GASERI 😎

Group for Applications and Services on Exascale Research Infrastructure

Faculty of Informatics and Digital Technologies, University of Rijeka, Rijeka, Croatia

GASERI 😎

Exascale supercomputing is almost here

A **supercomputer** is a computer with a high level of performance as compared to a general-purpose computer, also called a **high performance computer** (HPC).

The performance is measured in terms of **floating-point operations per second** (FLOPS), where modern HPCs offer 1 to 10 petaFLOPS (top 442 petaFLOPS), while near-future **exascale supercomputers** will offer more than 1 exaFLOPS.

GASERI 😎

More FLOPS means more complexity

Modern HPCs use different types of accelerators such as GPUs and FPGAs, but also specialized hardware accelerators such as tensor processing units (TPUs) and in-network/in-storage computational hardware such as **data processing units** (DPUs).

How to adapt the existing pile of scientific software for the near-future exascale supercomputing era?

GASERI 😎

Scientific software for exascale era

Our main interest is the research and development of algorithms in **scientific software** for **exascale supercomputers**.

The goal is to **design better-performing algorithms** and offer their implementations for academic and industrial use.

The specific focus in the present is the improvement of **molecular dynamics simulation algorithms** that enables efficient utilization of the **exascale HPC resources**.

GASERI 😎

GROMACS for the exascale era

Fast multipole method for general molecular dynamics simulation box types

The present design and implementation of the fast multipole method in GROMACS **only supports cubic** simulation boxes.

We are extending the method to also support approx. **30% smaller rhombic dodecahedron** simulation boxes, which would result in approx. **30% less computation time** per step required.

GASERI 😎

GROMACS for the exascale era

DPU offload of force reduction calculations in molecular dynamics simulations

The present implementation of accelerator offload in GROMACS uses CUDA or SYCL for force calculation on GPUs.

In this **NVIDIA-supported project**, our goal is to use the commercial off-the-shelf smart network cards with accelerators called **data processing units** (DPUs) for force reduction calculations.

GASERI 😎

Contact

Dr. Vedran Miletić, *Principal Investigator*

- Twitter, LinkedIn, GitHub, GitLab

Group for **A**pplications and **S**ervices on **E**xascale **R**esearch **I**nfrastructure

- LinkedIn, GitHub, GitLab