

European flagship  
supercomputer

**LUMI**

**LUMI Update**

**Kimmo Koski, Argonne September  
2024**



[Kimmo.koski@csc.fi](mailto:Kimmo.koski@csc.fi)

[www.lumi-supercomputer.eu](http://www.lumi-supercomputer.eu) #lumisupercomputer #lumieurohpc




**EuroHPC**  
Joint Undertaking

**LUMI**


A unique collaboration between eleven LUMI consortium countries and the EU to build and operate a world-class supercomputer.

LUMI research infrastructure provides a high-quality, cost-efficient and environmentally sustainable HPC ecosystem based on true European collaboration.





# LUMI



**EUROPEAN SUPERCOMPUTING**

How your company can accelerate innovations using LUMI supercomputer

Thursday 07:00 (CET) | 9:00 (EST)  
 Tampere, Finland | Register now  
 Register | 2024 | Tampere

Companies in Tampere region

**EURO<sup>2</sup>**

Accelerating innovations with Natural Language Processing

**Supercomputing Innovation Summit 2023**

4-5 October 2023, Nagasaki

How you can accelerate innovations using LUMI supercomputer

benefits: **Free access**  
 Free access to LUMI supercomputer resources  
 Free access to LUMI supercomputer resources  
 Free access to LUMI supercomputer resources

Supercomputing for chemistry and materials science companies

Register now

**DESTINATION EARTH**

**EuroHPC JU**

Updated 13.3.2024

**Users in Europe**

Information for all users in Europe how to apply for the resources allocated by the...

[Read more](#)

**The Netherlands**

Updated 13.3.2024

**Users in the Netherlands**

Information for users in the Netherlands. More information coming soon...

[Read more](#)

**Belgium**

Updated 6.3.2024

**Users in Belgium**

You can apply for the national LUMI resources either as an academic user affiliated with...

[Read more](#)

**Czech Republic**

Updated 12.2.2024

**Users in the Czech Republic**

Calls for LUMI Czech Republic allocations are issued by e-INFRA.CZ, and...

[Read more](#)

**Denmark**

Updated 11.3.2024

**Users in Denmark**

You can apply for the national LUMI resources either as an academic user affiliated with...

**Estonia**

Updated 29.11.2021

**Users in Estonia**

Information for users in Estonia on how to access LUMI. You...

**Finland**

Updated 19.1.2024

**Users in Finland**

You can apply for the national LUMI resources either as an academic user affiliated with...

[Read more](#)

**Iceland**

Updated 30.1.2024

**Users in Iceland**

You can apply for the national LUMI resources either as an academic user affiliated with...

[Read more](#)

**Norway**

Updated 30.1.2024

**Users in Norway**

You can apply for the national LUMI resources either as an academic user affiliated with...

[Read more](#)

**Poland**

Updated 8.2.2024

**Users in Poland**

You can apply for the national LUMI resources either as an academic user affiliated with...

[Read more](#)

**Sweden**

Updated 27.3.2024

**Users in Sweden**

Calls for LUMI Sweden allocations are issued by NAIS - National Infrastructure for Supercomputing.

[Read more](#)

