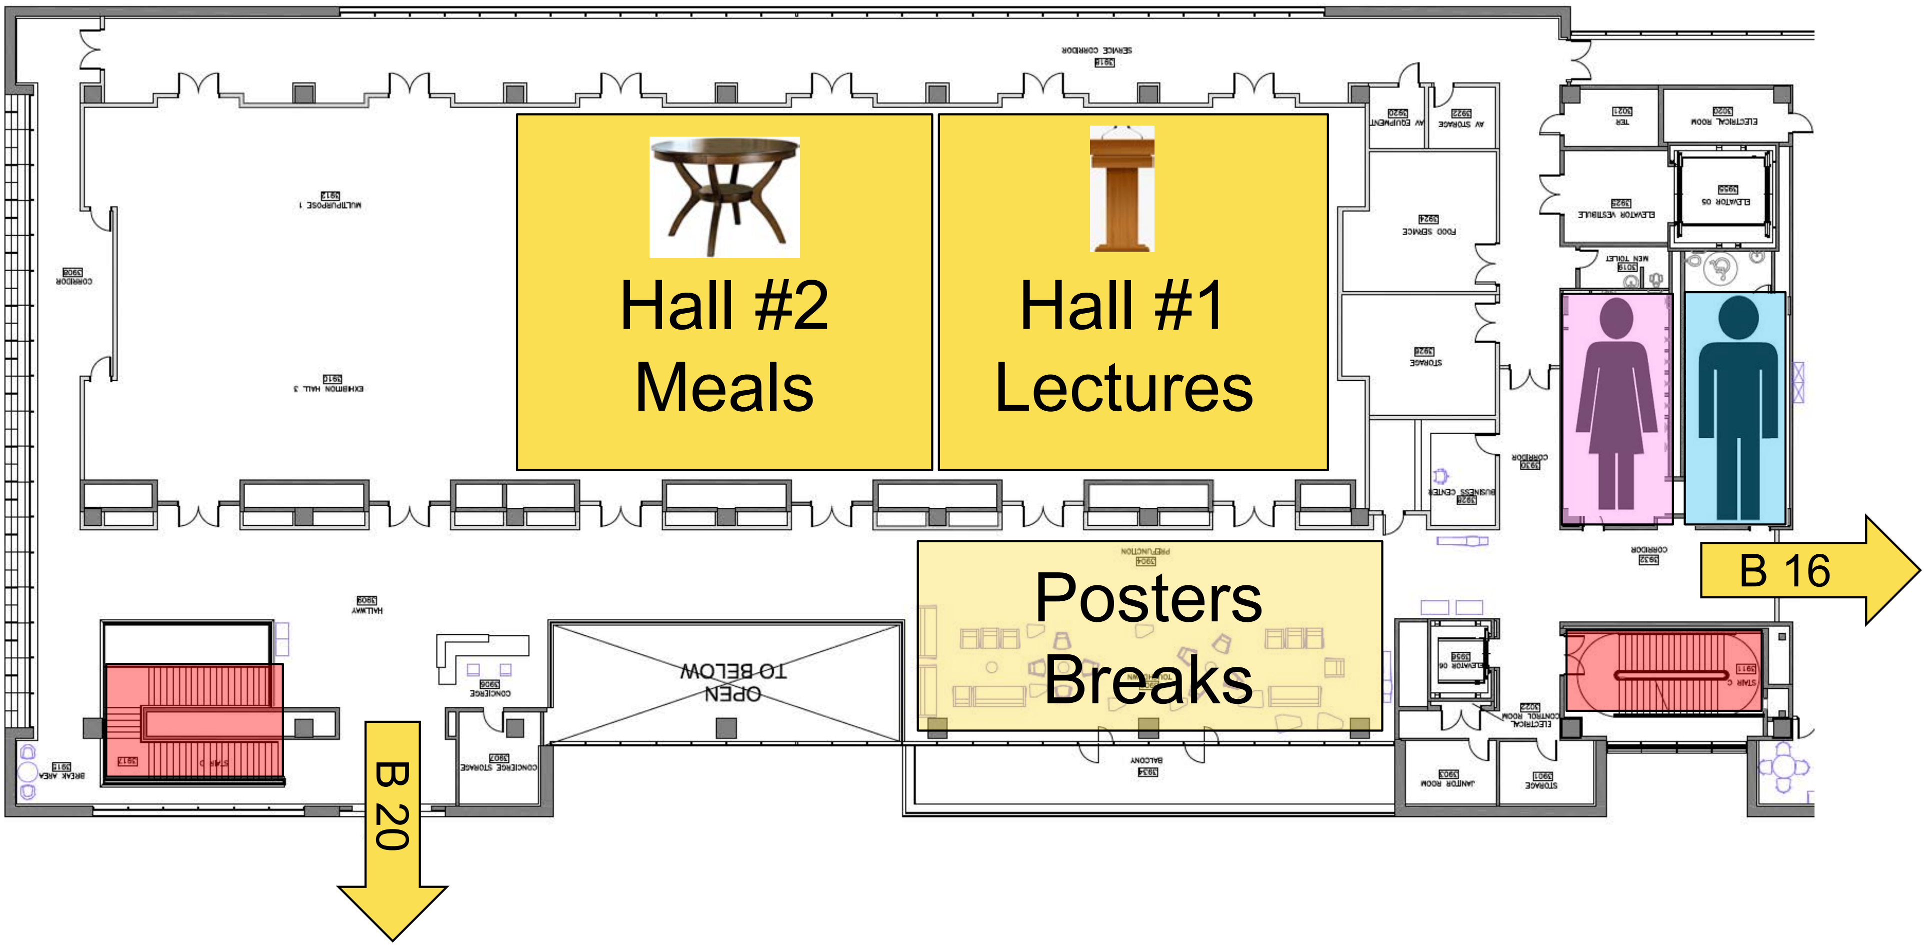

KAUST Conference Center
Building 19

What's across the bridge?

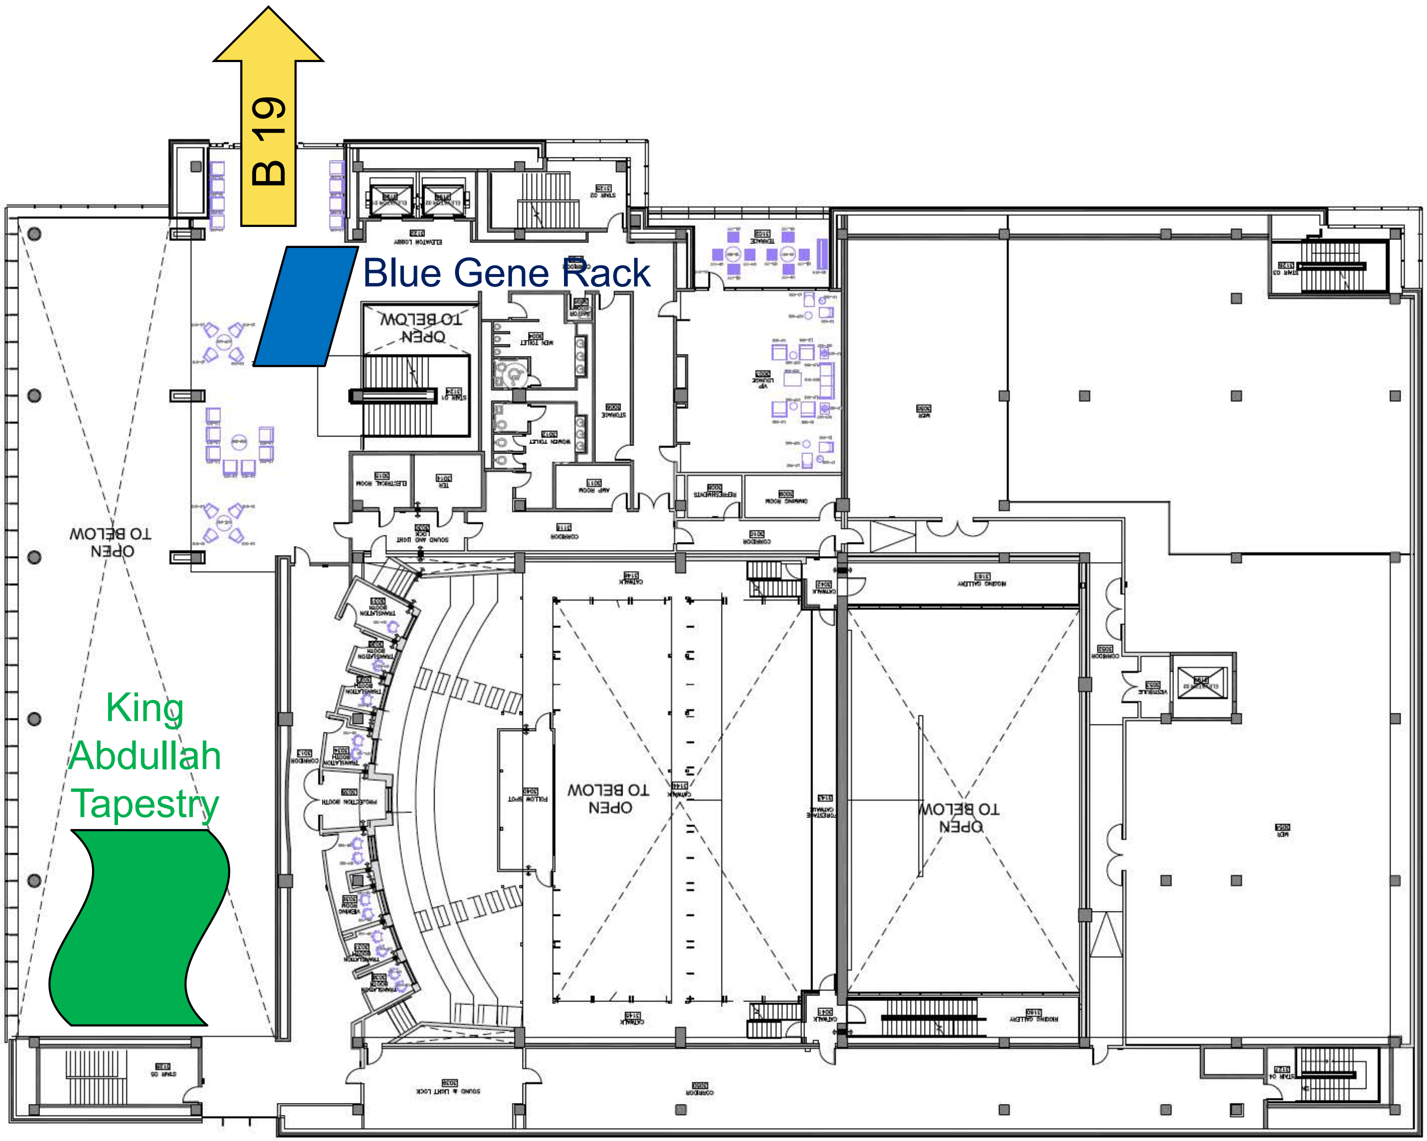

KAUST Auditorium Building 20

Al Gara's previous system

2008
IBM Blue Gene

King Abdullah tapestry

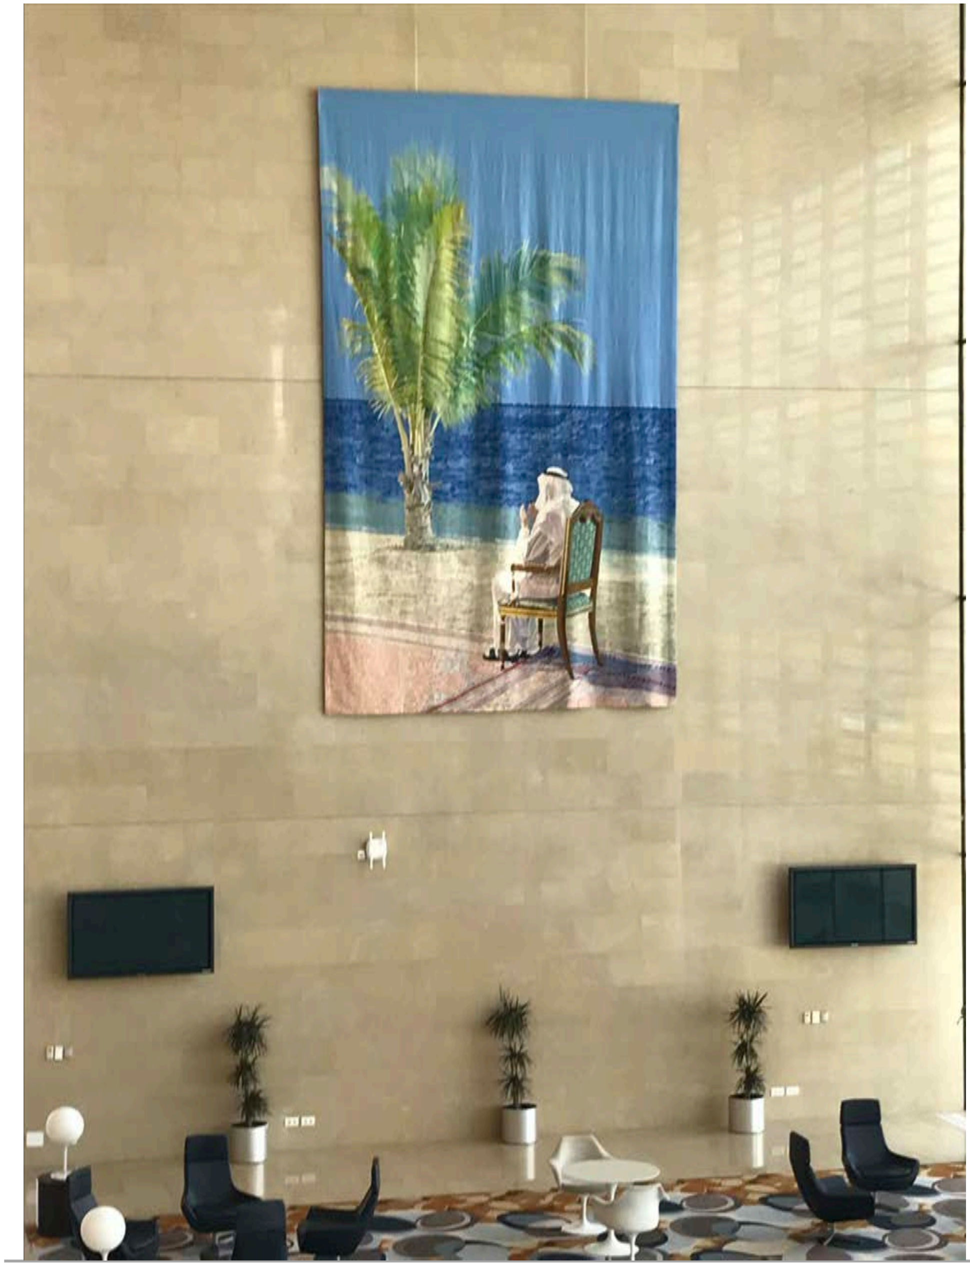

1924-2015
Abdullah bin Abdulaziz al Saud

Logistical questions?

