

TEAM LÖYLY



1 Meet Our Team



Figure 1: Our Team Photo.

1.1 Members



Camilla Kontto
University of Helsinki
Biology
Task force: **NAMD**

"I'm a third year biology student who specializes in genetics. The SCC is my first contact to supercomputing, essential to my study, because genomic data comes in packages so large your average computer cannot handle it at all. I have also studied mathematics, which sparked my interest in computer science. I will be focusing in NAMD at the competition, and I hope that will give me a benefit in my career someday."




Matteo Murdaca
Aalto University
Applied Physics
Task force: **HPL, MLPerf**

"As an aspiring physicist, I'm focusing my studies on semiconductor physics. This summer I had the opportunity to work in the emerging field of magnonics, designing and fabricating a magnonic convolutional encoder. Understanding the architecture of devices utilised for scientific computing and how the separation between memory and processing units can impact performance will be essential for my future career."




Hex Sun
University of Helsinki
Mathematics & Logic
Task force: **HPL, Mystery Challenges**

"I love computers, especially the beefiest ones. Computation is a core aspect of mathematics since math was born, and developments in computing technology continues to give rise to new theories in mathematics. Supercomputing will give answers to many problems we cannot even fathom today."




Santeri Toivonen
University of Helsinki
Computer Science
Task force: **HPL, Reproducibility Challenge**

Santeri is our expert in computer science. He has been working in CSC on containerization of supercomputing. His expertise in programming and optimization of software is crucial to the success of our team in benchmarking challenges.



Kim Takala
University of Helsinki
Computer Science
Task force: **ICON, MLPerf**

"I'm a second year computer science student with a background as an electrician. I am enthusiastic about getting a chance to optimize big machines on a big stage. My hope is to acquire skills especially pertaining to climate models, necessary for my future career in the field."



Sara Vakkilainen
University of Helsinki
Biology
Task force: **NAMD, ICON**

"I'm a third-year biology student with a passion for bioinformatics. I've also studied mathematics and computer science, and I hope these skills will help me in this competition. Supercomputing fascinates me because biological systems often generate massive datasets that require immense computational power to analyze. My current topic of interest is gene regulation networks, which can be computationally reconstructed to reveal underlying factors in diseases!"

1.2 About us

We are **Team Löyly**, representing both **Aalto University** and **University of Helsinki** from the beautiful country of **Finland**. We are trained by **CSC**, Finland's national supercomputing center. Team Löyly got its name from the sizzling hot temperatures of the servers that remind us of our favourite past time—the **sauna**.

Diversity is our key strength. We are proud to complement each other in a variety of categories: **gender, LGBT representation, ethnicity/nationality, neurodiversity, fields of study**, and more.

Despite being **first-time participants**, our team builds on the successes of past CSC teams, including Team Revontuli and Team Norppa, also guided by our experienced advisors, **Dr. Jussi Enkovaara** and **Mr. Leopekka Saraste**. We host bi-weekly training sessions at CSC where we utilize cutting-edge supercomputers **LUMI** and **Mahti** to familiarize us with proven methods in **industrial supercomputing**. With the continued support from **CSC expertise**, we are confident to be a **top contender** to SCC this year.

löyly NOUN
/löyly/

1. the steam that rises from the sauna stove (*kuuas*) or the heat of the sauna
2. (in the plural) sauna bath, sauna (act of taking a bath in a sauna)
3. (figuratively) heat, pressure, intensity

