

HPC User Forum@Tucson

# **HPC / AI Activities at KAUST & the Middle East**

Jysoo Lee

Facilities Director, Research Computing Core Labs  
King Abdullah University of Science and Technology

2023. 9. 7.

# Major HPC/AI Players in the Middle East




- **Saudi Arabia**

- KAUST
- Saudi Aramco
- KFUPM, KAU
- KACST
- NEOM

- **United Arab Emirates**

- G42
- NYU AD, MBZUAI
- ADIA Lab



**KAUST**  
SUPERCOMPUTING  
CORE LAB

# Saudi Arabia



# KAUST Shaheen II Supercomputer

<b>Compute</b>	<b>Node</b>	Processor type: Intel Haswell	2 CPU sockets per node @2.3GHz 16 processor cores per CPU
		6174 nodes	197,568 cores
		128 GB of memory per node	Over 790 TB total memory
	<b>Power</b>	Up to 3.5MW	Water cooled
	<b>Weight/Size</b>	More than 100 metrics tons	36 XC40 Compute cabinets, disk, blowers, management nodes
	<b>Speed</b>	7.2 Peta FLOPS peak performance	5.53 Peta FLOPS sustained LINPACK and ranked 113 <sup>th</sup> in the latest Top500 list
	<b>Network</b>	Cray Aries interconnect with Dragonfly topology	57% of the maximum global bandwidth between the 18 groups of two cabinets
<b>Storage</b>	<b>Storage</b>	Sonexion 2000 Lustre appliance	17.6 Peta Bytes of usable storage Over 500 GB/s bandwidth
	<b>Burst Buffer</b>	DataWarp	Intel Solid State Devices (SSD) fast data cache Over 1.5 TB/s bandwidth
	<b>Archive</b>	Tiered Adaptive Storage (TAS)	Hierarchical storage with 200 TB disk cache and 20 PB of tape storage, using a spectra logic tape library (Upgradable to 100 PB)

**Used by 115 KAUST faculty**



# Shaheen III Requirements


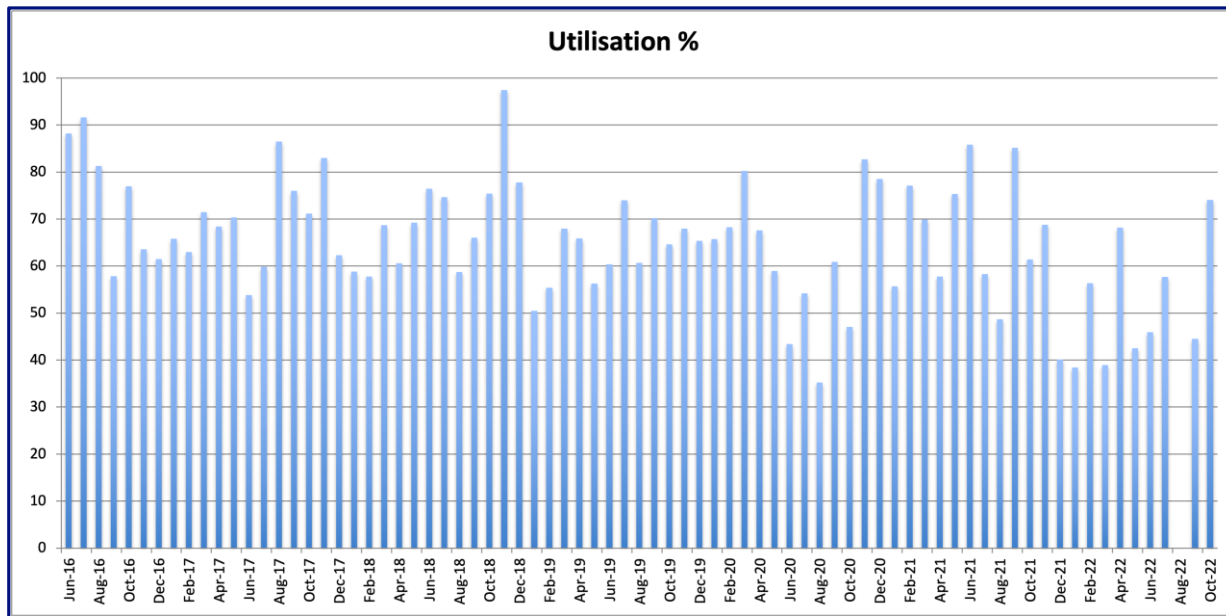
- Accommodate multiple workloads
  - HPC (Shaheen II)
  - AI (Ibex)
  - Hybrid
- Proposed solution
  - Single unified system
  - Modular architecture: CPU + GPU partition
  - Support multiple usage patterns
- Performance Requirements
  - CPU & GPU partition (real application performance)
  - Project & scratch storage
  - High-speed interconnect