see just outside:

Sherif, Frances, or Heba



Our IXPUG feast

Your program is around your neck ☺

…on your USB key

Agenda, 23-25 April

	23-Apr	24-Apr	25-Apr
	Monday	**Tuesday**	**Wednesday**
07:30	**Buffet Breakfast**	**Buffet Breakfast**	**Buffet Breakfast**
08:30	**Opening remarks**	**video: KAUST's Extreme Computing**	**videos: KAUST and Saudi Arabia**
09:00	KEYNOTE: Al Gara	INVITED TALK #3: Adrian Tate	TUTORIALS &
09:30		INVITED TALK #4: Taisuke Boku	INDUSTRIAL OUTREACH
10:00	INVITED TALK #1: Alex Heinecke	CONTRIB. TALK #5: Mohamed Ibrahim	[in parallel]
10:30	**Coffee & Networking**	**POSTERS**	**Coffee & Networking**
11:00	CONTRIB. TALK #1: Jahanzeb Hashmi	CONTRIB. TALK #6: Rabab Alomairy	TUTORIALS &
11:30	CONTRIB. TALK #2: Hatem Ltaief	CONTRIB. TALK #7: Mustafa AbdulJabbar	INDUSTRIAL OUTREACH
Noon	**Lunch**	**Lunch**	**Lunch**
13:30	INVITED TALK #2: Jeff Hammond	INVITED TALK #5: Vincent Etienne	
14:00	CONTRIB. TALK #3: Kadir Akbudak	CONTRIB. TALK #8: Matteo Parsani	
14:30	CONTRIB. TALK #4: Amani Alonazi	CONTRIB. TALK #9: Mohammed Farhan	
15:00	**Coffee & Networking**	**Coffee & Networking**	
15:30	LIGHTNING TALKS	SITE UPDATES	
16:00			
16:30	**POSTERS**	PANEL	
17:00			
17:30	**Evening Musicale (Library)**		
18:15			

HPC in Industry, Agencies & Business

Industrial Outreach

David MARTIN

High Performance Computing in Industry, Agencies, and Small Business

High performance computing (HPC) in industry includes simulation and data analytics. Today, "HPC" means more than the speed of numerical processing. It refers as well to the vast storage resources and high-bandwidth communication that come proportionally with fast processing resources. This program on the impact of HPC in industry will contain with a keynote, a series of short contributed talks, and a closing panel. The workshop is intended to be interactive and allow time for discussion and sharing experiences. Industries that are established in HPC are encouraged to present a short case study of a foray into HPC – either positive or negative. Industries that are only exploring the potential of HPC are welcome to raise questions. Topics to be addressed include:

- How industry uses HPC to increase competitiveness
- Examples of HPC driving innovation and improving product cycles
- Models for how academia and national laboratories collaborate through computing
- Strategies for increasing pre-competitive collaboration among like industries
- Strategies for increasing collaboration between complementary industries
- Ideas about how exascale computing will impact industry

Biosketch

David Martin is Manager, Industry Partnerships and Outreach at the Argonne Leadership Computing Facility at Argonne National Laboratory, where he works with industrial users to harness high performance computing and take advantage of the transformational capabilities of modeling and simulation. David brings broad industry and research experience to ALCF. Prior to joining ALCF, David led IBM's integration of internet standards, grid and cloud computing into offerings from IBM's Systems and Technology Group. Before IBM, David managed networks and built network services for the worldwide high-energy physics community at Fermilab. David began his career at AT&T Bell Laboratories, doing paradigm-changing work in software engineering and high-speed networking. David has a BS from Purdue and an MS from the University of Illinois at Urbana-Champaign, both in Computer Science.

Industrial Outreach

42

Tutorial

George
MARKOMANOLIS

Saber
FEKI

The Burst Buffer: a New Frontier in Hierarchical Memories

While processor performance has been increasing significantly over the past several years, hard disks (I/O) performance has not improved to keep pace. In order to reduce this performance gap, the memory hierarchy includes nowadays additional layers, the latest being the usage of SSD-based storage as a Burst Buffer for I/O acceleration. The Cray solution is called DataWarp technology, and it is available in many HPC sites including the Intel-Haswell based Shaheen XC40 supercomputer operated by the KAUST Supercomputing Core Laboratory. Shaheen has a large installation of Burst Buffer using 568 Intel P3608 SSDs providing a total capacity of 1.5 PB and a throughput exceeding 1.5TB/s. In this tutorial, we present an efficient approach on how to use Burst Buffer technology to accelerate I/O for various applications, including weather forecast and seismic migration. We focus on optimizing massively parallel I/O on Burst Buffer, a relatively new problem compared to well-established optimizations for parallel I/O on disk-based parallel file systems. The tutorial will conclude with a demonstration of a complex workflow executed on the Burst Buffer, including simulation, analysis, and visualization.

Biosketch

George Markomanolis is a Computational Scientist in the KAUST Supercomputing Laboratory. His research is on performance optimization of applications with a focus on weather models. He is working on new technologies such as Burst Buffer to improve I/O performance and leads the KAUST Burst Buffer Early User program. Moreover, he is one of the developers of the IO-500 benchmark. He obtained his M.S. in Computational Science from the department of Informatics and Telecommunication, National and Kapodistrian University of Athens, Greece in 2008 and his Ph.D. in Computer Science from the Ecole Normale Supérieure de Lyon, France in 2014.

Saber Feki leads the Computational Scientist team at the KAUST Supercomputing Laboratory, providing scientific and application support and training to over 500 users of the Shaheen XC40 supercomputer. Saber received his M.Sc. and Ph.D. degrees in Computer Science from the University of Houston in 2008 and 2010, respectively. He then joined TOTAL in 2011 as HPC Research Scientist for about 2 years. His research interests include parallel I/O, parallel programming models, and automatic performance tuning.

44

Panel

Panel

Rashid MEHMOOD

Towards a "Saudi Top 10" List

Is the time ripe to establish and commit to curating a Top 10 (or Top 25, etc.) supercomputing list for Saudi Arabia? Such a list would give visibility to the importance of supercomputing as the Kingdom accelerates its economic transformation, make supercomputing easier to track, call attention to the growing HPC workforce, and lead better networking and the sharing of access and expertise. In 2002, China started its first national Top 50 list and sixteen years later, China has more supercomputers in the international Top500 than any other country, while the sum of its top two supercomputers outranks the sum of the top ten of every other country. Saudi Arabia could improve on the Top500 list by introducing a new best practice benchmark offering more insight than HPL, for instance one customized to energy-related simulations, data-intensive and smarter society applications. This panel will summarize and continue a vibrant discussion begun in March at HPC Saudi 2018, hosted by King Abdulaziz University, and seek to advance the concept and draw in new participants.

Biosketch

Rashid Mehmood is the Research Professor of Big Data Systems and the Director of Research, Training, and Consultancy at the High Performance Computing Centre, King Abdulaziz University, Saudi Arabia. He has gained qualifications and academic work experience from universities in the UK including Cambridge, Swansea, Birmingham, and Oxford. Rashid has over 20 years of research experience in computational modelling and simulation systems coupled with his expertise in high performance computing. His broad research aim is to develop multi-disciplinary science and technology to enable a better quality of life and smarter economy with a focus on real-time intelligence and dynamic system management. He has published over 150 research papers including 4 edited books. He has led and contributed to academia-industry collaborative projects funded by EPSRC, EU, UK regional funds, and Technology Strategy Board UK with the value over £50 million. He is a founding member of the Future Cities and Community Resilience (FCCR) Network. He is a member of ACM, OSA, Senior Member IEEE and former Vice-Chairman of IET Wales SW Network.
https://uk.linkedin.com/in/rashid-mehmood-87b841

Panel

Site Updates

Site Update #1

Richard GERBER

National Energy Research Scientific Computing Center (NERSC). Lawrence Berkeley National Laboratory, Berkeley, California, USA

The National Energy Research Scientific Computing Center (NERSC) is the primary scientific computing facility for the Office of Science in the U.S. Department of Energy. As one of the largest facilities in the world devoted to providing computational resources and expertise for basic scientific research, NERSC is a world leader in accelerating scientific discovery through computation. More than 6,000 scientists use NERSC to perform basic scientific research across a wide range of disciplines, including climate modeling, research into new materials, simulations of the early universe, analysis of data from high energy physics experiments, investigations of protein structure, and a host of other scientific endeavors. A collaboration between the Department of Energy's National Energy Research Scientific Computing Center (NERSC), Intel and five Intel Parallel Computing Centers (IPCCs) has resulted in a new Big Data Center (BDC) that will work both on code modernization and tackle real science challenges. The BDC is investigating how current HPC systems can support data-intensive workloads that require analysis of over 100 terabytes datasets on 100,000 cores or greater. The BDC will optimize and scale the production data analytics and management stack on NERSC's Cori system.

Biosketch

Richard Gerber is Senior Science Advisor and High Performance Department Head at the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory in Berkeley, CA. Richard has a Ph.D. in Physics (computational astrophysics) from the University of Illinois at Urbana-Champaign and held a National Research Council Postdoctoral Fellowship at NASA Ames Research Center before joining NERSC in 1996. He has more than 30 years' experience in high performance scientific computing and specializes in helping scientists productively use supercomputers.

Site Update

Site Updates

Site Update #2

David MARTIN

Argonne Leadership Computing Facility (ALCF). Argonne National Laboratory, Argonne, Illinois, USA

The Argonne Leadership Computing Facility (ALCF) provides supercomputing capabilities to the scientific and engineering community to advance fundamental discovery and understanding in a broad range of disciplines. Supported by the U.S. Department of Energy's (DOE) Office of Science, Advanced Scientific Computing Research (ASCR) program, the ALCF is one of two DOE Leadership Computing Facilities in the nation dedicated to open science. Through awards of supercomputing time and user support services, the ALCF enables large-scale modeling and simulation research aimed at solving some of the world's largest and most complex problems in science and engineering. The ALCF slated to host the USA's first exascale system when it is delivered in 2021: Intel-based Aurora will be capable of performing a quintillion calculations per second. The system is expected to have more than 50,000 nodes and more than 5 petabytes of total memory, including high-bandwidth memory.

Biosketch

David Martin is Manager, Industry Partnerships and Outreach at the Argonne Leadership Computing Facility at Argonne National Laboratory, where he works with industrial users to harness high performance computing and take advantage of the transformational capabilities of modeling and simulation. David brings broad industry and research experience to ALCF. Prior to joining ALCF, David led IBM's integration of internet standards, grid and cloud computing into offerings from IBM's Systems and Technology Group. Before IBM, David managed networks and built network services for the worldwide high-energy physics community at Fermilab. David began his career at AT&T Bell Laboratories, doing paradigm-changing work in software engineering and high-speed networking. David has a BS from Purdue and an MS from the University of Illinois at Urbana-Champaign, both in Computer Science.

Site Update #3

Taisuke BOKU

Joint Center for Advanced High Performance Computing (JCAHPC). University of Tsukuba, Tsukuba, Japan

The Joint Center for Advanced High Performance Computing was established in 2013 under agreement between the Center for Computational Sciences at the University of Tsukuba and the Information Technology Center at the University of Tokyo in order to design, operate and manage a next-generation supercomputer system. JCAHPC promotes advanced computational sciences, and contributes to the fields of academia, science, and technology. JCAHPC designs and develops large-scale HPC systems in which an operating system, programming languages, numerical libraries for many-core architectures are exploited. The Oakforest-PACS supercomputer ranked #1 in the first IO-500 list released in November 2017. The IO-500 is a world ranking list of storage performance, which is evaluated by a benchmark that measures the storage performance using read/write bandwidth for large files and read/write/listing performance for small files. Storage performance in supercomputers is critical for large-scale simulation, big data analysis, and artificial intelligence. The IO-500 list facilitates to improve the storage performance that greatly helps to improve the CPU performance.