**Switzerland**

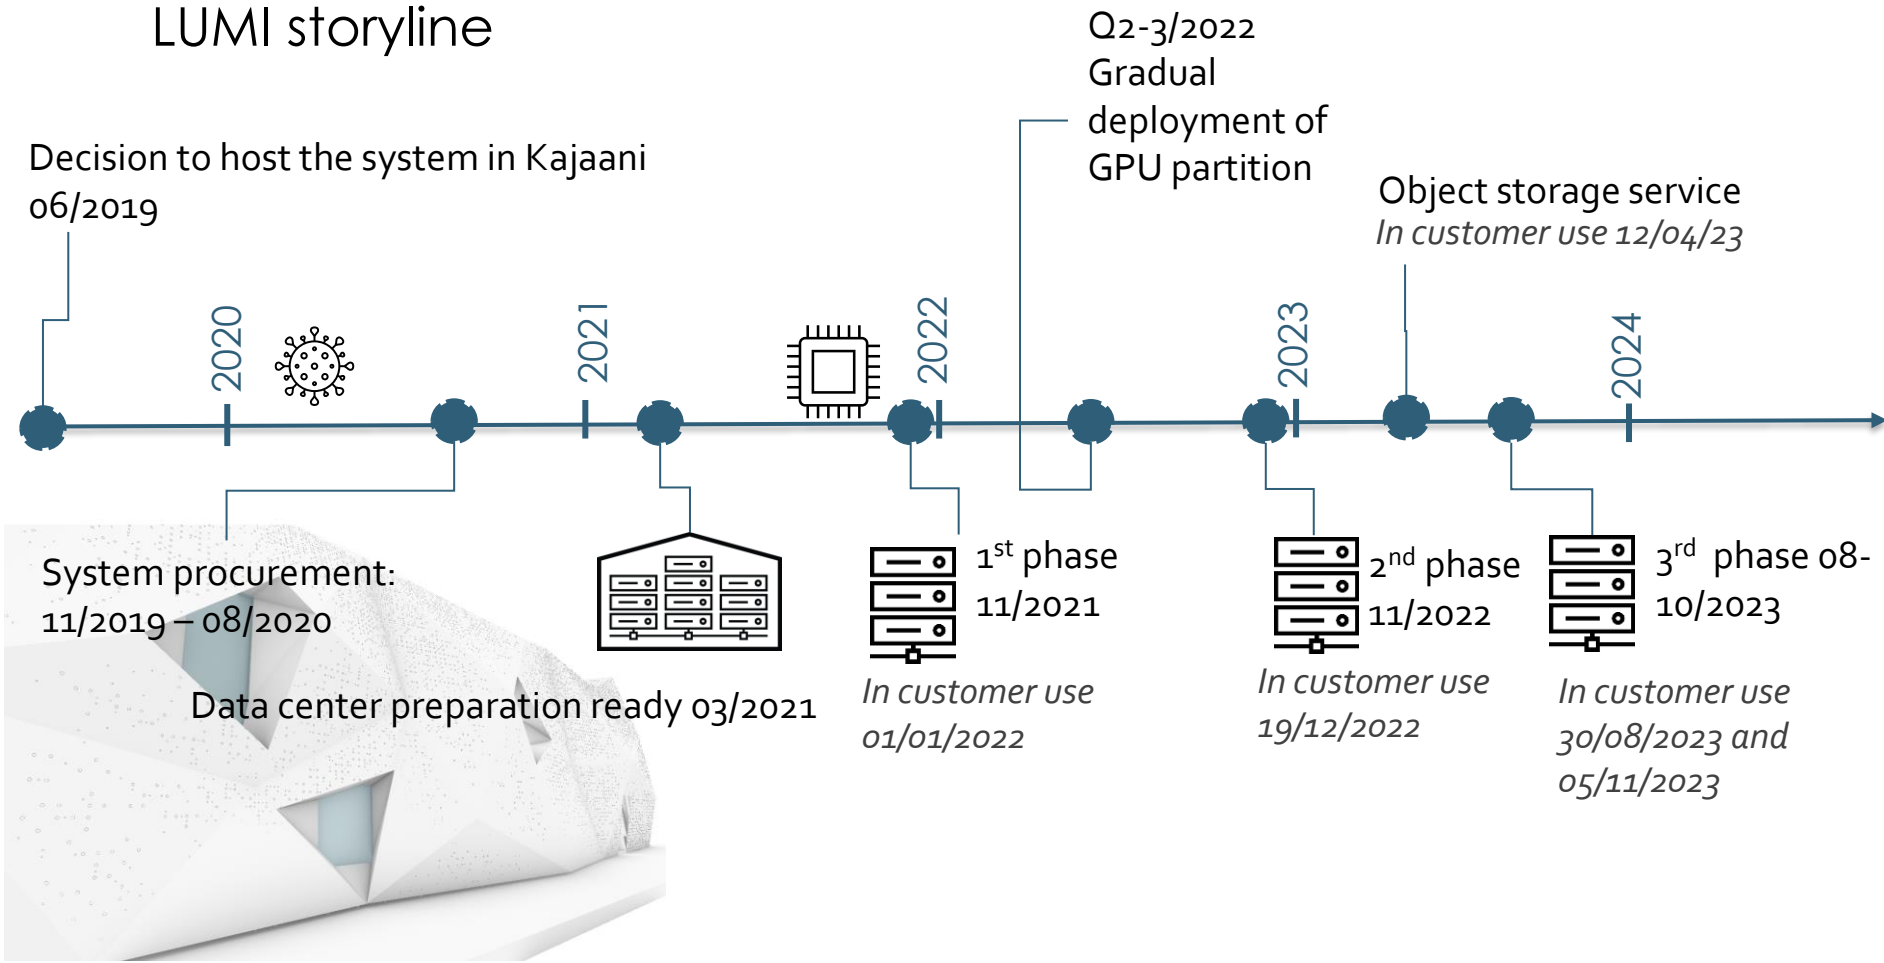
Updated 26.4.2022

**Users in Switzerland**

You can access the national LUMI resources as an academic user by applying to CSCS.

[Read more](#)

# LUMI storyline



# LUMI: one of the fastest supercomputers in the world

- LUMI is an **HPE Cray EX** supercomputer manufactured by **Hewlett Packard Enterprise**
- Sustained performance of **380 petaflop/s** makes the system one of the world's fastest
  - #5 in Top500 (#1 in Europe)
  - #3 HPCG, #3 HPL-MxP



1 system

**380**


**Pflop/s**

Sustained performance

Computing power  
equivalent to

**1 500 000**

Modern laptop computers



Size of two tennis  
courts

Modern platform for

High-performance  
computing,  
Artificial intelligence,  
Data analytics

Based on GPU technology

# Modern architecture

**LUMI-C:**  
x86 Partition  