# Shaheen III Specifications

- HPE Cray EX System
- CPU Partition
  - 4608 nodes (2 x 96 Cores AMD Genoa)
- GPU Partition
  - 704 nodes (4 x NVIDIA Grace Hopper Superchips)
- Storage
  - ClusterStor E1000: 20PB (scratch), 57PB (project)
- Interconnect: HPE Slingshot
- Cooling: closed-loop liquid cooling
- Schedule
  - Operational in 2023 (CPU), 2024 (GPU)

# Shaheen II Utilization

- Provided 7.4 Billion Core Hours
- Served 1467 Users





# Collaboration with Prof. Matteo Parsani: Full Simulation of Flows Around Racing Cars



Significantly enhance simulation accuracy at the full vehicle scale

- Detect and improve performance bottleneck
- Scaling the code up to full Shaheen II

## Performance Study of Sustained Petascale Direct Numerical Simulation on Cray XC40 Systems

Bilel Hadri<sup>1</sup> | Matteo Parsani<sup>2</sup> | Maxwell Hutchinson<sup>3</sup> | Alexander Heinecke<sup>4</sup> | Lisandro Dalcin<sup>2</sup> | David Keyes<sup>2</sup>

<sup>1</sup>KAUST Supercomputing Lab, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia

<sup>2</sup>Extreme Computing Research Center, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia

<sup>3</sup>Citrine Informatics, Redwood City, California, USA

<sup>4</sup>Intel Corporation, Santa Clara, California, USA

### Correspondence

\*Bilel Hadri, KAUST Supercomputing Lab, Al Khawarizmi Bldg. (1) Office 126, 4700 King Abdullah University of Science and Technology, Thuwal 23955-6900. Email: bilel.hadri@kaust.edu.sa

### Abstract

We present in this paper a comprehensive performance study of highly efficient extreme scale direct numerical simulations of secondary flows, using an optimized version of Nek5000. Our investigations are conducted on various Cray XC40 systems, using a very high-order spectral element method. Single-node efficiency is achieved by auto-generated assembly implementations of small matrix multiplies and key vector-vector operations, streaming lossless I/O compression, aggressive loop merging and selective single precision evaluations. Comparative studies across different Cray XC40 systems at scale, Trinity (LANL), Cori (NERSC) and ShaheenII (KAUST), show that a Cray programming environment, network configuration, parallel file system and burst buffer all have a major impact on the performance. All three systems possess a similar hardware with similar CPU nodes and parallel file system, but they have different theoretical peak network bandwidths, different OSs and different versions of the programming environment. Our study reveals how these slight configuration differences can be critical in terms of performance of the application. We also find that with 9216 nodes (294,912 cores) on Trinity XC40 the applications sustains petascale performance, as well as 50% of peak memory bandwidth over the entire solver (500 TB/s in aggregate). On 3072 Xeon Phi™ nodes of Cori, we reach 378 TFLOP/s with an aggregated bandwidth of 310 TB/s, corresponding to time-to-solution 2.11× faster than obtained with the same number of (dual-socket) Xeon® nodes.