Biosketch

Taisuke Boku received Masters and PhD degrees from Department of Electrical Engineering at Keio University. After his career as Assistant Professor in Department of Physics at Keio University, he joined the Center for Computational Sciences (formerly Center for Computational Physics) at University of Tsukuba where he is currently the Deputy Director, the HPC Division leader and the system manager of supercomputing resources. He has been working more than 20 years on HPC system architecture, system software, and performance evaluation on various scientific applications. He has played the central role of system development on CP-PACS (ranked as number one in the TOP500 in 1996), FIRST (hybrid cluster with gravity accelerator), PACS-CS (bandwidth-aware cluster) and HA-PACS (high-density GPU cluster), among representative supercomputers in Japan. Also, He is a member of the system architecture working group of Post-K Computer development. He is a coauthor of ACM Gordon Bell Prize paper in 2011. http://www.hpcs.cs.tsukuba.ac.jp/~taisuke/index-e.html

Lightning Talks & Posters

Dalal Sukkari	Massively Parallel Polar Decomposition on Distributed-Memory Systems
Alexander Ukhov	Large-scale Rigid Body Dynamics Simulation on HPCs with Distributed Memory Architecture
Li Luo	Parallel Simulation of Bloodflow in 3D Patient-specific Cerebral/Coronary Arteries
Sameh Abdulah	Tile Low-Rank Approximation of Large-scale Likelihood Estimation on Manycore Architectures
Samar Aseeri	Benchmarking Distributed Memory Fast Fourier Transforms
Rashid Mehmood	Big Data and HPC Convergence
Aleksandr Mikhalev	STARS-H: a High-Performance H-Matrix Market Library on Large-Scale Systems
Samuel Kortas	Decimate: a Portable and Fault-tolerant Scheduler Extension
Abdulhamit Subasi	Epileptic Seizure Prediction using Rotation Forest in a Parallel Environment
Essam Algizawy	Highly Scalable Flexi-trace-based Pre-stack Time Migration
Khaled El-Amrawi	Revisited RTM IO Strategies with New Memory Hierarchies
Steffen Hirschmann	Load-balancing and hybrid parallelization for simulations with large particle numbers

http://epostersonline.com/ixpug-me2018/

12 posters, 30 authors, 15 institutions

IXPUG-ME 2018

APRIL 22-25, 2018, KING ABDULLAH UNIVERSITY OF SCIENCE AND TECHNOLOGY, THUWAL, SAUDI ARABIA

Home | Posters by ID | Posters by Title | Authors | Affiliations

Enter terms then hit Search...

Posters

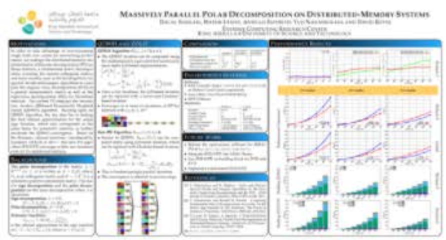

P01

Massively Parallel Polar Decomposition on Distributed-Memory Systems

Dalal Sukkari

Poster Presenter: Dalal Sukkari

No votes yet

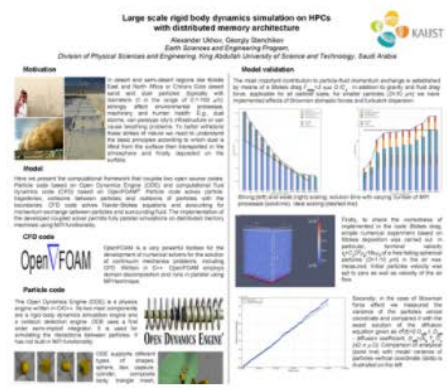

P02

Large-scale Rigid Body Dynamics Simulation on HPCs with Distributed Memory Architecture

Alexander Ukhov, Georgiy Stenchikov

King Abdullah University of Science and Technology

Poster Presenter: Alexander Ukhov

No votes yet

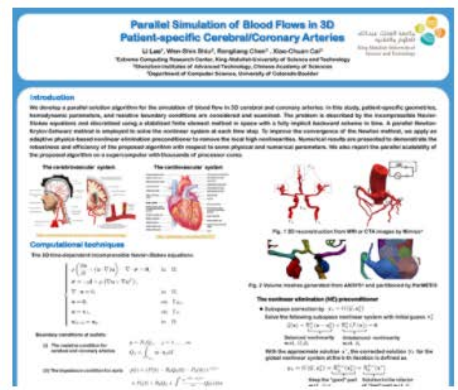

P03

Parallel Simulation of Bloodflow in 3D Patient-specific Cerebral/Coronary Arteries

Li Luo, Wen-Shin Shiu, Rongliang Chen, Xiao-Chuan Cai

Extreme Computing Research Center, King Abdullah University of Science and Technology, Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, Department of Computer Science, University of Colorado Boulder

Poster Presenter: Li Luo

No votes yet

E-poster schedule

- **Monday at 4:30pm**
  - **Abdulhamit Subasi**
  - **Aleksander Mikhalev**
  - **Alexander Ukhov**
  - **Khaled El-Amrawi**
  - **Li Luo**
  - **Rashid Mehmood**
  - **Sameh Abdulah**
- **Tuesday at 10:30am**
  - **Dalal Sukkari**
  - **Essam Algizawy**
  - **Samar Aseeri**
  - **Samuel Kortas**
  - **Steffen Hirschmann**

At other times E-posters are available on demand in their digital frames

Contributed talks

Jahanzeb Hashmi	Designing High-Performance and Scalable Collectives for the Many-core Era: The MVAPICH2 Approach
Hatem Ltaief	Real-Time Massively Distributed Multi-Object Adaptive Optics Simulations for the European Extremely Large Telescope
Kadir Akbudak	Cholesky Factorization on Tile Low-Rank Matrices for Distributed-Memory Systems
Amani Alonazi	Asynchronous Task-Based Parallelization of Iterative Algebraic Solvers
Mohamed Ibrahim	Screen-Space Normal Distribution Function Caching for Consistent Multi-Resolution Rendering of Large Particle Data
Rabab Alomairy	ALTANAL: Abstraction Layer for Task bAsed NumericAl Libraries
Mustafa AbdulJabbar	Extreme Scale FMM-Accelerated Wave Scattering Solver for a Complex 3D Helmholtz Boundary Integral Equation
Matteo Parsani	Sustained Petascale Direct Numerical Simulation of Secondary Flows in Square Ducts
Mohammed AlFarhan	Optimizations of Unstructured Aerodynamics Computations for Intel Xeon Phi (Knights Landing) Architecture

Invited Talks

Alex Heinecke	Deep Learning Hardware Accelerates Fused Discontinuous Galerkin Seismic Simulations
Jeff Hammond	Evaluating Data Parallelism in C++ Programming Models using the Parallel Research Kernels
Adrian Tate	Explicit Data Transport Machinery Supporting Producer-Consumer Workflows on Cray Systems
Taisuke Boku	Oakforest-PACS: Japan's Fastest Intel Xeon Phi Supercomputer and Its Applications
Vincent Etienne	Optimization of Finite-difference Kernels on Multicore Architectures for Seismic Applications

Keynote

5 | IXPUG Workshop at KAUST 2018

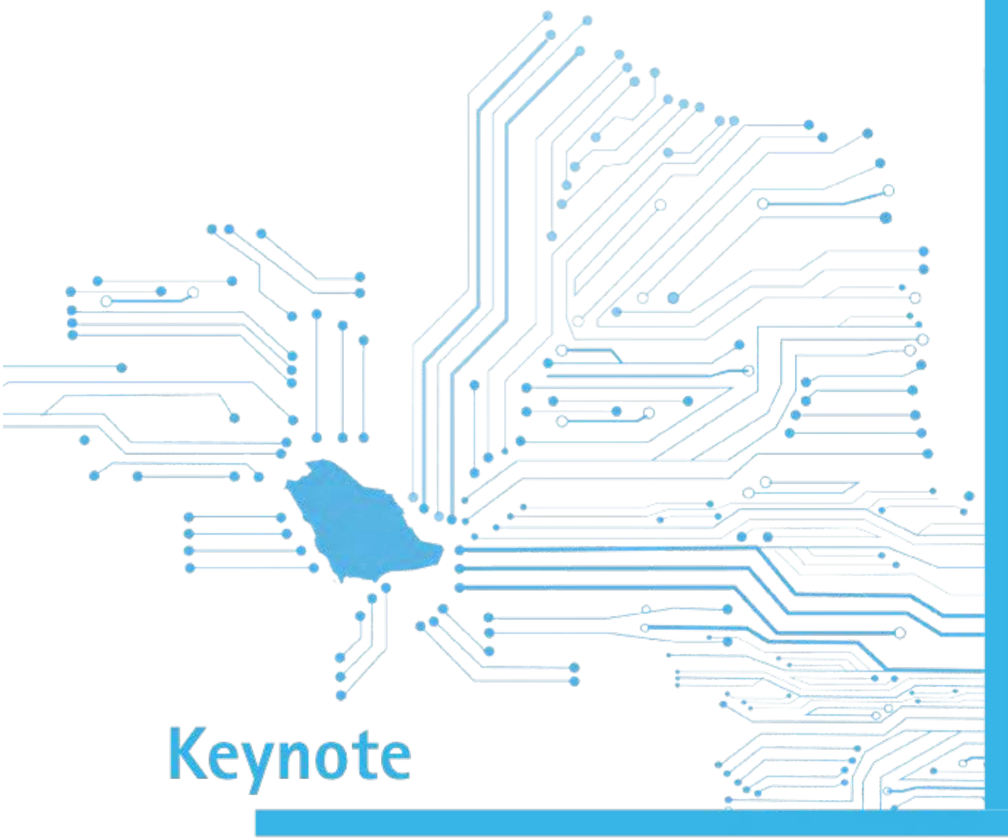

Dr. Alan GARA

The Intertwined Futures of High Performance Computing and Artificial Intelligence

We are seeing both the emergence of new workloads involving artificial intelligence as well as the maturation of new core technologies. These technologies and the new workloads present us with tremendous opportunities as well as challenges, for we now are architecting systems that can cover a broader range of applications. Namely, future architectures will target artificial intelligence and data analytics at scale, as well as traditional modeling and simulation applications. This talk will cover opportunities and challenges of AI and HPC and how the emerging new technologies will be developed to address the workloads and push computing into exascale in the near future.

Biosketch

Dr. Alan Gara is an Intel Fellow and Chief Architect of Enterprise and Government Advanced Development, within the Data Center Group at Intel Corporation. In this capacity he is leading a team of Intel architects in system pathfinding the future for high performance computing. Additionally, he is leading the Intel team responsible for delivering the CORAL system, a Collaboration of Oak Ridge, Argonne, and Livermore National Laboratories.

Prior to joining Intel in 2011, Dr. Gara was an IBM fellow and Chief Architect for three generations of the Blue Gene platforms, awarded the U.S. National Medal of Technology and Innovation in 2008. Al has been the Chief Architect on more than one-third of the top10 systems over the last 10 years as measured by the Top500. Al has received two Gordon Bell prizes (1998 and 2006) and the Seymour Cray award in 2010. He has over 70 publications in computer science and physics and more than 130 US patents in the area of computer design and architecture.

Gara received his PhD in physics from the University of Wisconsin, Madison.

6

Shaheen-1 (2009) → Shaheen-2 (2015)

IBM Blue Gene/P	Cray XC40
Speed: **.222 Petaflop/s (peak)** Launch ranking: **#14 HPL R_{max}**	Speed: **7.3 Petaflop/s (peak)** **↑ ~33X** Launch ranking**: #7 HPL R_{max}**
Power: **.5 MW (0.44 GF/s/W)**	Power: **2.8 MW (2.6 GF/s/Watt)** **↑ ~6X** Lag relative to flop/s **↓ ~5.5X**
Memory: **65 TeraBytes**	Memory: **793 TeraBytes (↑ ~12X)** Lag relative to flop/s: **↓ ~3X**
Storage: **2.7 PetaBytes**	Storage: **17.6 PetaBytes (↑ ~6.5X)** Lag relative to flop/s: **↓ ~5X**
Bandwidth: **25GB/s**	Bandwidth: **500GB/s (↑ ~20X)** Lag relative to flop/s: **↓ ~1.7X**