Supplementary CPU partition:  
over **262,000**  
AMD EPYC CPU cores.



**LUMI-K:**  
Container Cloud Service



**LUMI-O:**  
Object Storage Service  
**30 PB**  
encrypted object storage  
(Ceph) for storing, sharing  
and staging data.



**LUMI-Q:**  
Quantum Computing



## High-speed interconnect

Possibility for combining  
different resources within  
a single run. HPE  
Slingshot technology.



**LUMI-G:**  
GPU Partition  
Sustained performance  
**380**  
Pflop/s powered by AMD  
Radeon Instinct™ MI250X GPUs.



**LUMI-D:**  
Data Analytics Partition  
Interactive partition with  
**32 TB**  
of memory and graphics GPUs for  
data analytics and visualization.



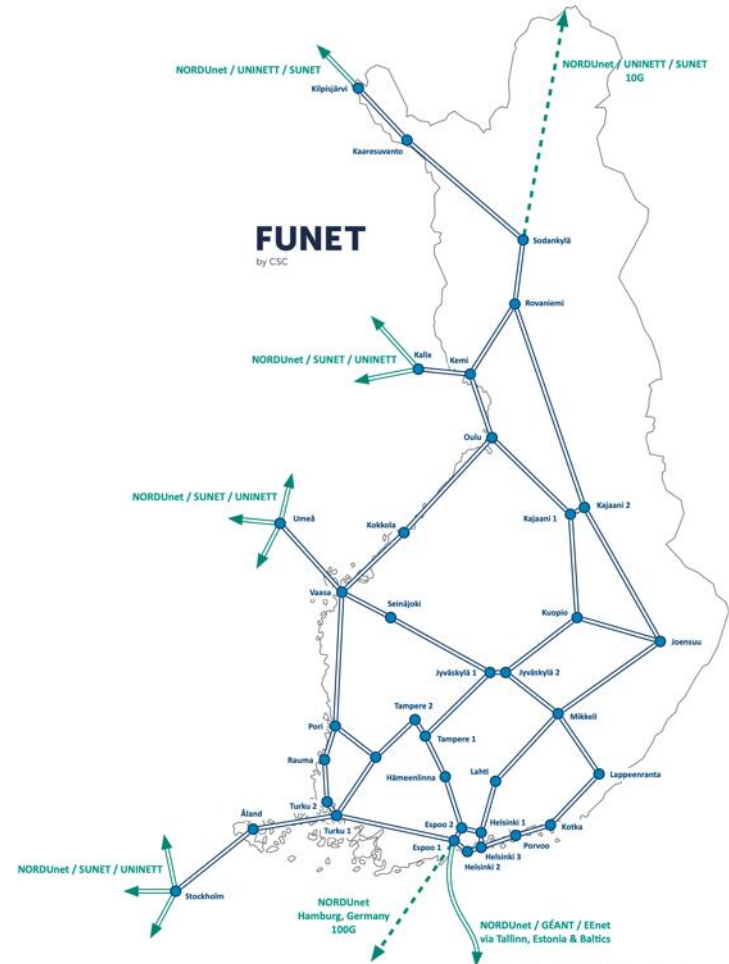
**LUMI-F:**  
Accelerated Storage  
**10 PB**  
Flash-based storage layer with  
extreme I/O bandwidth of  
2 TB/s and IOPS capability.