### KEYWORDS:

Cray XC40, Haswell, KNL, Nek5000, Performance Analysis, Regression, Energy Efficiency



# Introduction to Data Science Workshop Series

- **Mission:** Equip students from KAUST and Saudi Arabia to build capacity in core data science tools and methods so they can begin building their data science applications.
- Widely popular series since 2018
- Consisted of hands-on single-day workshops,
  - Introduction to Python for Data Science
  - Introduction to Shell for (Data) Scientists
  - Introduction to Conda for (Data) Scientists
  - Introduction to SQL for Data Science
  - Introduction to Git for (Data) Scientists
- Statistics
  - More than 5000 participants registered
  - 40% In-Kingdom Participants
  - Since 2020, More than 6000 YouTube Views



# Building Communities

- Goal: Establishment of National HPC / AI Ecosystem
  - National HPC / AI Infrastructure
  - National HPC / AI Community
- Providing service to in-Kingdom organizations
  - Computing cycle (157 users from 18 organizations)
  - Consulting
  - Training
- Community event: e.g., HPC / AI Saudi




HPC Saudi 2017 Home Attendees Program Committee

# HPC SAUDI ARABIA

Building the HPC Community in the Kingdom

Conference March 13th - 15th, 2017  
Plenary Sessions March 13th - 14th, 2017  
Tutorials March 15th, 2017  
email contact@hpcsaudi.org  
Social f | t | in

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

REGISTRATION

Keynote Speakers

Dr. Steven E. Koonin  
Director, NYU Center for  
Urban Science and  
Progress

Dr. Thomas Schultness  
Director, Swiss National  
Supercomputing Centre  
(SCS) at Lugano


Dr. Robert G. Voigt  
Senior Member of  
Technical Staff, Intel  
Institute, US DOE

Important Dates

January 7th, 2017  
Tutorial Proposal Submissions  
Poster Submission  
February 1st, 2017  
KAUST Inn Hotel Reservations  
Application for Student Travel Awards  
March 1st, 2017  
Registration Deadline

# Advanced Support for Saudi Aramco: Trillion Node Reservoir Simulation

- TeraPOWERS – new trillion node reservoir simulation to model oil migration problems in the Kingdom in a fraction of the time of previous run
- Shaheen II serves as the ONLY platform in the Kingdom for TeraPOWERS for developing capability and performing large scale production runs
- *" We simulated an oil migration problem in the Kingdom from the source rock to the trap with millions of years of history in 10 hours using 1 trillion active computational cells", Ali Dogru*
- *" We could not have achieved this incredible milestone without the expertise and resources from KAUST, which provided superb support," Larry Fung*



The EXPEC Advanced Research Center (EXPEC ARC) TeraPOWERS Technology Team, under the leadership of Saudi Aramco fellow Ali Dogru, achieved a major breakthrough with the industry's first trillion cell reservoir simulation run on October 28, 2016

**Aramco news, Dhahran, November 23, 2016**

- **Multiple Top500 Systems**




- Dammam-7: #23 (seismic analysis)
- Ghawar-1: #33 (reservoir simulation)
- Unizah-2 (#188), Makman-2 (#393), Makman-3 (#369)

- **Relevant Division**

- Upstream (reservoir, seismic)
- Downstream
- R&D

## Gravity Separator

Largest Scale Simulation of Engineering Code (~200,000 Cores)