Design constraints have changed

Old Constraints	New Constraints
Performance limited by **clock**	Performance limited by **power**
Flops biggest cost: **minimize computation**	Data motion biggest cost: **minimize data movement**
Concurrency from **adding nodes with own memory**	Concurrency from **more cores within a node sharing memory**

c/o W. Harrod, US Department of Energy

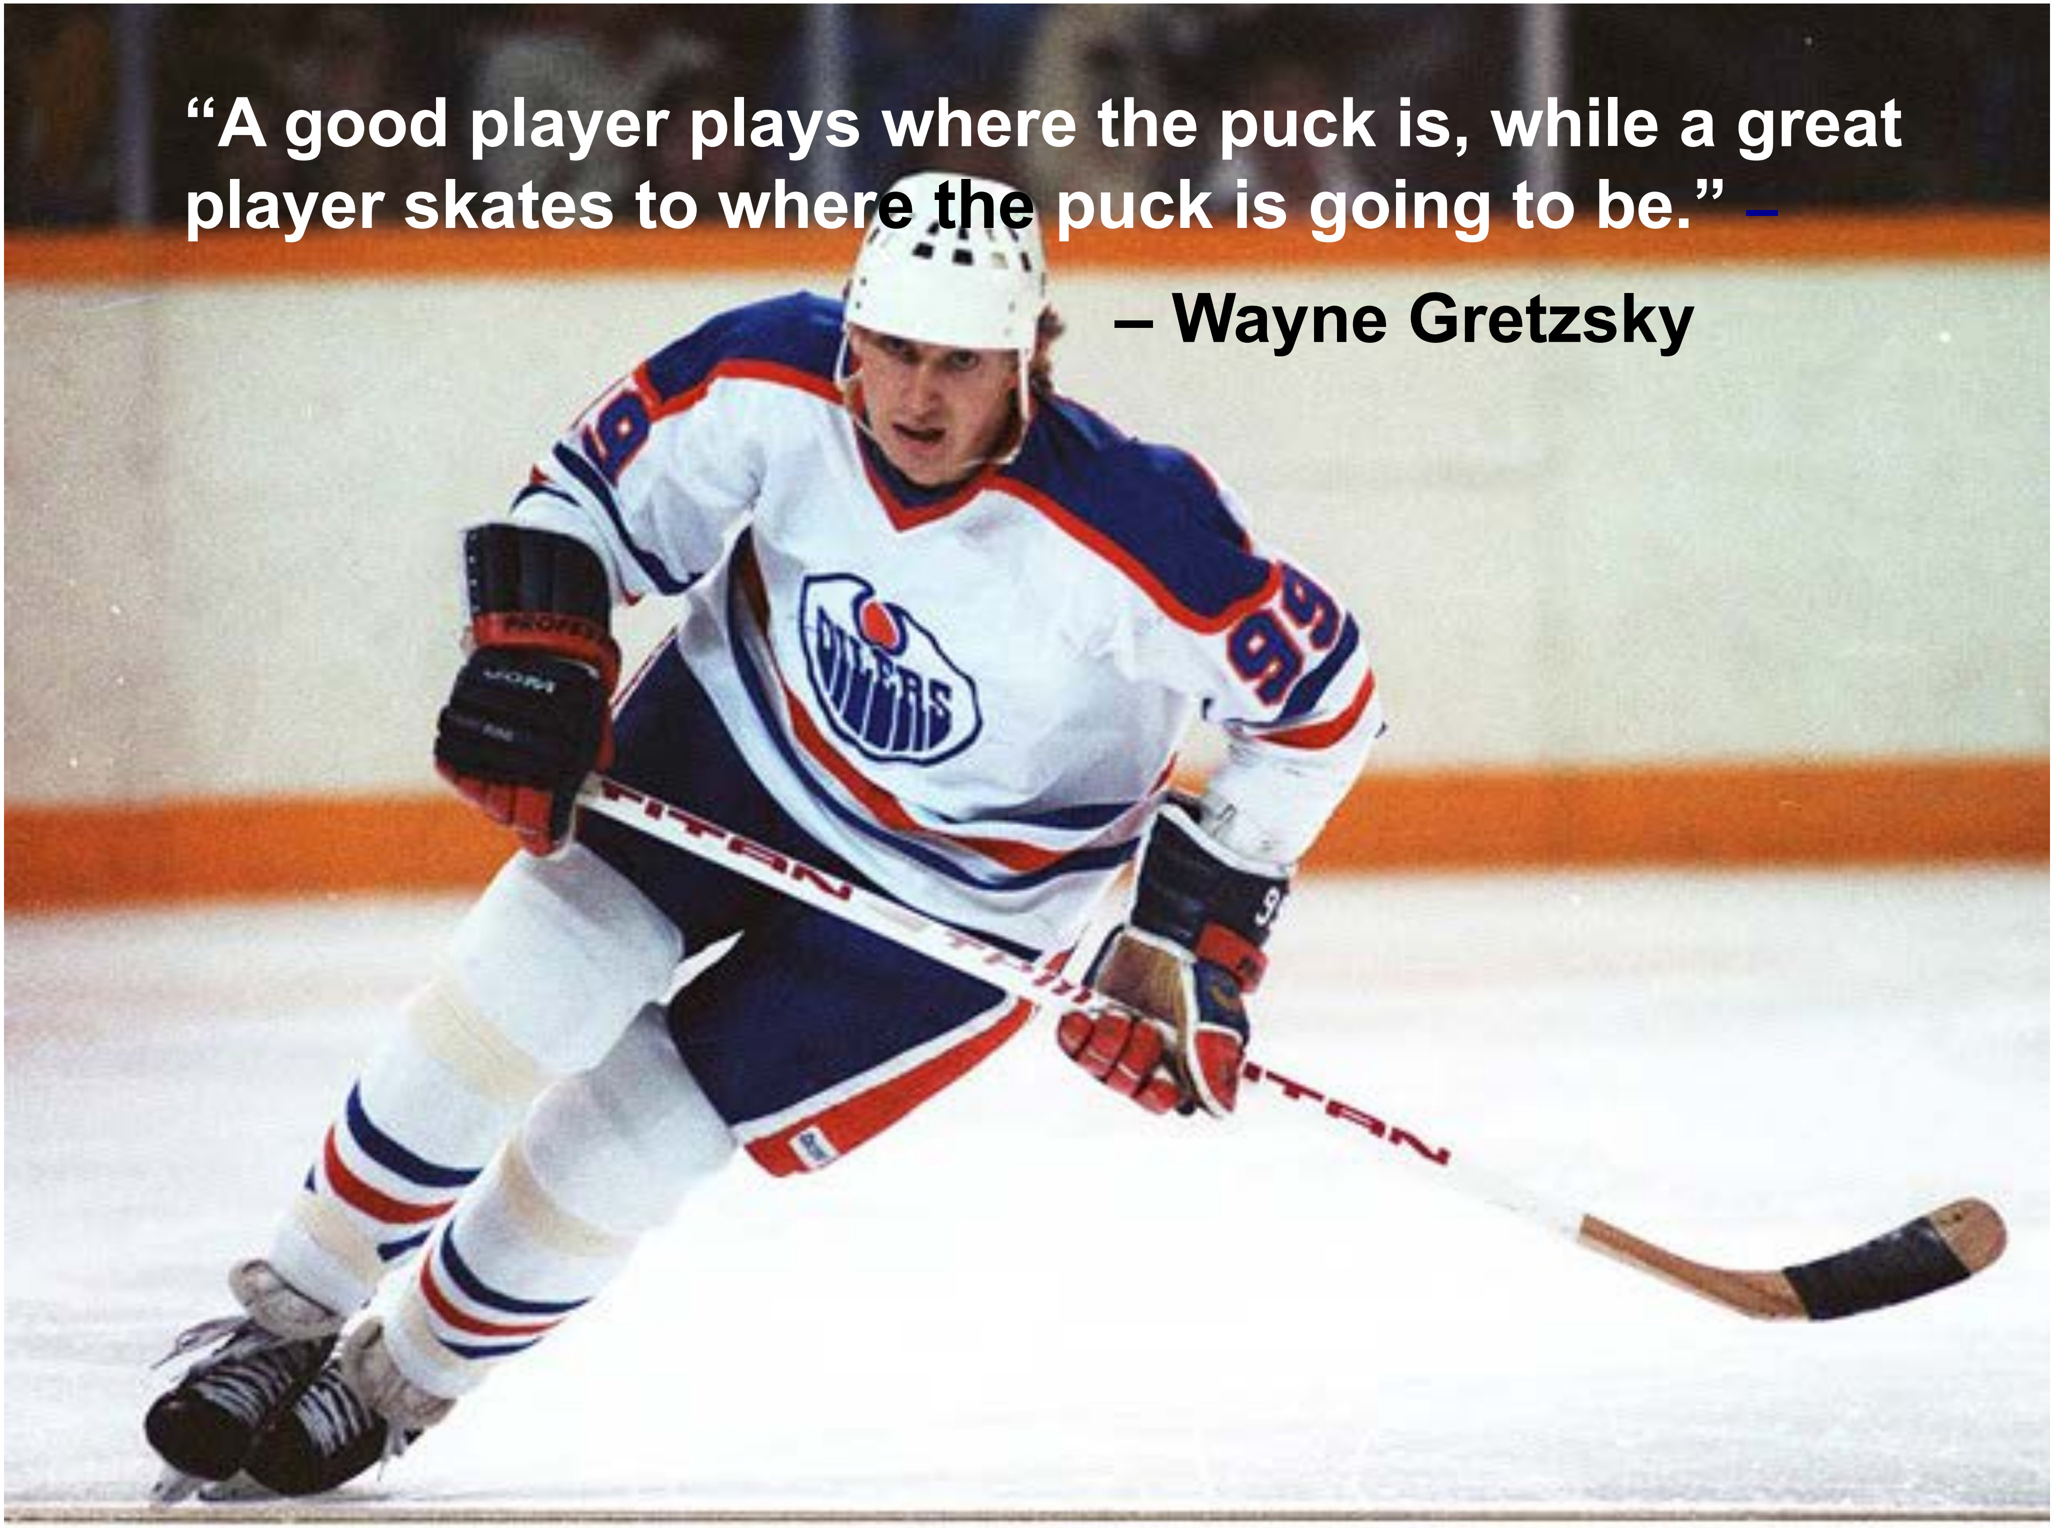
“A good player plays where the puck is, while a great player skates to where the puck is going to be.” –
– Wayne Gretzky

KAUST'S Extreme Computing Research Center

Embracing the opportunities of exascale

Extreme Computing Research Center in one slide

We develop and demonstrate core algorithms for scientific simulation and data analytics on emerging architectures.

Today's scientific libraries are designed for a 25-year-old programming model and need to be evolved for: ↑more concurrency (cores), ↓less memory per core, ↓less memory bandwidth per core, and ↓less global synchronization. Primary motivating applications include energy and environment.

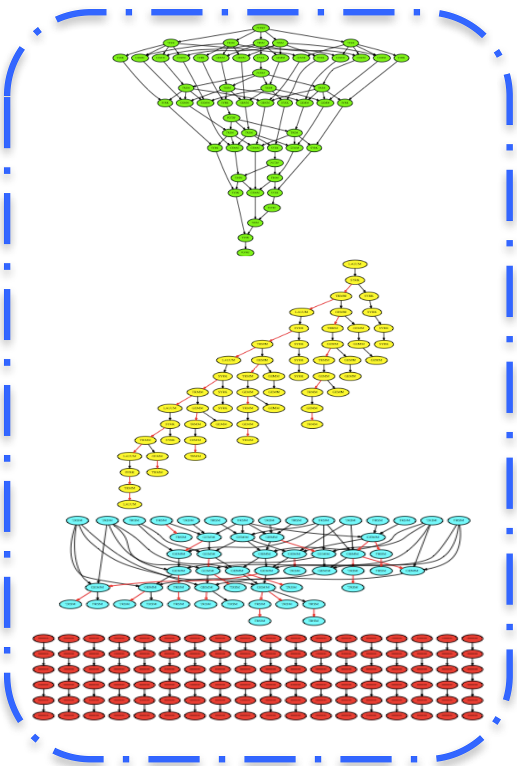

Example: four sequential task graphs of an eigensolver are merged into one that is ***shorter*** *(faster) &* ***wider*** *(more concurrent)*

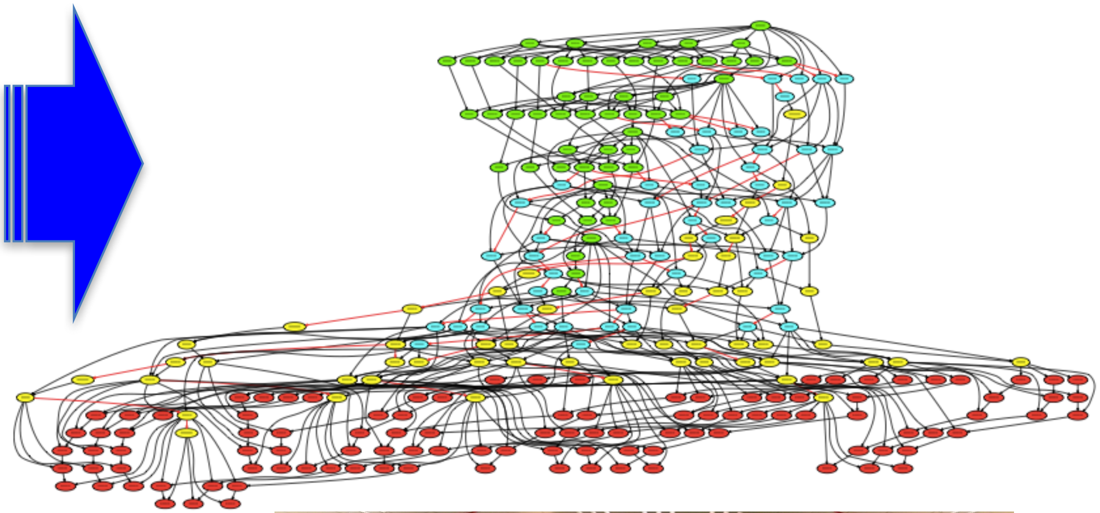

Best-in-class algorithms raise impact and increase the return on investment of scientific simulation across KAUST.

