



جامعة الملك عبدالله
للعلوم والتقنية
King Abdullah University of
Science and Technology

Core Labs and
Research Infrastructure

How to get access to Shaheen ?

Bilel Hadri

bilel.hadri@kaust.edu.sa

Computational Scientist

KAUST Supercomputing Lab

[@mnoukhiya](#) [@KAUST_HPC](#)



جامعة الملك عبدالله
للعلوم والتقنية
King Abdullah University of
Science and Technology

SHAHEEN
SUPERCOMPUTING LABORATORY



Why do I need an account on Shaheen?



What can we achieve with a supercomputer?

<https://www.youtube.com/watch?v=TGSRvV9u32M>



Impact of Shaheen2

- **Shaheen2 Usage**

- Has supported 617 projects of 218 distinct PIs
- Has provided over 6.8 billion CPU core hours
- Has produced over 1000 publications, three world records and patents
- Several best papers awards at major HPC conferences
- Economic impact: Saving \$\$\$.



KAUST Ph.D. graduate wins best paper award at prestigious Euro-Par 2020



<https://arab.news/22tph>

- **Shaheen2 User**

- 26 Organizations in the Kingdom
- Over 1400 Users


Impact of Shaheen2



<https://www.hpcwire.com/off-the-wire/supercomputing-bridges-saudi-arabia-from-coast-to-coast/>

<https://www.kaust.edu.sa/en/news/forecasting-solar-radiation-in-the-kingdom>

KAUST and Lean Business Services sign collaboration to advance digital healthcare in Saudi Arabia



<https://www.kaust.edu.sa/en/news/kaust-agreement-advances-healthcare-data-services>



KAUST team finalist for the Gordon Bell Prize

Paper coauthored by three KAUST PIs, David Keyes, Marc Genton and Ying Sun, as well as the 2021 Turing Award laureate, Jack Dongarra, has been nominated for the **2022 Gordon Bell Prize**, a prestigious award for computational sciences.



<https://www.hpcwire.com/2022/09/07/inside-the-gordon-bell-prize-finalist-projects/>

“For our exploratory science runs, and to demonstrate the acceptable accuracy of our algorithmic variations on Cholesky factorization and further manipulation of massive covariance matrices, we used Shaheen-2 at KAUST,” [...] Fugaku has 158,976 nodes, about 25 times more than Shaheen-2, and each node has 48 cores, 1.5 times more than a Shaheen-2 node. However, each Fugaku node is equipped with only 32GB of memory, one-quarter as much as Shaheen-2’s 128GB per node, thus only one-sixth as much per core, which required us to make software adaptations.”

Geostatistics get a boost from Shaheen-2 and Fugaku

The exascale-enabled research only constitutes half the list. Another finalist paper — **Reshaping Geostatistical Modeling and Prediction for Extreme-Scale Environmental Applications** — used Shaheen-2 as well as Fugaku.




The Shaheen-2 supercomputer.


Abstract: We extend the capability of space-time geostatistical modeling using algebraic approximations, illustrating application-expected accuracy worthy of double precision from majority low-precision computations and low-rank matrix approximations. We exploit the mathematical structure of the dense covariance matrix whose inverse action and determinant are repeatedly required in Gaussian log-likelihood optimization. Geostatistics augments first-principles modeling approaches for the prediction of environmental phenomena given the availability of measurements at a large number of locations; however, traditional Cholesky-based approaches grow cubically in complexity, gating practical extension to continental and global datasets now available. We combine the linear algebraic contributions of mixed-precision and low-rank computations within a tile-based Cholesky solver with on-demand casting of precisions and dynamic runtime support from PaRSEC to orchestrate tasks and data movement. Our adaptive approach scales on various systems and leverages the Fujitsu A64FX nodes of Fugaku to achieve upto 12X performance speedup against the highly optimized dense Cholesky implementation.