**LUMI-P:**  
Lustre Storage  
**80 PB**  
parallel file system.


The reliable high-speed data communications networks of the data center are designed specifically for HPC

- LUMI research infrastructure is a **direct part of the Nordic backbone**
- Scalability for multi-terabit transmission needs already today, and readiness for future transmission technologies
- The Funet 2020 network supports the EuroHPC installation perfectly **without a need for additional investments**
- The next-generation NORDUnet connects the Kajaani LUMI site to GÉANT, ensuring European-wide availability of any HPC resources installed in Kajaani





The convergence of High-performance Computing, Artificial Intelligence and Data Analytics will be key for solving the great scientific and societal challenges.

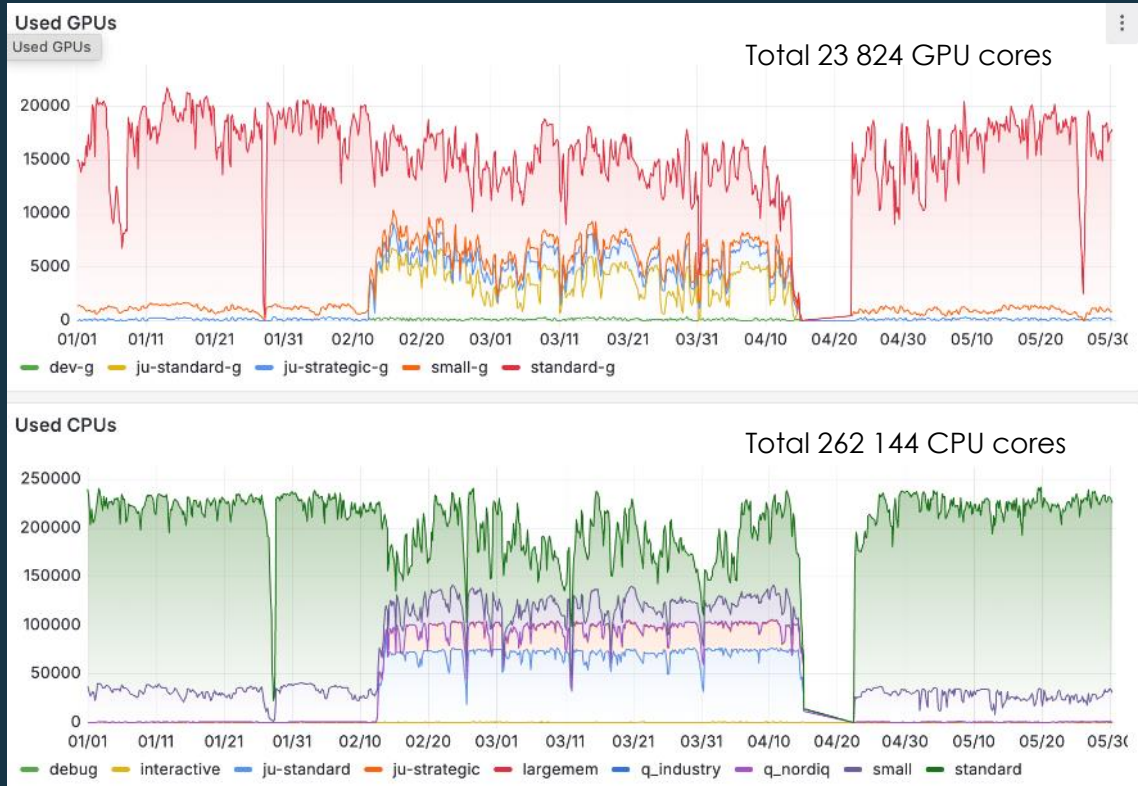


Enabling top research and scientific breakthroughs

# Key Numbers for LUMI as of May 30 Production during 2024


3969 Users consumed:

- 57.57 M GPUh
- 1.35 M CPUh
- 800 836 QPU seconds



# Largest GPU projects 2023 by usage

- The continuum limit of lattice QCD for high-precision tests of the standard model of particle physics from domain wall fermions: (JU, 5M GPUh)
- Deep Generative Language Modeling for Finnish (FI, 5M GPUh)
- Destination Earth / DestinE climate DT and associated use cases (JU, 4M GPUh)
- SSSI: Studying Small-Scale dynamo action in the Sun (FI, 3,3M GPUh)
- Open source large language models (FI, 2,4M GPUh)
- LUMI-G benchmarks of Grid and GPT (JU, 2,9M GPUh)
- LumiLLM - LUMI Large Language Models (FI, 2,3M GPUh)
- GRaCoS Gravitational radiation from Cosmic Strings (FI, 2,2M GPUh)





## Destination Earth key dates

**2022**

Launch of Destination Earth (DestinE) initiative.

**2024**


Demonstration of open core digital platform and the first two digital twins: weather-induced extremes and climate change adaptation.

**2027**

Completion of the core platform, the first two digital twins and integration of additional digital twins.

**2030**

A full digital replica of the Earth through the convergence of the digital twins available.



2022

2024

2027

2030



# CLIMATE CHANGE

How will living conditions change when the climate is warming?



## Life Sciences

Advanced computing plays a key role in all levels of modern medicine and health, and will have tremendous impact for personalised medicine.