The diagram that launched SciDAC in the US

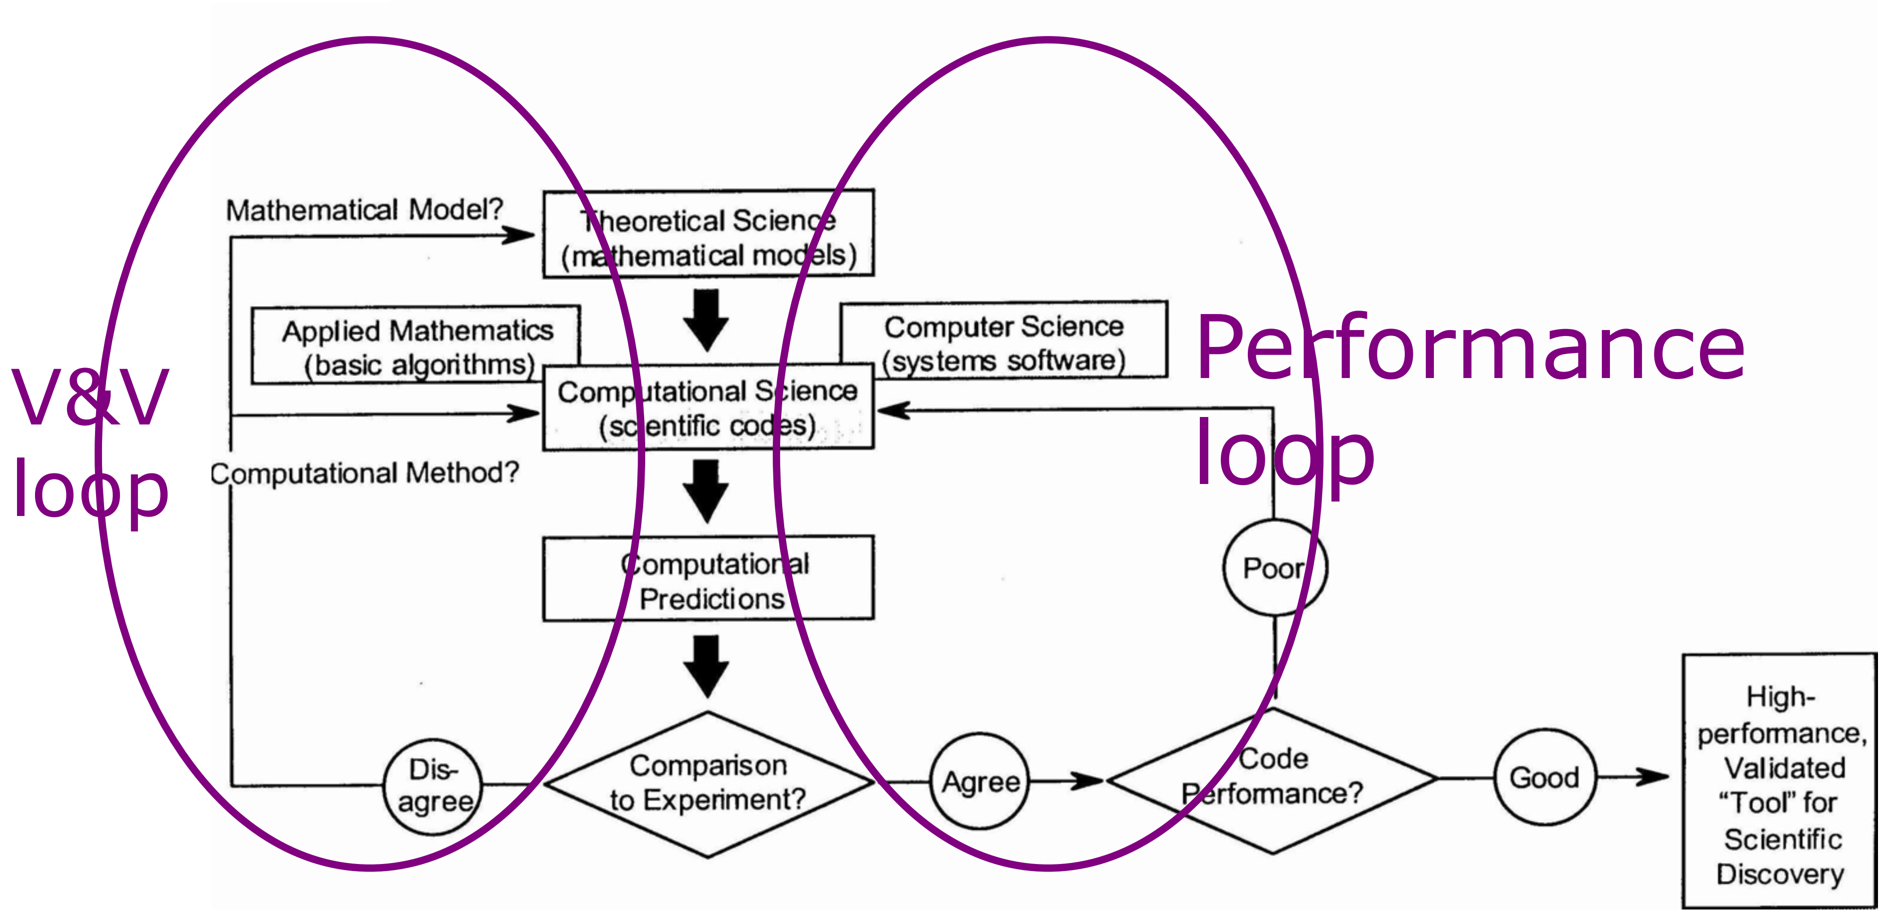

c/o T. Dunning, UIUC/NNSA, 2000 SciDAC report

Scientific Computing at KAUST

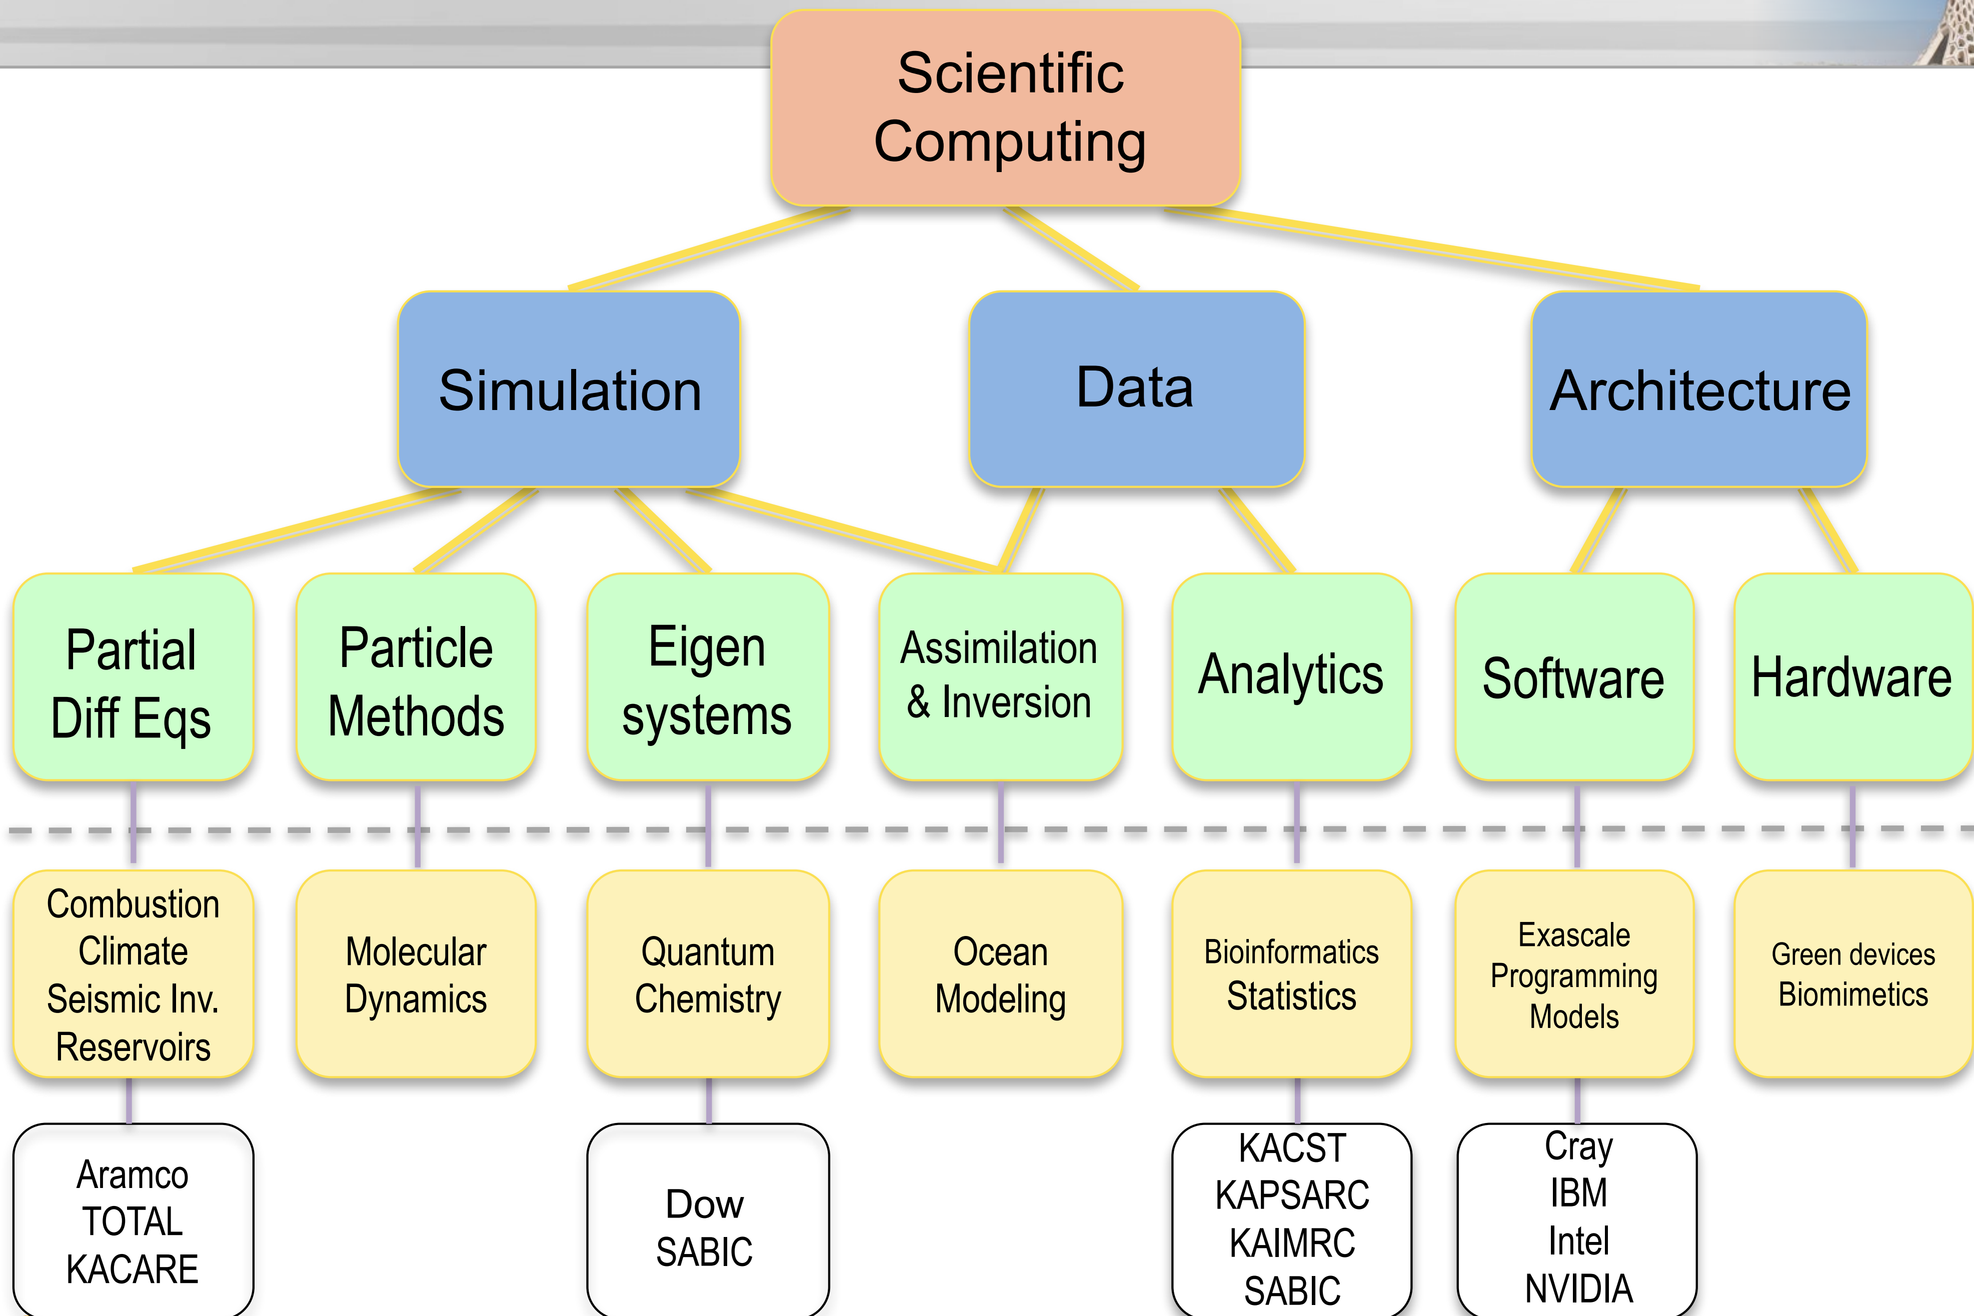

Applications – the visible impact

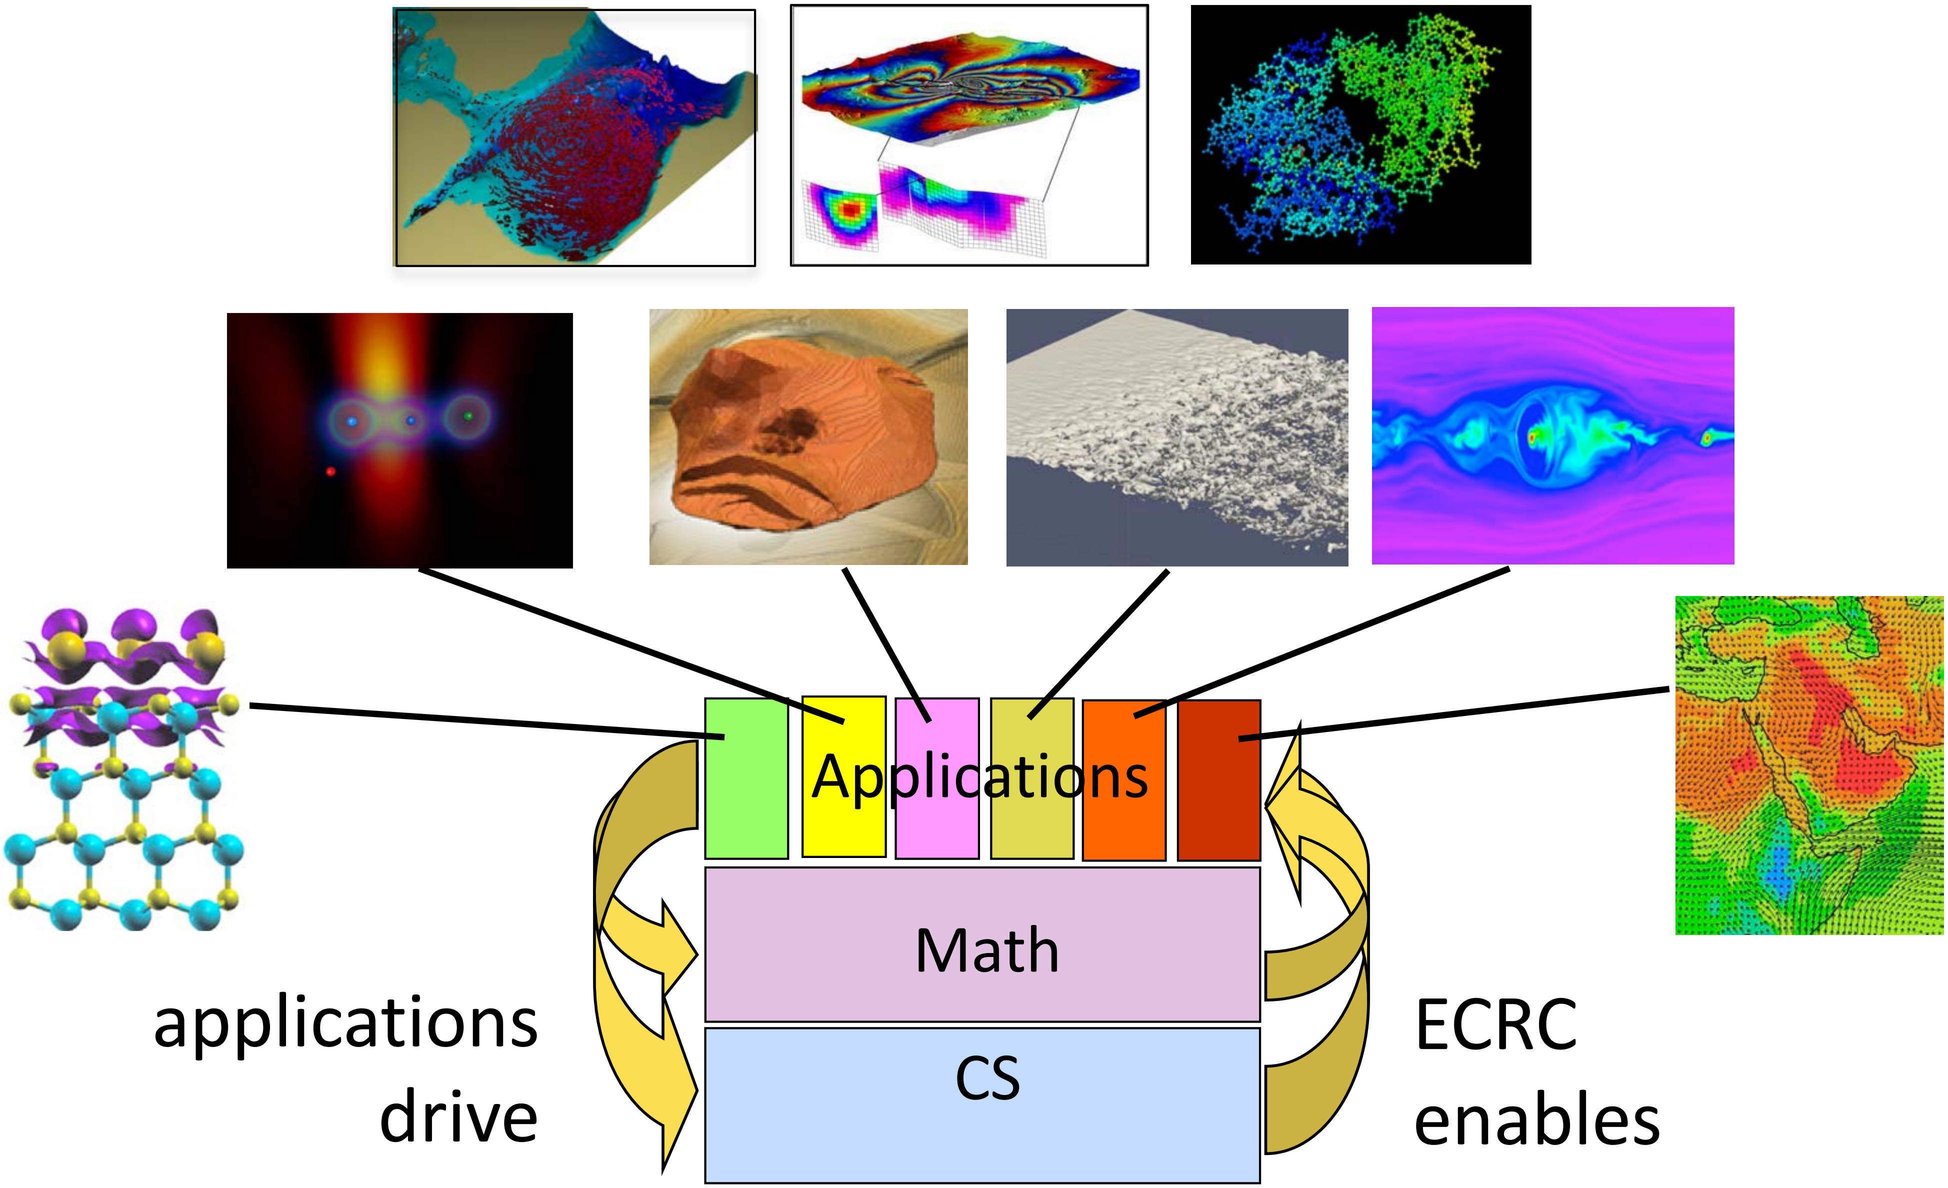

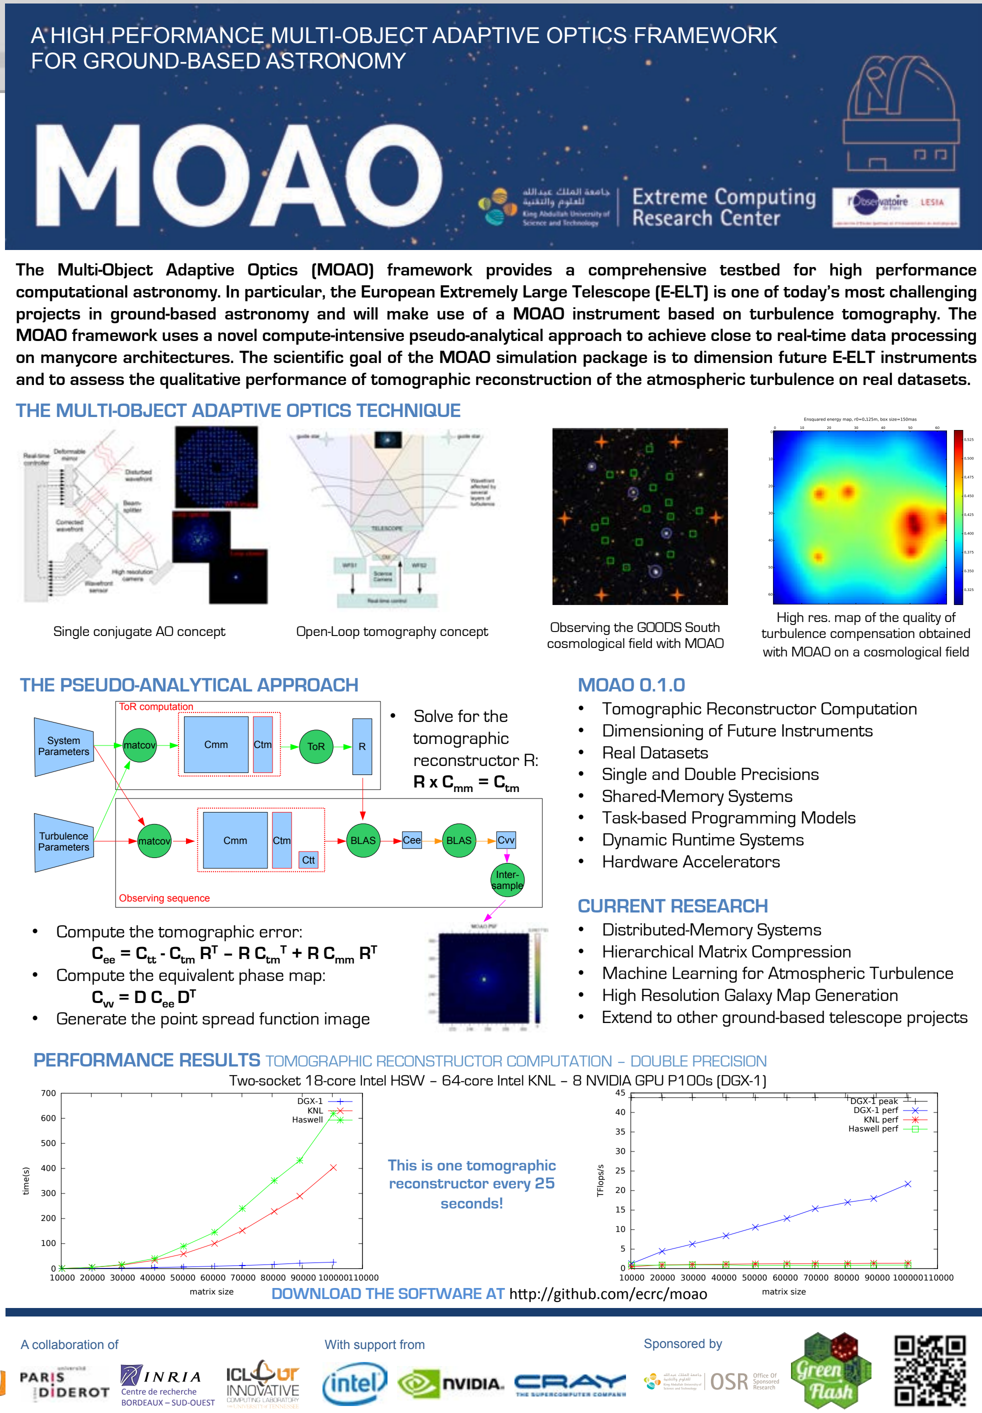

https://github.com/ecrc

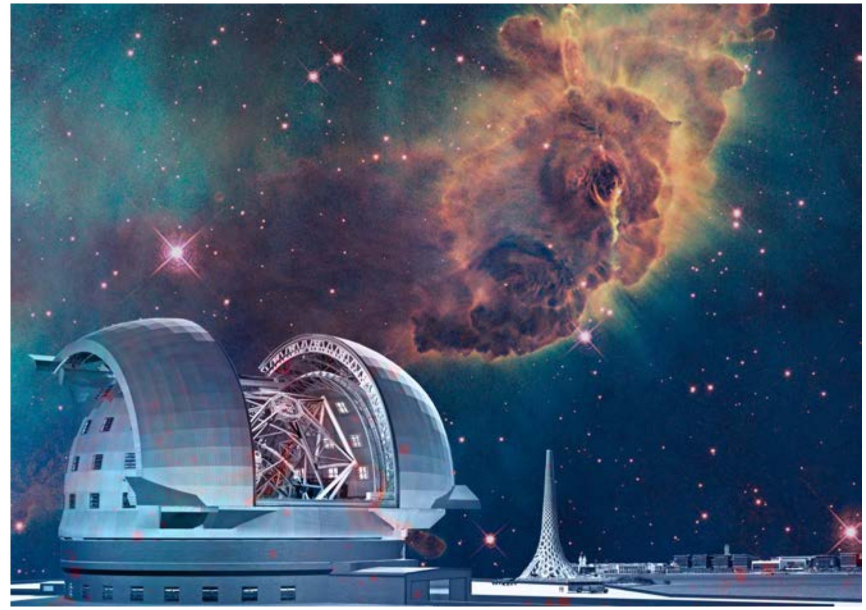

Used in MOSAIC instrument in Europe's Extremely Large Telescope (E-ELT, Chile) design and in Japanese Subaru telescope (Hawaii)

Hatem Ltaief

Fields of View

Twinkling stars are cute
But MOSAIC can freeze them
So we see the past.

https://github.com/ecrc

HIERARCHICAL COMPUTATIONS ON MANYCORE ARCHITECTURES

HiCMA

Extreme Computing Research Center

The Hierarchical Computations on Manycore Architectures (HiCMA) library aims to redesign existing dense linear algebra libraries to exploit the data sparsity of the matrix operator. Data sparse matrices arise in many scientific problems (e.g., in statistics-based weather forecasting, seismic imaging, and materials science applications) and are characterized by low-rank off-diagonal tile structure. Numerical low-rank approximations have demonstrated attractive theoretical bounds, both in memory footprint and arithmetic complexity. The core idea of HiCMA is to develop fast linear algebra computations operating on the underlying tile low-rank data format, while satisfying a specified numerical accuracy and leveraging performance from massively parallel hardware architectures.