Per the SC22 schedule, this team includes researchers from KAUST, ORNL, and the University of Tennessee. Perhaps notably, the team also includes Jack Dongarra, one of SC22’s keynote speakers.

Usage by Science Field



% Usage in core hours



% of number of projects

In Kingdom PIs on Shaheen Supercomputer

Academy

القطاع الأكاديمي

جامعة الجوف
Jouf University



جامعة حفر الباطن
University of Hafr Al Batin



Princess Nourah Bint
Abdulrahman University



KFU
جامعة الملك فيصل
KING FAISAL UNIVERSITY



جامعة
الملك سعود
King Saud University



كلية الجبيل الجامعية
JUBAIL UNIVERSITY COLLEGE



جامعة الفيصل
Alfaisal University



جامعة الإمام عبد الرحمن بن فيصل
IMAM ABULRAHMAN BIN FAISAL UNIVERSITY



جامعة الطائف
TAIF UNIVERSITY TU

Agency and Industry

القطاع الحكومي و الصناعي

سابك
SABIC



الهيئة الوطنية للأمن السيبراني
National Cybersecurity Authority

المركز الوطني للأرصاد
National Center for Meteorology
الهيئة العامة للأرصاد الجوية



مدينة الملك عبد الله للطاقة
الذرية والمتجددة K.A-CARE



أرامكو السعودية
saudi aramco



لين
Lean



مدينة الملك عبدالعزيز
للعلوم والتقنية KACST



مستشفى الملك فيصل التخصصي ومركز الأبحاث
King Fahd Specialist Hospital & Research Centre
مركز أبحاث

UNITIX®



الشركة الكويتية لتفط الخليج
KUWAIT GULF OIL COMPANY

KJO

أرامكو لأممال الخليج
Aramco Gulf Operations



جامعة الملك عبدالعزيز
العلوم والتكنولوجيا
King Abdulaziz University of
Science and Technology

Core Labs and
Research Infrastructure

Who is next to join us?





جامعة الملك عبدالعزيز
العلوم والتكنولوجيا
King Abdulaziz University of
Science and Technology

Core Labs and
Research Infrastructure

How I can get access on Shaheen?





<http://hpc.kaust.edu.sa>

The screenshot shows the website interface for the SHAHEEN SUPERCOMPUTING LABORATORY. At the top left is the KAUST logo and name in Arabic and English. The main header features the 'SHAHEEN SUPERCOMPUTING LABORATORY' logo. Below this is a navigation bar with dropdown menus for 'My KSL', 'Resources', 'Documentation', 'Training', 'Services', and 'About Us'. The 'Documentation' menu is open, showing options for 'Shaheen II User Guide', 'Shaheen II Get Started Flyer', 'Account Applications', and 'Tips'. The main content area includes a 'Welcome to KAUST Supercomputing Laboratory' message and two event banners: 'HPC Saudi 2017' (High Performance Computing Saudi Arabia, March 13 to 15, 2017) and 'GPU Hackathon 2017' (featuring a group photo in front of a 'SHAHEEN' banner).



Account Applications

- Getting an account on Shaheen is a three-step process:
 1. **Submit a Project Proposal:** User must be associate with a project lead by faculty and submit a new project or to be added on an existing project.
 2. **Apply for an Individual Access Application:** This application will need to countersigned by the PI and send it with a scanned copy of the user organisation ID and passport.
 3. **OAA – Organisational Access Application** (This application for non KAUST users)
- The applications can be found in our website at <https://www.hpc.kaust.edu.sa/account-applications>
- Send all documents on via email projects@hpc.kaust.edu.sa or use the [Contact Us](#) form to submit credentials or other private information.

Type of Project Allocation

- **Development Project (Only for new code by KAUST faculty)** : To test your code on Shaheen (performance, check portability,....).
 - Can be granted up to 2M core hours.
 - Short project proposal, similar to an abstract
- **Production Project:** Production run after applications have been ported and optimized.
 - More details are needed
 - No limit of core hours, as long it is justified and approved by RCAC.
 - Only up to one year is granted by RCAC which can be renewed
- **Extension:** RCAC may grant extension of duration and compute core hours after progress report(including publications, success stories...)