Researchers are already able to rapidly identify genetic disease variants, and it will become possible to identify diseases that are caused by combinations of variants, with treatments and drugs tailored both to the individual patient and the exact state of the disease.



## TIME CRITICAL MODELLING

Fast-track for urgent computing needs in time- and mission-critical simulations, e.g., related to national or EU threat or other major crisis, e.g., pandemics.



# ARTIFICIAL INTELLIGENCE

LUMI is a one of the worlds leading research platforms for AI.

LUMI enables the convergence of high-performance computing, artificial intelligence, and high-performance data analytics.



LUMI



#### Poro 34B Model Card

Poro is a 34B parameter decoder-only transformer pretrained on Finnish, English and code. It is being trained on 1 trillion tokens. Poro is a fully open source model and is made available under the Apache 2.0 License.

Poro was created in a collaboration between [SiloGen](#) from [Silo AI](#), the [TurkuNLP group](#) of the University of Turku, and [High Performance Language Technologies \(HPLT\)](#). Training was conducted on the [LUMI supercomputer](#), using compute resources generously provided by [CSC](#) - IT Center for Science, Finland.

# Examples of LUMI use cases

# DestinE flagship initiative: Climate Digital Twin

Digital Modelling, Climate Research | Led by the European Centre for Medium-Range Weather Forecasts

LUMI

## Project description

- **The EU's Destination Earth initiative (DestinE) aims to create highly-accurate digital models (twins) of different aspects of the Earth system on a global scale.** The initiative will contribute to achieving the objectives of both green and digital transitions as part of the European Commission's Green Deal and Digital Strategy.
- The LUMI supercomputer functions as one of the platforms to run the Climate Change Adaptation Digital Twin (Climate DT), a high-priority digital twin of DestinE. **The Climate DT will be a state-of-the-art climate simulation and information system that can be used to support decision-making on the impacts of climate change and different adaptation strategies at local and regional levels over multiple decades.**

## Objectives & Outcomes

- As necessary simulations require extreme computing power and data handling capacities, **the LUMI supercomputer is a key element for DestinE's Climate DT to succeed.** LUMI is well positioned to host digital twins due to its computing power and data processing capability, combined with versatile value-add features.