TILE LOW-RANK ALGORITHMS CHOLESKY FACTORIZATION

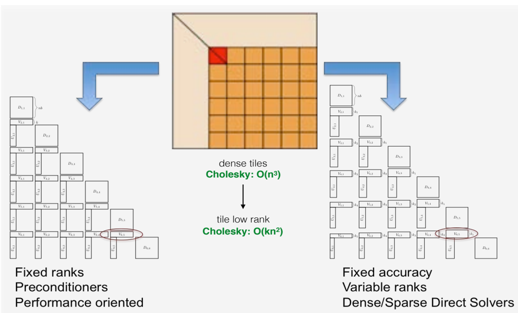

SOFTWARE STACK

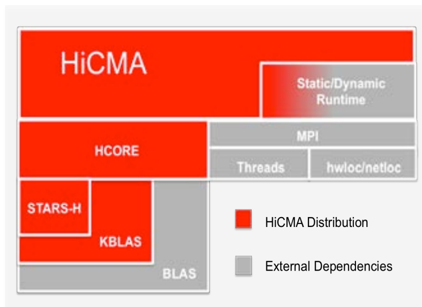

GEOSPATIAL STATISTICS

N = 20000, NB = 500, acc=10^{9}, 2D problem sq. exp.

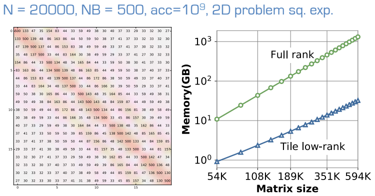

HiCMA 0.1.0

- Matrix-Matrix Multiplication
- Cholesky Factorization/Solve
- Double Precision
- Task-based Programming Models
- Shared and Distributed-Memory Environments
- Support for StarPU Dynamic Runtime Systems
- Testing Suite and Examples

CURRENT RESEARCH

- LU Factorization/Solve
- Schur Complements
- Preconditioners
- Hardware Accelerators
- Support for Multiple Precisions
- Autotuning: Tile Size, Fixed Accuracy and Fixed Ranks
- Support for OpenMP, PaRSEC and Kokkos
- Support for HODLR, H, HSS and H^2

PERFORMANCE RESULTS CHOLESKY FACTORIZATION – DOUBLE PRECISION – CRAY XC40 WITH TWO-SOCKET, 16-CORE HSW

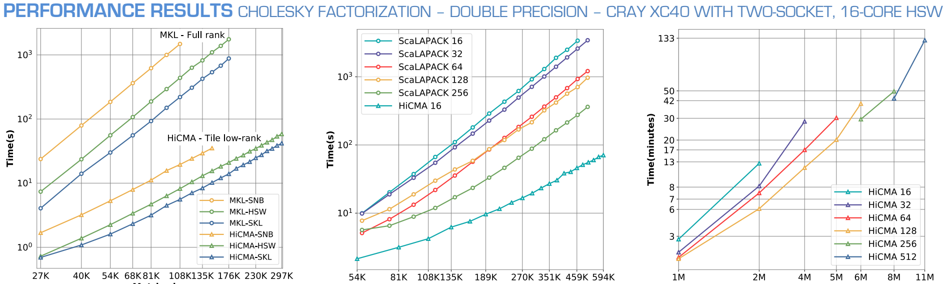

DOWNLOAD THE SOFTWARE AT http://github.com/ecrc/hicma

A collaboration of

Tokyo Tech

INRIA Centre de recherche BORDEAUX – SUD-OUEST

ICL INNOVATIVE COMPUTING LABORATORY

With support from

intel NVIDIA CRAY THE SUPERCOMPUTER COMPANY

Sponsored by

OSR Office Of Sponsored Research

Kadir Akbudak

HiCMA*

Energy dictates
Structure and storage limits –
HiCMA's prime targets.

* "Hikmah" is the Arabic word for wisdom

https://github.com/ecrc

PARALLEL HIGH PERFORMANCE UNIFIED FRAMEWORK FOR GEOSTATISTICS ON MANY-CORE SYSTEMS

ExaGeoStat

Extreme Computing Research Center

The Exascale GeoStatistics project (ExaGeoStat) is a parallel high performance unified framework for computational geostatistics on many-core systems. The project aims at optimizing the likelihood function for a given spatial data to provide an efficient way to predict missing observations in the context of climate/weather forecasting applications. This machine learning framework proposes a unified simulation code structure to target various hardware architectures, from commodity x86 to GPU accelerator-based shared and distributed-memory systems. ExaGeoStat enables statisticians to tackle computationally challenging scientific problems at large-scale, while abstracting the hardware complexity, through state-of-the-art high performance linear algebra software libraries.

ExaGeoStat Dataset Generator

- Generate 2D spatial Locations using uniform distribution.
- Matérn covariance function:

$$C(r;\theta) = \frac{\theta_1}{2^{(\theta_3-1)}\Gamma(\theta_3)}\left(\frac{r}{\theta_2}\right)^{\theta_3} \mathrm{K}_{\theta_3}\left(\frac{r}{\theta_2}\right)$$

- Cholesky factorization of the covariance matrix:

$$\Sigma(\theta) = V \,.\, V^T$$

- Measurement vector generation (Z):

$$\vec{Z} = V \,.\, \vec{e}, \quad \vec{e_i} \sim N(0,1)$$

Figure: An example of 400 points irregularly distributed in space, with 362 points (o) for maximum likelihood estimation and 38 points (×) for prediction validation.

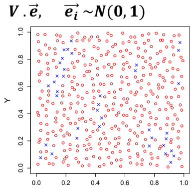

ExaGeoStat Maximum Likelihood Estimator

- Maximum Likelihood Estimation (MLE) learning function:

$$\ell(\theta) = -\frac{n}{2}\log(2\pi) - \frac{1}{2}\log|\Sigma(\theta)| - \frac{1}{2}Z^T\,\Sigma(\theta)^{-1}Z$$

Where $\Sigma(\theta)$ is a covariance matrix with entries

$$\Sigma_{ij} = C(s_i - s_j;\theta), i,j = 1,\dots,n$$

Figure: Two different examples of real datasets (wind speed dataset in the middle east region and soil moisture dataset coming from Mississippi region, USA).

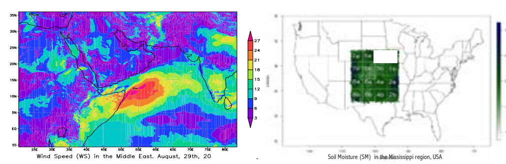

ExaGeoStat 0.1.0

- Large-scale synthetic geo-spatial datasets generator
- Maximum Likelihood Estimation (MLE)
  - Synthetic and real datasets
- A large-scale prediction tool for unknown measurements on known locations

Current Research

- ExaGeoStat R-wrapper package
- Tile Low Rank (TLR) approximation
- NetCDF format support
- PaRSEC runtime system
- In-situ processing

ExaGeoStat Predictor

- MLE prediction problem

$$\begin{bmatrix} Z_1 \\ Z_2 \end{bmatrix} \sim N_{m+n}\left(\begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}, \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}\right)$$

With $\Sigma_{11} \in \Re^{m\times m}$, $\Sigma_{12}\Re^{m\times n}$, $\Sigma_{21} \in \Re^{n\times m}$, and $\Sigma_{22} \in \Re^{n\times n}$

- The associated conditional distribution where Z_1 represents a set of unknown measurements :

$$Z_1|Z_2 \sim N_m(\mu_1 + \Sigma_{12}\Sigma_{22}^{-1}(Z_2 - \mu_2), \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21})$$

Figure: Mean square error for predicting large scale synthetic dataset.

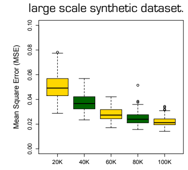

Performance Results (MLE)

Two-socket shared memory Intel x86 architectures

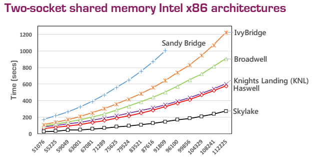

Intel two-socket Haswell + NVIDIA K80

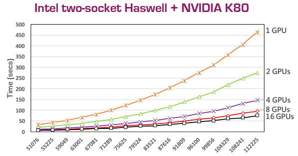

Cray XC40 with two-socket, 16 cores Haswell

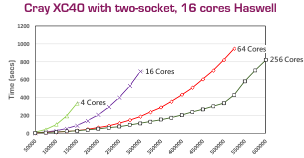

DOWNLOAD THE LIBRARY AT http://github.com/ecrc/exageostat

A collaboration of INRIA Centre de recherche BORDEAUX – SUD-OUEST, ICL UT INNOVATIVE COMPUTING LABORATORY

With support from intel, NVIDIA, CRAY THE SUPERCOMPUTER COMPANY

Sponsored by OSR Office Of Sponsored Research

Sameh Abdulah

ExaGeoStat

Covariances
In the billions require
ExaGeoStat.

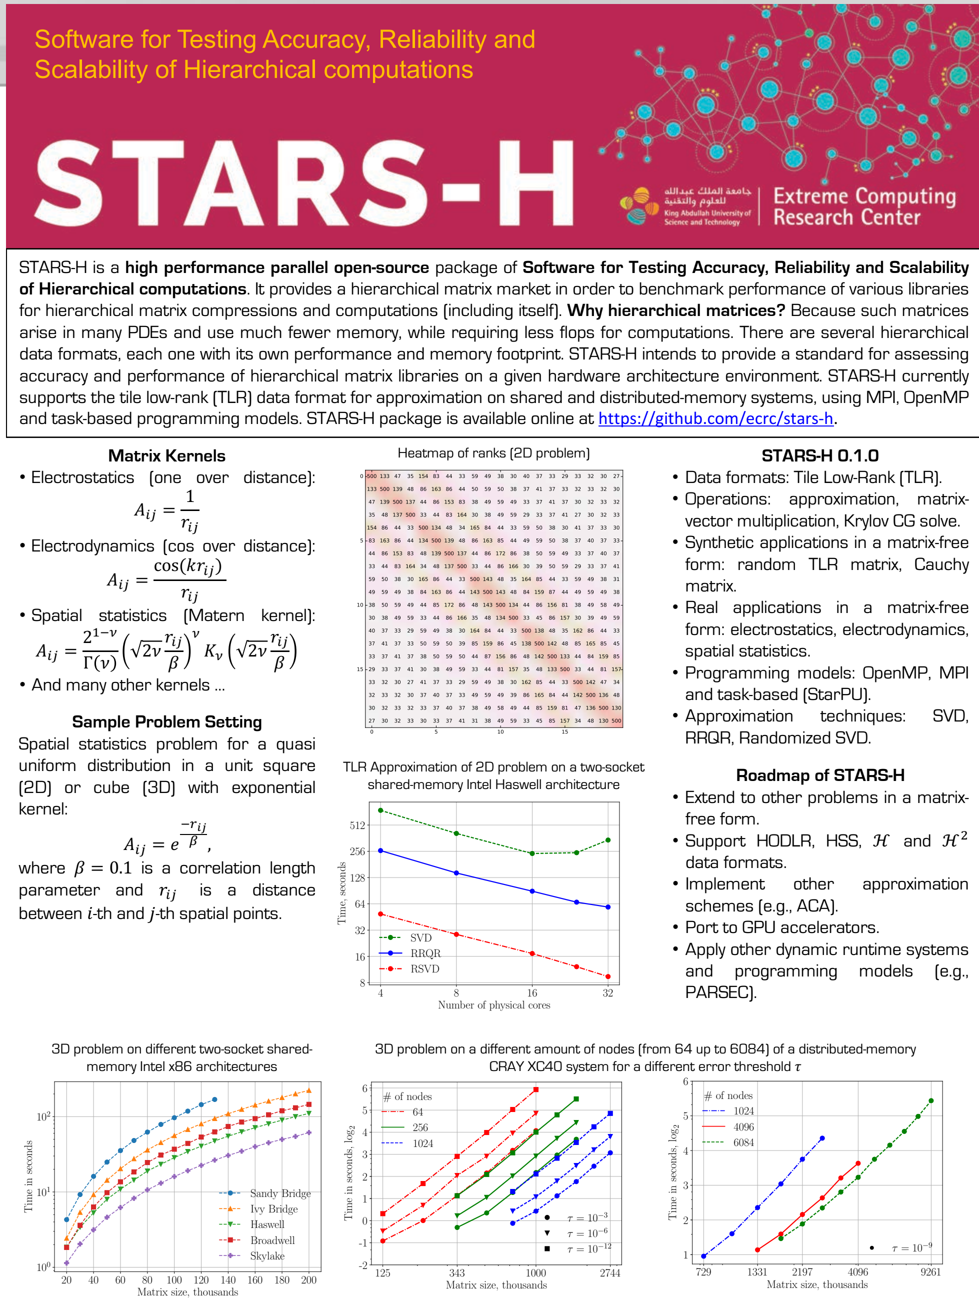

https://github.com/ecrc

STARS-H: Software for Testing Accuracy, Reliability, and Scalability of Hierarchical matrix computations

Aleksandr Mikhalev

STARS-H

Data sparsity
Hides everywhere in plain sight
STARS-H unmasks it.

https://github.com/ecrc

A QDWH-Based SVD Software Framework on Distributed-Memory Manycore Systems

KSVD

Extreme Computing Research Center

The KAUST SVD (KSVD) is a high performance software framework for computing a dense SVD on distributed-memory manycore systems. The KSVD solver relies on the polar decomposition using the QR Dynamically-Weighted Halley algorithm (QDWH), introduced by Nakatsukasa and Higham (SIAM Journal on Scientific Computing, 2013). The computational challenge resides in the significant amount of extra floating-point operations required by the QDWH-based SVD algorithm, compared to the traditional one-stage bidiagonal SVD. However, the inherent high level of concurrency associated with Level 3 BLAS compute-bound kernels ultimately compensates the arithmetic complexity overhead and makes KSVD a competitive SVD solver on large-scale supercomputers.