2 And Our Gear

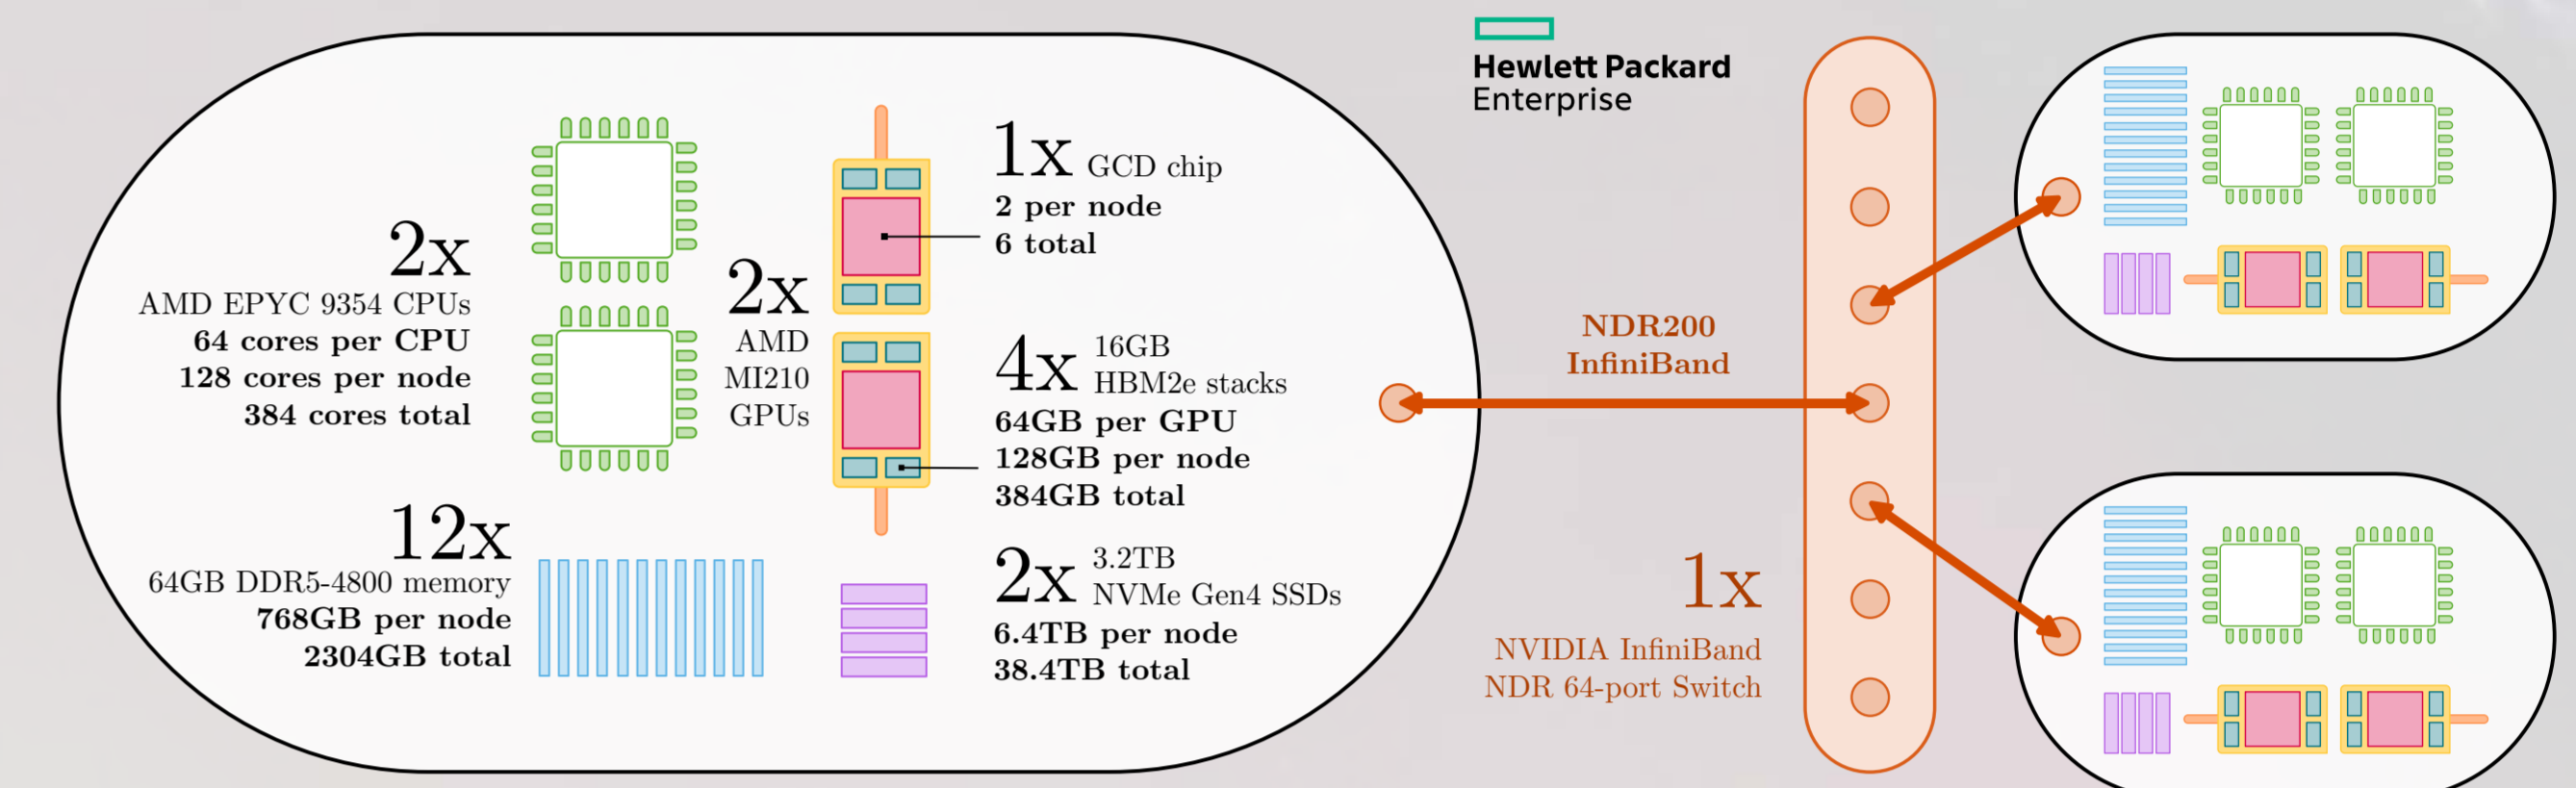


Figure 2: Composition of the Cluster.

2.1 Hardware

Our cluster consists of **3 nodes** with **identical hardware**. Each node is equipped with:

- **HPE DL385 Gen11 platform**
- **2 AMD EPYC 9354 CPUs**
- **2 AMD MI210 GPUs**
- **768GB of DDR5-4800 RAM** in total
- **6.4TB of SSD storage**

All of them are connected with a **200Gb/s InfiniBand** interconnect to achieve **low latency** and **high bandwidth**.

Our entire cluster is fine-tuned with the **power budget** in mind: under full utilization the power consumption is **4.0kW**.

The cluster is provided to us by **Helwett Packard Enterprise**. This configuration was chosen to feature a very similar hardware platform to that of **LUMI**, Europe's **largest supercomputer** ranking in **TOP500** at #5 and our training platform of choice.

2.2 Software

We are utilizing an **HPE Cray Supercomputing** software stack that is **highly optimized** and **thoroughly tested** with our setup:

- Operating System: **SUSE Enterprise Linux 15**
- Scheduler: **SLURM**
- Compilers Toolchain:
 - GNU Compiler Collection (**GCC**)
 - AMD Optimizing C/C++ Compiler (**AOCC**)
 - Cray Compiler Environment (**CCE**)
 - Linking against Cray libraries: **Cray MPICH**, **Cray LibSci...**
- Building & Deploying:
 - **Ansible** for reproducible deployment
 - **Nix** for reproducible building
 - Manual **Make/CMake** tweaking

As our system has **AMD GPUs**, our team is trained to convert **CUDA** code into **HIP** for the purpose of the mystery challenge and the reproducibility challenge.

3 Solving All Challenges

General Strategy:

- Each challenge has a small **task group** associated, who leads the challenge with **tight communication** saving time
- All team members are **cross-trained** to handle challenges and system management tasks for optimal **flexibility** and **efficiency**

3.1 MLPerf



MLPerf is a benchmark that measures how fast a system can use an existing machine learning model to make inferences for inputs.