Image: European Commission.

# Largest Finnish language model thus far

Natural Language Processing | University of Turku

LUMI

## Project description

- A Finnish research team on Natural Language Processing, TurkuNLP, created a Generative Pre-trained Transformer 3 (GPT-3) level language model with 13 billion parameters based entirely on Finnish. **This is the largest Finnish language model ever.**

## Objectives & Outcomes

- Massive language models are needed because they lay the foundation for next-generation AI applications. **It is vital to enable the development of language models for smaller languages that are often outside of large companies' commercial interests.**
- TurkuNLP's language model was published on the Internet for open access, enabling building on the model.
- The development of language models is based on huge data sets, which are used by deep-learning neural networks to create new language models. **The computing power of the LUMI supercomputer accelerates the creation of language models enormously.**



*TurkuNLP researchers. Image: CSC – IT Center for Science.*

# Automating sleep research with massive datasets

Sleep research, Neural Networks | Aarhus University, Denmark

LUMI

## Project description

- Two graduate students from the Aarhus University aimed to automate a laborious task of determining sleep stages of every individual sleeper every 30 seconds throughout a night's recording **with a sustainable and scalable solution that could be used by others in the future.**

## Objectives & Outcomes

- To create an analysis model that could replicate specialised sleep experts' work, **The researchers trained neural networks to perform sleep scoring based on 20,000 nights' worth of data.**
- The data-heavy project was made possible by the LUMI supercomputer, which allowed to train the models with such a large dataset.
- Thanks to LUMI-specific tools developed by the project, other researchers can easily load more data into the neural network training in the future. With more data, neural networks will become increasingly skilled at automatically reading sleep stages correctly, leading **to better sleep scoring models and more accurate interpretations.**



Photo: Kaare Mikkelsen

# SARS-CoV-2 variant surveillance pipelines

Generating data analysis pipelines | European Molecular Biology Laboratory



## Project description

- The European COVID-19 Data Platform was launched to provide infrastructure for countries aligning strategically to combat the pandemic and drive forward scientific progress for preparedness against infectious diseases. **SARS-CoV-2 data analysis pipelines, utilized for updating the European COVID-19 data portal, are effectively deployed and tested on LUMI.**

## Objectives & Outcomes

- Big viral data is a technical challenge for the present EuroHPC supercomputers. LUMI is used to migrate and optimize the data processing pipelines that, together with efficient mirroring of relevant datasets, **will allow timely analysis of open pathogen sequence data, manage the rapid growth in SARS-CoV-2 genomes, and support the analysis of future viral variants.**
- This endeavour and computing collaboration between life science data infrastructures with the EuroHPC remains relevant task for European preparedness for infectious diseases until a computational process for the periodic update of the European COVID-19 data portal have been successfully established.



*Image: Adobe Stock.*

Read more [here](#)

# Information from under the Sun's surface

Solar Research | Aalto University & Max Planck Institute for Solar System Research

L U M I

## Project description

- Solar magnetism is responsible for generating solar activity, including solar storms. **A solar storm's effects on Earth may cause large-scale problems for electrical grids, telecommunications and positioning systems.** The mechanisms of solar magnetic fields are still very poorly understood.
- **Predictions about the space weather are currently based on monitoring the changes in the solar surface.** This means that predictions can only be made when the active region has already appeared on the Sun's surface. It can often be too late to react by then.
- Verifying Small-Scale dynamo action in the Sun (VISSI) project focused on a previously unproven theory about mechanisms that exist under the Sun's surface.

## Objectives & Outcomes

- With LUMI, a multidisciplinary research group **was able to verify the existence of a fluctuating dynamo process in the outer-most layer of the Sun's interior by using simulation models that were the most accurate of their own type so far to be created.** In a follow-up project, the group will focus on how the verified process affects the larger scale solar phenomena. This could improve predictions of space weather in the future, which in turn would give more time to prepare for possible impacts.



Image: Adobe Stock.