**KAUST Supercomputing Laboratory (KSL)
Production Project Proposal (PPP)**

Project Title	
Principal Investigator (PI)	
PI Signature	
Date of Proposal	
CRG, CCF, KACST, industrially and/or governmentally funding (award details with funds and duration)	

	<input type="checkbox"/> New	<input type="checkbox"/> Extension	Project ID: k_
System	<input type="checkbox"/> Shaheen II	<input type="checkbox"/> Ibox	<input type="checkbox"/> Nesper
Core Hours Requested	Shaheen II : _____	Ibox : _____	Nesper : _____
Storage TB Requested	Shaheen II : _____	Ibox : _____	

Available Systems:

- Shaheen II Supercomputer:** 36-cabinets Cray XC40 system, comprising 6174 nodes, each with 32 Haswell cores and 128GB of memory, for a total of 197,568 cores along with 17.4 PB of Luster storage with a maximum data transfer bandwidth of 0.5 TB/s
- Ibox Cluster:**
 - Heterogeneous cluster with a total of 842 nodes of Intel, AMD and GPUs
 - <https://www.hpc.kaust.edu.sa/ibex/computing>
- Nesper Cluster:**
 - 19 nodes composed of two Skylake Intel Xeon(R) Gold 6138 CPU 2.0 GHz with 40 physical cores and 12TB of local disk per node (37 nodes with 192 GB of memory and 2 nodes are equipped with 768 GB of memory).
 - One node composed of 2 Intel Xeon(R) Haswell 16 cores 2.3 GHz with 256 GB of RAM and 16 Nvidia K80 Tesla cards.
 - <https://www.hpc.kaust.edu.sa/neser>

Submission

Please send a scanned copy of the completed Project Proposal to:
Projects@hpc.kaust.edu.sa

Definitions:

- Production Project** - A production project requires that applications have been ported and tuned, and that performance assessments have been completed. Production Projects will be allocated significant computing resources.

Terms and Conditions regarding Research Publications

Whenever the results of research conducted on the HPC systems at KAUST are published, or the research involved personnel from KAUST Supercomputing Laboratory (KSL), Principal Investigators (PIs) are required to acknowledge the usage of the HPC systems at KAUST and/or the involvement of KSL personnel in their research in their publications. For example, the following statement could be used: "For computer time, this research used the resources of the Supercomputing Laboratory at King Abdullah University of Science & Technology (KAUST) in Thuwal, Saudi Arabia."

Principal Investigator (PI):

Name:	
Email:	
Tel:	
Organisation:	
Department:	
Organisation Address:	

Additional Investigators

1	Name:	
	Email:	
	Tel:	
	Organisation:	
	Org Address:	
2	Name:	
	Email:	
	Tel:	
	Organisation:	
	Org Address:	

Collaborators (External):

1	Name:	
	Email:	
	Tel:	
	Organisation:	
	Org Address:	
2	Name:	
	Email:	
	Tel:	
	Organisation:	
	Org Address:	

Project Description:

Please describe the activities proposed, including current state of art, research work proposed, expected milestones, and deliverables, as well as a summary description in the box below, and include the scientific field of the investigation as part of the description.

Note: Citations of the scientific literature are encouraged in order to show where the proposed simulations stand with respect to the 'state of the practice' as model generality, resolution, and advantages of simulation versus experiments.

Detail your scientific project. Should be similar to long abstract.



Project Background:

Please describe the background to this project, including (i) what is the existing work in this area; (ii) what is the novelty of the proposed project; (iii) why is the proposed work significant; (iv) what is the common methodology to tackle similar problems; (v) does the proposed project follow a similar/different methodology; (vi) what are the expected outcomes of this project

Scientific Impact:

Please detail the expected scientific impact of the proposed research.

