Team RACKlette

Department of Computer Science, ETH Zürich

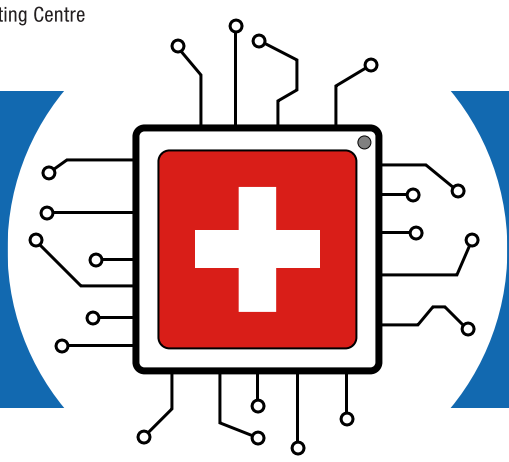

1 Team

Team **RACKlette** consists of a smaller group of enthusiastic Computer Science undergraduate and graduate students at **ETH Zürich** (Swiss Federal Institute of technology in Zurich). Our university was founded in 1854 by the Swiss Government with the mission to provide excellent education with the focus on science, technology, engineering and mathematics.

We formed team RACKlette in the beginning of 2019. During ISC19 we were able achieve the highest *LINPACK* score and placed 3rd overall. Motivated by our success we continued and have since competed at ISC19, SC19, ISC20 and now SC20. Everybody who is keen to learn more about HPC is welcome in our team. Through the presence on our website and in the introductory course of HPC we recruit new team members. Most of our knowledge is built on this course which is thought by our mentor and Professor **Torsten Hoefler**. We gratefully receive advisory support from him, **Timo Schneider** and **Hussein Harake** (*CSCS*). Our goals are: sharing gathered knowledge, learning by doing, compete to win and having fun.

The team competing at SC20 consists of six second- and third-year undergraduate Computer Science students from different backgrounds:

- **Emir İşman**, TUR – Our cluster management expert and veteran. Competed at SC19. Main focus: Cluster setup and *IO500*.
- **Simon Jacob**, DE – Our *CUDA* and math specialist. Competed at SC19. Main focus: *GROMACs* and Reproducibility Challenge.
- **Jan Kleine**, DE – Head of the team and veteran. Competed at three previous competitions. Main focus: Reproducibility Challenge and cluster setup.
- **Dario Romaniello**, CH – Novice. Main focus: *CESM* and *GROMACs*.
- **Rahul Steiger**, CH – Joined the team recently. Has achieved very good results with *LINPACK* on our cluster.
- **Simon Wachter**, CH – Also joined the team recently and has already learned a lot. Main focus: *GROMACs* and *LINPACK*.

2 Our cluster

Hardware at home

CSCS (Swiss National Supercomputing Centre) in Lugano is home to our 4-node cluster called **Emmentaler** (Swiss cheese). Each node has:

- two *AMD EPYC 7742* 64C/128T (Zen 2) CPUs
- 512GB of DDR4 memory
- four *Nvidia V100* 32GB GPUs
- two SSDs with each 2TB of storage capacity
- *Mellanox ConnectX-6 InfiniBand* HDR (200 Gb/s).

Hardware at CS20

We plan to use *AMD*, *Intel* and *Nvidia* GPU nodes. The *AMD* nodes will be interesting for the *MemXCT* application, as they offer a higher memory bandwidth, while the *Intel* nodes offer higher computational power for CPU intensive tasks. The GPU nodes will drive our benchmarks and GPU accelerated application (*GROMACs*). In general we are looking for the fastest connection available between the nodes. Furthermore we plan to use a separated storage node for data.

Software

We run *Bright Cluster Manager system 9.0* which is based on *CentOS Linux 7*. As scheduler we use *Slurm,* which has excellent integration with the *CycleCloud* autoscale feature, allowing us to boot nodes automatically when needed and shut them down after. As package manager we use *Spack*, as we are familiar wit the workflow from our own cluster. It allows us to compile all dependencies in all possible configurations we need. We found that it eases the management of having many different versions of the same library and that it greatly reduces resulting conflicts.

3 Preparation

- We assigned three students to each application. This partitioning is overlapping and redundant.
- In our weekly *Zoom* meeting we synchronize, discuss problems that arise and strategies to solve them. In a separate weekly meeting with Hussein Harake, who is a HPC expert at *CSCS* and our mentor, we address more specific problems. Hussein has helped us a lot and brought us in contact with specialists from *CSCS*, that are working with applications from SC20 professionally.
- Before we run the applications in the *Microsoft Azure* cloud, we have run and tested them on our cluster to gain experience.
- In order to keep the overview and not solving problems twice, we maintain a wiki to document the build and tuning process.

Strategy

We are developing a performance model for all applications and run *reframe* tests. Utilizing this data, we can solve the optimization problems. We want to optimize for gained points per spent cost. With such a goal it is probably not the best strategy to use the maximum resources available. However this also depends heavily on the grading of the applications. Our gathered data will enable us to make the best decision when it counts.

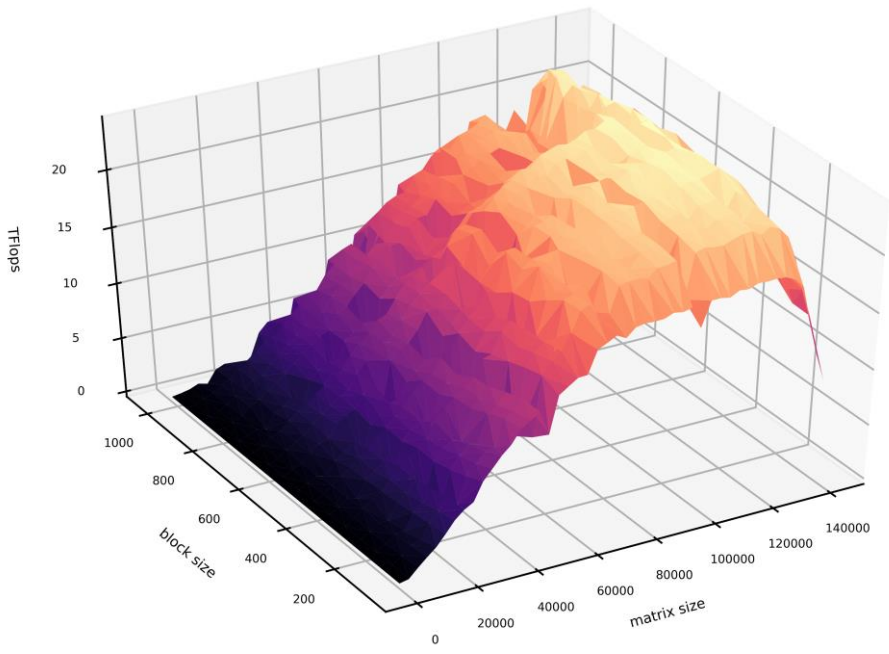

The graph displays *LINPACK* performance on one node of our cluster with different block and matrix sizes. It reached about 85% of the maximum theoretical performance.

4 We win because

- We are all highly motivated and want to learn as much as possible.
- Our own 4-node cluster at *CSCS* gives us a big advantage over other teams because it enables us to test without spending cloud budget.
- We receive great support from our mentors Hussein Harake and Torsten Hoefler and from former RACKlette members
- In the end, confidence is our strategy and good planning will ensure that we end up on top.

Contact

Website:	https://racklette.ethz.ch/
Email:	racklette@lists.inf.ethz.ch
CS Department ETH:	https://inf.ethz.ch/
Twitter:	https://twitter.com/ETHRacklette