# Resolving ice stream formation over Greenland

Computational glaciology | ETH Zurich, Switzerland (EuroHPC Extreme Scale project)


LUMI

## Project description

- A Swiss research team in computational glaciology developed a massively GPU-parallel thermo-mechanical ice flow model in Julia, FastIce.jl, to understand what mechanisms control the spontaneous rearrangement of ice motion in ice sheets (STREAM). **The aim is to resolve ice stream formation over Greenland in 3D.**

## Objectives & Outcomes

- Greenland, Antarctica and smaller ice sheets lose most of their mass via a few corridors --or streams-- of rapidly flowing ice, **Understanding why ice flow localization occurs spontaneously in narrow streams is crucial to more accurately predict sea-level rise.**
- FastIce.jl is a new massively parallel multi-physics 3D model. **It is written in the Julia language and runs in parallel on thousands of GPUs** using scalable in-house developed iterative solvers.
- Resolving ice flow localization into streams in 3D is numerically challenging and requires enormous spatial and temporal resolution. **The computing power of the LUMI supercomputer allows to resolve multi-physics interaction within the ice at unprecedented scales.**



*Ice stream formation due to strain localization over a rough bedrock, Figure: Utkin, I. & Röss, L., ETH Zurich.*

# Quantum integration

# Why are we (as an HPC centre) interested in quantum computing?

## Quantum computers will not replace supercomputers


- Instead, the two will merge
- Important to stimulate **co-creation** between the HPC and QC communities

## Our end-users will need early-access to HPC+QC in order to make their workflows quantum-ready

- The **transition from classical CPUs/GPUs to quantum QPUs** requires fundamental rethinking of problems and algorithms
- Requires time, resources, and support






## HPC centres serve as natural starting points

- All the necessary basic services already in place:
  - User authentication and administration
  - Resource allocation and control
  - Storage, high-speed network, classical resources, ...
- **End-users!**



Quantum computer by IQM

# LUMI-Q

-  LUMI-Q consortium
-  LUMI consortium
-  LUMI-Q quantum computer
-  quantum computer
-  supercomputer

## Inclusive

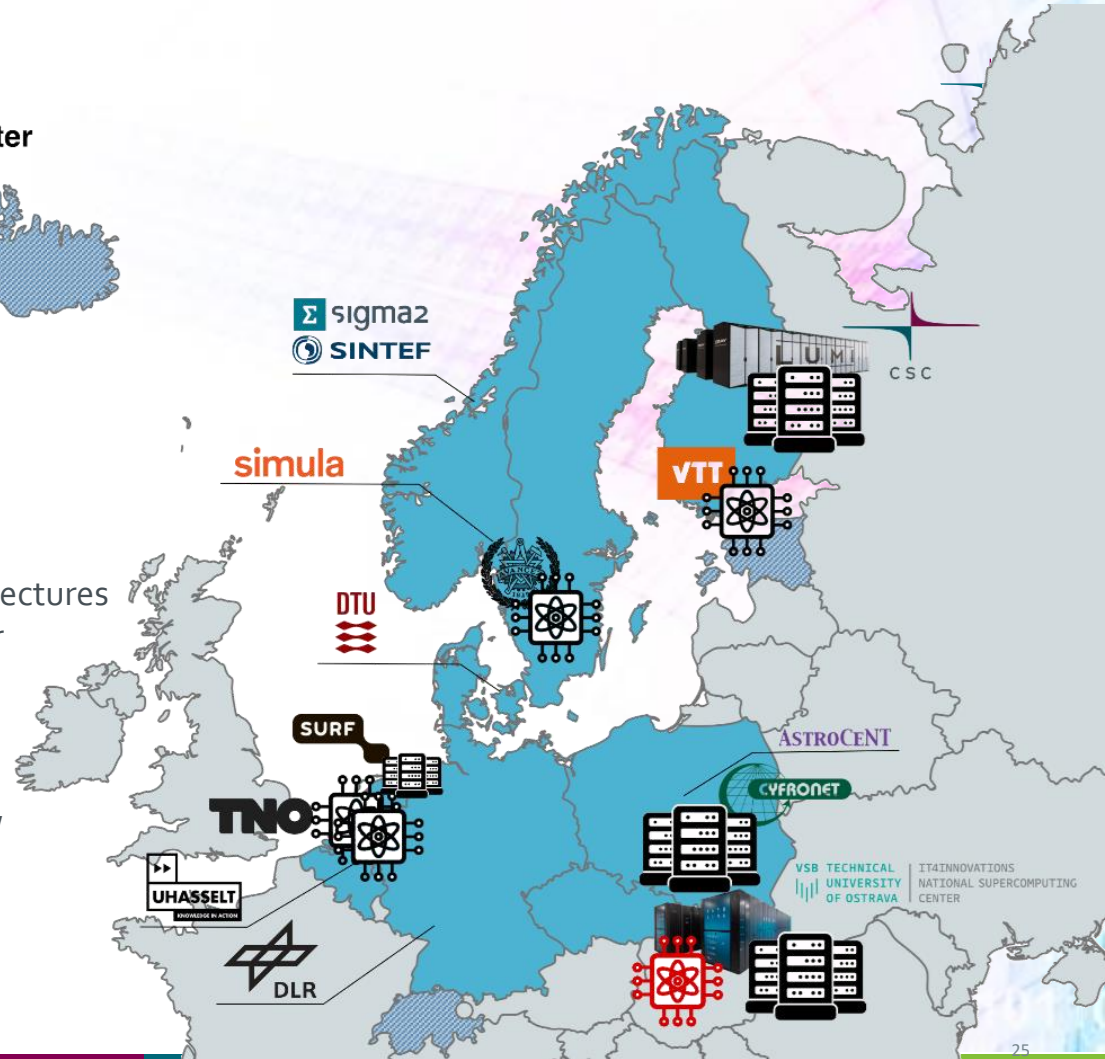
- Builds on the 10-country pan-European LUMI consortium and Germany