List of Publications:

Please provide a list of publications that have appeared **publicly**, or been accepted for journals, conference proceedings, posters, etc., resulting from KAUST HPC resources from any previous allocations granted. If more convenient, append to this proposal an existing list of publications, with those attributed to past KAUST HPC allocations of PIs on this proposal marked.

In Preparation and Future Publications:

To enhance the review committee's understanding of the potential impact of your computations, please list a paper or papers that you contemplate arising from this research. Please mention a tentative title, a potential venue, and a list of likely co-authors from inside and outside of KAUST. This information will not be circulated beyond the committee in any form and is understood not to constitute a part of any core-hour award agreement, given the unpredictability of the fruitful directions of research.

Please remember that all papers benefiting from KSL resources must make an attribution to KSL.



Codes & Libraries:

- Please provide the following information for each code or library that will be used.
- if needed, please include the same information for any other codes or libraries to be used in 'Additional information' at the end of this proposal, or attached on a separate sheet.

1	Name of Code/Library:	
	Ownership / Licensing:	
	URL (for Open Source codes):	
	Function:	
2	Name of Code/Library:	
	Ownership / Licensing:	
	URL (for Open Source codes):	
	Function:	

Code Readiness:

- Please provide details of code performance and scalability achieved, and note any known issues that might impact production execution.
- if possible, please provide a simple table/graph showing the required 'wall time' versus the number of cores used.
- if needed, please include the same information for any other codes or libraries to be used in 'Additional information' at the end of this proposal, or attached on a separate sheet.

**I'm using XYZ software and I need to run it 500 different parameters (explain briefly the 500 parameters for example).
10 node x 32 cores x 24 hours x 500 simulations = 3.84M core hours.**



1	Name of Code/Library:	
	Scalability:	
	Known Issues:	
2	Name of Code/Library:	
	Scalability:	
	Known Issues:	

Resource Requirements:

Compute Resource	Requirement (core hours)	Duration (in Days)
Shabreen II		

Resource Requirement Justification:

Please detail how the above requirements were calculated.

Typical problem description:

Please describe typical problem size and duration e.g. typical job will use 100 nodes for 2 hours.

Notes:

Notes:

- Jobs are submitted to the systems via the SLURM Job Scheduler that are scheduled and run as soon as possible.
- The charging is allocated in 'core-hours', i.e. one core for one hour.

Storage Requirements:

Storage Resource	Requirement (TB)	Duration (in Days)
Shahen II /project		
Ibox		

Storage Request Justification(beyond 80TB)

Please describe both short-term and long-term storage requirements, including requirements for the number of files and data volume. How long the data needs to be stored after the project is completed? Will it be moved out from Shahen parallel filesystem for subsequent analysis in other systems?

- Please describe short-term (up to 1 year) storage requirements, including requirements for the number of files and data volume and corresponding justification.
- Please describe long-term (beyond 1 year) storage requirements, including requirements for the number of files and data volume and corresponding justification.
- How long the data needs to be stored after the project is completed?
- Will it be moved out from Shahen parallel filesystem for subsequent analysis in other systems?
- Please provide your data management plan after completion of a given project.

Notes:

- Policy states that all project data stored on disk will be removed 6 months after the completion of the project.
- Policy states that all project data stored on the tape archive will be removed upon the completion of the project unless special arrangements have been requested and granted.

Other Resource Requirements:

e.g. prep time required to characterise scalability; any human/machine interaction required during the computation, etc.:

1	
2	

Consultancy Support Required

Please indicate the number of man-days and type of any support required from staff (e.g. 0.2FTE for 3 months), which can include:

- Code development
- Code porting
- Code performance tuning
- Algorithm development
- Pre- and Post-Processing code development
- Data analysis and visualisation support
- Research program development
- Project management support

Please note that KAUST may be entitled to a share in the Intellectual Property Rights to any research results produced as a result of support provided by KAUST.