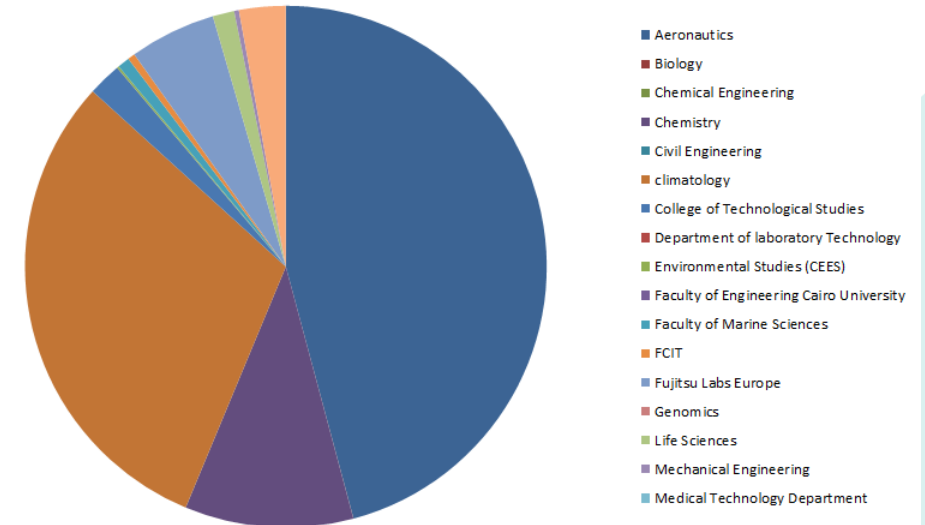
# Saudi Universities

- **KAU (King Abdulaziz University)**

- Aziz Supercomputer (2015)
  - Fujitsu Primergy CX400
  - 230TF Peak Performance (11,904 cores in 496 nodes)
- Major Research Area
  - Aeronautics, climate modeling, genomics

- **KFUPM (King Fahd University of Petroleum & Minerals)**

- Alfahd Cluster (2020)
  - Huawei Fusion / Dell PowerEdge Servers
  - 153TF Peak Performance (3,944 cores in 132 nodes)
- Major Research Area
  - Geoscience, computational chemistry, material science



- **KACST (King Abdulaziz City for Science and Technology)**
  - Three Major Functions
    - Laboratories and Technology Development
    - Support for Scientific Research
    - Innovation Support
  - Sanam Supercomputer
    - Adtech
    - 1.1PF Peak Performance (38,400 cores in 210 nodes)
    - #52 in top500 (2012)





- **Pillars:** Three Engines for a Truly Cognitive Infrastructure

---

HPC explains the past

**QUANTUM:**  
FOR ANALYSIS OF LARGE SCALE STATISTICAL DATA

- Industrial development
- Financial modelling
- Traffic prediction

---

Quantum predicts the future

**NEUROMORPHIC:**  
FOR ANALYSIS OF ANY REAL TIME DATA

- Robotic control
- Real time visual analytics
- Real time drone fleet management
- Autonomous everything


---

Neuromorphic enables the present



**HPC:**  
FOR ANALYSIS OF LARGE SCALE MULTI-TYPE DATA ON DEMAND (BUT NOT REAL TIME)

- Weather analysis
- Macroeconomics analysis
- Geographic modelling




**KAUST**  
SUPERCOMPUTING  
CORE LAB

# United Arab Emirates











# Computing Resource

- **G42 (Group 42) Company**
  - Artemis Supercomputer
    - 81 NVIDIA DGX-2
    - 7.26 PF Rmax (#90 in top500)
  - Condor Galaxy 1
    - 64 Cerebras CS-2
    - 4 EF (FP16) from 56 Million AI Cores
    - Plan to deploy two more systems CG systems
  - G42 Cloud maintains several other systems




## Condor Galaxy 1 AI Supercomputer

 <b>64</b> CS-2 nodes	 <b>54 million</b> AI cores	 <b>4 exaFLOPS</b> AI compute at FP16	 <b>82 TB</b> parameter memory	 <b>388 Tbps</b> internal bandwidth	 <b>72,704</b> AMD EPYC™ cores	 <b>10 days</b> to first training run
-----------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------





- **NYU AD (New York University Abu Dhabi)**
  - Jubail Supercomputer
    - Dell PowerEdge
    - 1.5 PF (33,000 cores, 300 nodes)
  - Key Research Areas
    - Astrophysics: Center for Space Science and Center for Astro Physics and Particle Physics
    - Chemistry (Simulation, Modeling & NMR): including the Center for Smart Engineering Materials
    - Computer Science: NLP & Machine Learning
- **MBZUAI (Mohamed bin Zayed University of Artificial Intelligence)**
  - Campus Supercomputing Center
    - 25 DGX-2, 18 DGX, 64 GPU nodes (4 A100), 32 CPU nodes



# Thank you!

[jysoo.lee@kaust.edu.sa](mailto:jysoo.lee@kaust.edu.sa)