## Diverse

- Getting several QCs to the fingertips of researchers and developers is crucial for catalysing software development
- Different problems will fit different architectures and software stack infrastructure better

## Accessible

- By being available through several platforms distributed throughout Europe, LUMI-Q provides a familiar interface to a uniquely large user base



## Promoting the Green Deal

# A vibrant data center ecosystem in Kajaani



# Green Deal for HPC

100% hydroelectric energy up to 200 MW


Very reliable power grid: Only one outage in 40 years

100% free cooling available, PUE 1.02 (air cooled)

Waste heat reuse: effective energy price reasonable,  
negative CO<sub>2</sub> footprint: 13500 tons reduced every year

Extreme connectivity: Kajaani DC is a direct part of the Nordic backbone. 4x100 Gbit/s to GÉANT, can be easily scaled up to multi-terabit level

Elevated security standards guaranteed by ISO27001 compliancy



# World-class environmental sustainability and cost-efficiency

**LUMI**

LUMI runs on

**100%**  
hydropower

Negative  
operational

**CO<sub>2</sub>**  
**footprint**

Waste heat is used to  
heat

**hundreds**

of households  
in the city of  
Kajaani


# LUMI data center

## LUMI facility overview

- 5800m<sup>2</sup> in three floors
- 800m<sup>2</sup> whitespace for IT devices
- Power capacity 15MW at full buildout
- 14 000m<sup>2</sup> free space for future expansions

## Data center cooling

- Mechanical cooling with 3 heat pumps with 7.2MW total cooling capacity and 9MW of total heating capacity
- Free cooling, total capacity 10MW
- Air cooling installed capacity 1MW with N+1 redundancy
- 17°C inlet to CRAH units in whitespace
- Free cooling on chillers >15°C outside temperature



LUMI

# Kajaani Future Scalability

LUMI

10 MW

Option 1  
10 MW

Option 2  
10 MW

Option 3  
10 MW

Reservation  
50-100 MW



# LUMI-Next

# Continuation of LUMI Story

- Finnish government has allocated max. 250 MEUR funding for the follow-up of LUMI
- EuroHPC AI Factory concept, max. 200 + 15 MEUR
- Possible LUMI consortium countries (10) contributions
- Building a new datacenter (30 MW?) in paper machine hall next to current LUMI
- Running both LUMI and LUMI-Next in parallel for (?) years/months