Confidentiality and Legal Issues:

Please provide details of any potential confidentiality or legal issues, e.g.:

- Is the project proposal confidential? If so, how? Does it contain human data?
- Is the data confidential? If so, how?
- Are any other aspects of the project confidential? If so, how?
- If the project is successful, could it be the subject of publicity?
- Do any third parties have ownership of any codes or data being used?

Other information in Support of the Proposal:

Please include any other considerations you feel would support of this proposal, e.g.:

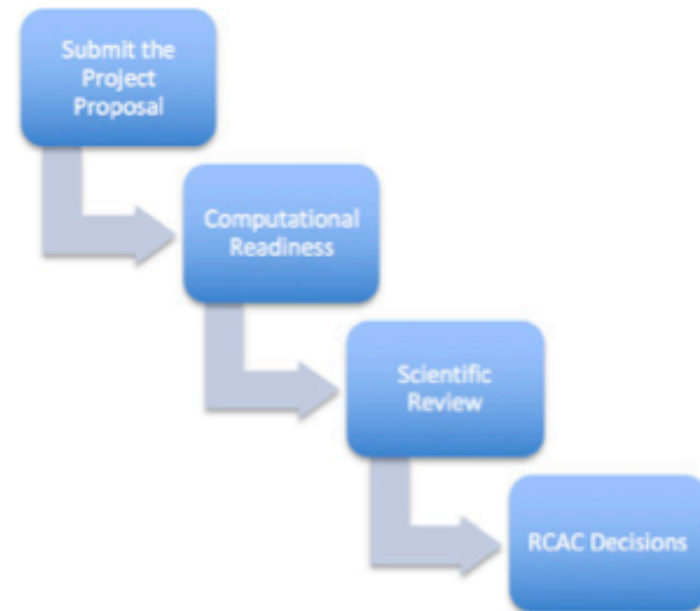
- Would the proposed project have any social impact?
- Would the proposed project have potential for generating good publicity for KAUST and/or partner organisations?
- Would the proposed project develop any useful tools that might be shared with others inside or outside of KAUST?

Potential reviewers:

Proposed production projects undergo a two-stage review process. The first stage, a Computational Readiness Review, is performed internally by KSL staff members possessing expertise in the application code proposed for execution, on the basis of applicant's documentation under "Code Readiness" above or on the basis of earlier known uses of the code. The second stage is a Scientific Readiness Review, performed by scientific peers in the discipline of the proposed project, at KAUST or externally. The selection of reviewers is the responsibility of the RCAC and KSL staff; however, suggestions for

Computing Resource Allocation

- Project proposals reviewed monthly by RCAC (Resource Computing Allocation Committee)
- Three-step process
 - **Computational Readiness Review:** Performed by KSL. Justification for core hour, portability, scalability, impact on execution
 - **Scientific Review:** Performed by scientific peers in the discipline of the proposed project
 - **RCAC** final review and recommendation



Guidance for proposals

- Justify the resource requirements
- Review your project with your PI
- Feel free to contact us before submitting it.
- Projects Request send it to projects@hpc.kaust.edu.sa
- Don't forget to acknowledge KAUST Supercomputing in your future publications.
 - **“For computer time, this research used the resources of the Supercomputing Laboratory at King Abdullah University of Science & Technology (KAUST) in Thuwal, Saudi Arabia.”**



bilel.hadri@kaust.edu.sa
[@mnoukhiya](https://twitter.com/mnoukhiya) [@KAUST_HPC](https://twitter.com/KAUST_HPC)

Thank You ! شكراً !





TITLE

SUBTITLE

Lorem Ipsum is simply dummy text of the printing and typesetting industry. Lorem Ipsum has been the industry's standard dummy text ever since the 1500s, when an unknown printer took a galley of type and scrambled it to make a type specimen book. It has survived not only five centuries.