The Polar Decomposition

- $A = U_pH$, A in $R^{m \times n}$ $(m \geq n)$, where U_p is orthogonal Matrix, and H is symmetric positive semidefinite matrix

Application to SVD

- $A = U_pH$
  $= U_p(V\Sigma V^T) = (U_pV)\Sigma V^T = U\Sigma V^T$

QDWH Algorithm

- Backward stable algorithm for computing the QDWH-based SVD
- Based on conventional computational kernels, i.e., Cholesky/QR factorizations (≤ 6 iterations for double precision) and GEMM
- The total flop count for QDWH depends on the condition number κ of the matrix

Advantages

- Performs extra flops but nice flops
- Relies on compute intensive kernels
- Exposes high concurrency
- Maps well to GPU architectures
- Minimizes data movement
- Weakens resource synchronizations

Performance Results

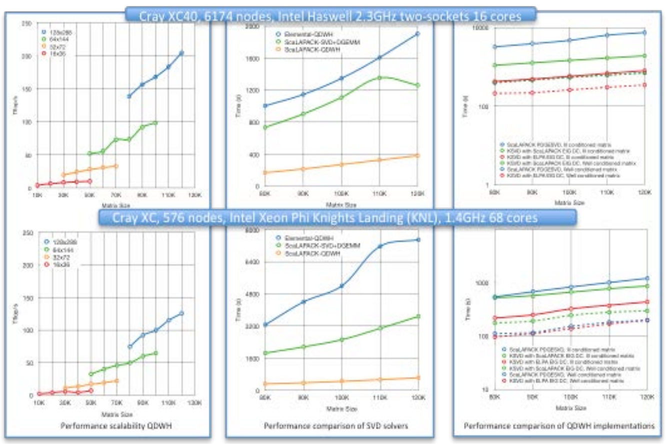

KSVD 1.0

- QDWH-based Polar Decomposition
- Singular Value Decomposition
- Double Precision
- Support to ELPA Symmetric Eigensolver
- Support to ScaLAPACK D&C and MR^3 Symmetric Eigensolvers
- ScaLAPACK Interface / Native Interface
- ScaLAPACK-Compliant Error Handling
- ScaLAPACK-Derived Testing Suite
- ScaLAPACK-Compliant Accuracy

Cray LibSci 17.11.1

State-of-the-Art

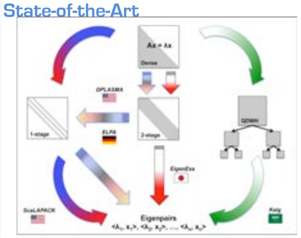

Current Research

- Asynchronous, Task-Based QDWH
- Dynamic Scheduling
- Hardware Accelerators
- Distributed Memory Machines
- Asynchronous, Task-Based QDWH-SVD
- QDWH-based Eigensolver (QDWH-EIG)
- Integration into PLASMA/ MAGMA

Chameleon 1.9

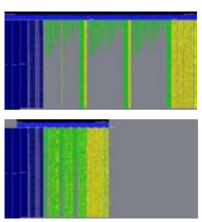

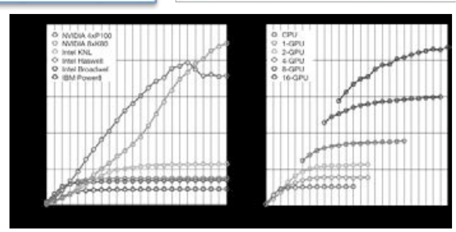

Download the software at http://github.com/ecrc/ksvd

A collaboration of — University of Oxford, INRIA Centre de recherche Bordeaux – Sud-Ouest, ICL UT Innovative Computing Laboratory

With support from — Intel, NVIDIA, Cray The Supercomputer Company

Sponsored by — KAUST OSR Office of Sponsored Research

Dalal Sukkari

Complexity Made Simple

It's not the flop count
But the data location;
The paradigm shifts.

https://github.com/ecrc

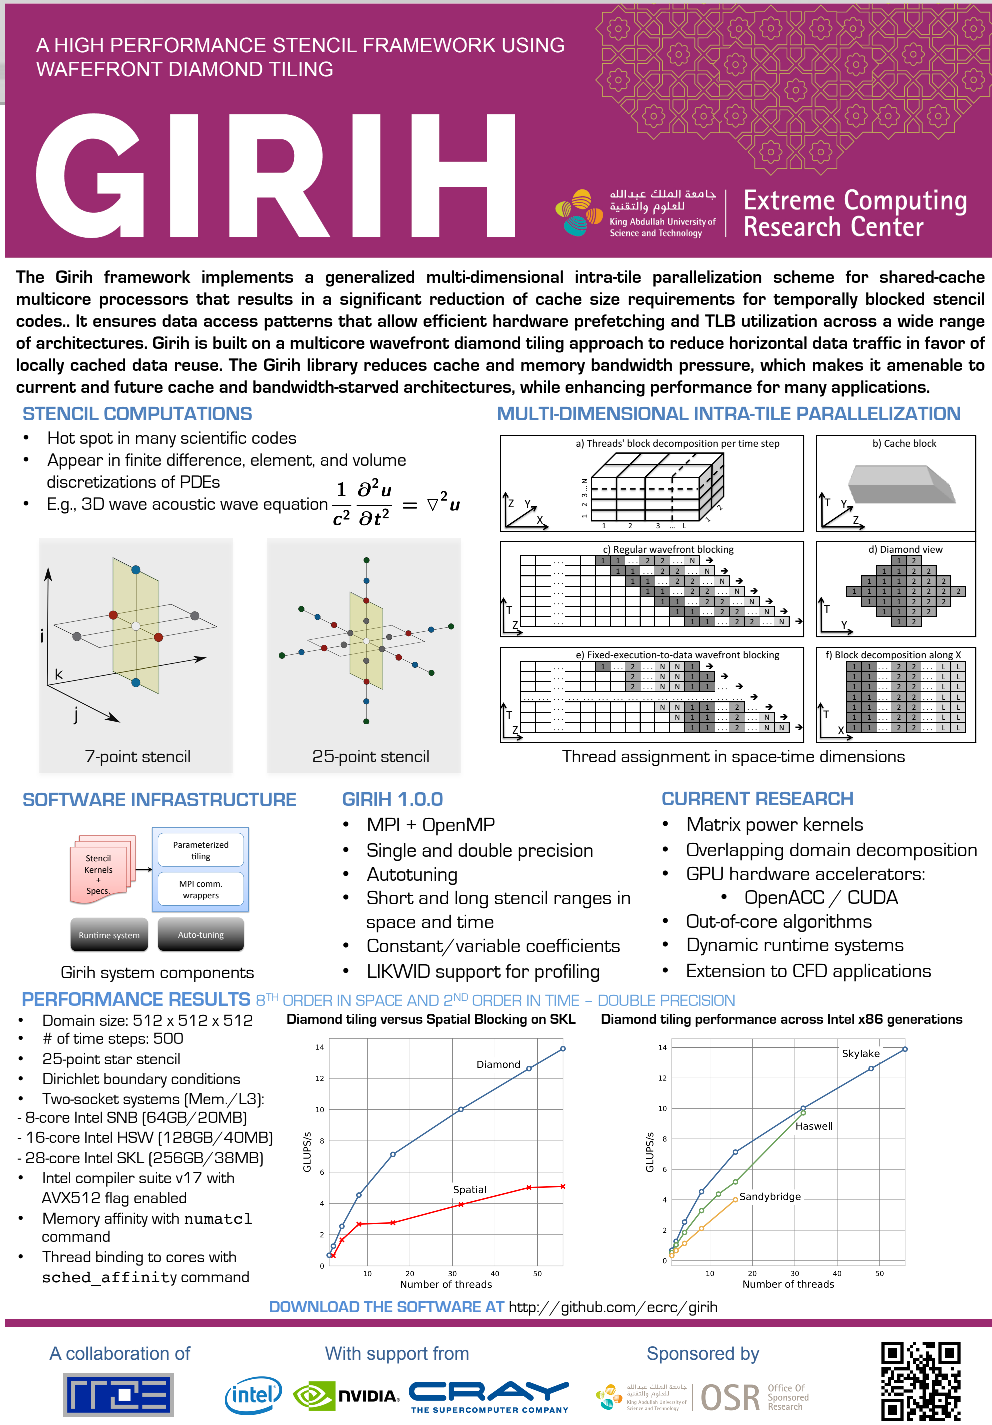

Now employed in seismic forward modeling steps in Aramco's next-gen seismic inversion code

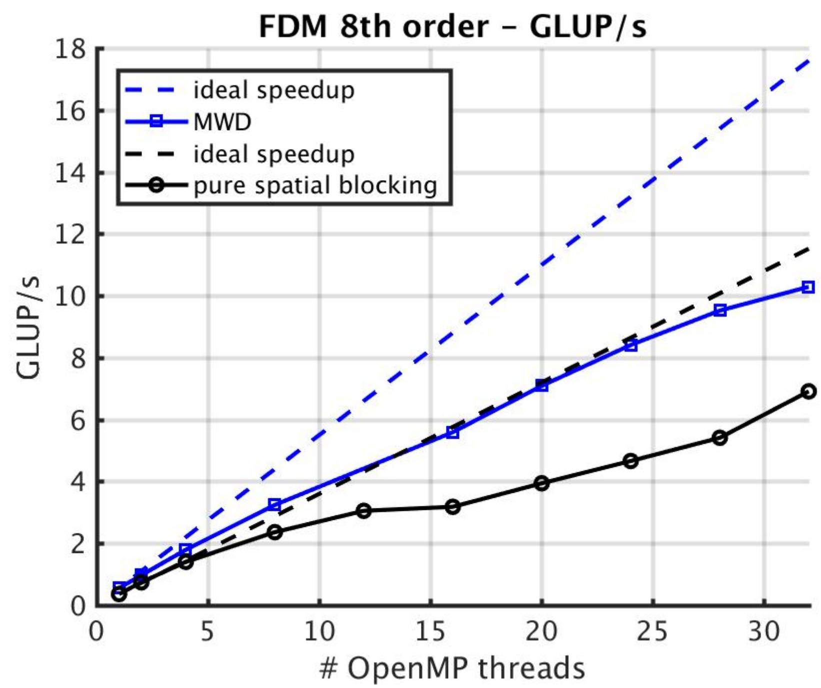

Tareq Malas

GIRIH

Many cores learn to share!
Cooperation is good
For limited cache.

https://github.com/ecrc

Ali Charara

Houdinis of Time

Houdinis of time,
We batch up concurrent tasks
From sequential streams.

Mohammed Farhan

Unstructured Grids

Compute rate suffers
When for a mesh we must ask:
Who is my neighbor?

Mustafa Abduljabbar

Space-filling curves

A curve fills 3-D,
And now submanifolds, too,
Touching every point.


Amani Alonazi

Taskification

Bulk synchrony worked,
But exascale demands more:
Dynamic future!

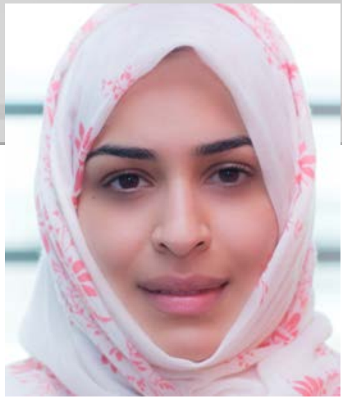

Rabab AlOmairy

Al Tanal

Like C-3-P-0,
Al Tanal **masters all tongues.**
Choose runtime at run!

Li Luo

Nonlinear Preconditioning

When stagnation strikes,
Isolate the stubborn DOFs
And update them first.

Yours Truly

Curse of Dimension

Curse of dimension,
Can you be mitigated
By blessing of rank?

Thanks to our sponsors!

Intel Parallel Computing Center

Cray Center of Excellence

KAUST Innovation & Economic Development

Thanks to our organizers!

IXPUG
INTEL EXTREME PERFORMANCE USERS GROUP


ECRC Extreme Computing Research Center



Group photo 10:30am, tiered on steps outside

Sunset concert in your honor, 5:30pm, Library