Strategy: We will use the most optimized version of MLPerf that is designed for our AMD hardware and the ROCm stack, and **fine-tune** the **CPU and GPU bindings**. We will **profile** the code and detect **bottlenecks**, potentially improving the **communication schemes** and optimize further.

3.2 HPL



High Performance Linpack (HPL) is one of the most popular benchmark utilized for supercomputers.

Strategy: As our cluster is built with AMD components, we are using **roCHPL** to optimize performance; this will be combined with a **sapient optimization of core bindings** and inputs derived from rigorous testing.


3.3 ICON



ICON (Icosahedral Nonhydrostatic) is a Climate, Weather and Environmental modelling application based on Navier-Stokes equations.

Strategy: We are tuning the **blocking sizes** to optimize performance on the **GPU**, taking careful consideration to manage our **I/O rate** to find the perfect balance between speed and efficiency.

3.4 NAMD



NAMD is a scalable molecular dynamics application that is used for simulating large biomolecular systems with high-performance.

Strategy: We use the **GPU-enabled** version that utilizes our hardware effectively, focus on the limits that enforce **bonded and non-bonded** interactions, in which the **biology expertise** is going to come in handy. We will **visualize** our problems and make sure our simulations are running correctly.

3.5 Mystery Challenges

The instructions for a mystery benchmark and a mystery application will only be revealed when the competition starts. We are opting to have a **dedicated small task group** to focus on the challenges efficiently, and have other team members help if they are waiting for results.

3.6 Reproducibility Challenge

We are analyzing the paper **Data Flow Lifecycles for Optimizing Workflow Coordination** and running the DataLife tool to reproduce case studies featured in the paper. We will be validating the result stated in the paper against our cluster configuration.