EuroHPC Summit Week 2022

#PRACEdays

PRACE-6IP WP8: Forward-looking Software Solutions

Making software ready for the European (pre-)exascale HPC landscape

Joost VandeVondele[1], Fabio Affinito[2], Alex Upton[1]
[1]ETHZ-CSCS; [2]CINECA

1 Overview of PRACE-6IP WP8

- Work package 8 (WP8) of PRACE-6IP focuses on 'Forward-looking Software Solutions' and has the objective to deliver high quality, transversal software that addresses the challenge posed by the rapidly changing HPC (pre-)exascale landscape
- WP8 has successfully initiated ten independent projects developing forward-looking software solutions, covering a wide range of scientific domains
- All ten projects have delivered a first production quality release of the software, including documentation and testing results, and have incorporated industry standard tools such as issue tracking, continuous integration, validation and verification

2 NB-LIB

- The goal of NB-LIB is to implement an API enabling performance-portable GROMACS non-bonded force calculation routines to be called as a library
- NB-LIB offers a complete MD program:

```
SimulationState simState(coordinates, velocities, forces, box, topology);
NBKernelOptions options = NBKernelOptions();
ForceCalculator forceCalculator(simState, options);
LeapFrog integrator(simState);
for (int step = 0; step < 100; step++)
{
        forceCalculator.compute(simState.coordinates(), simState.forces());
        integrator.integrate(1.0, simState.coordinates(), simState.velocities(),
simState.forces());
}
```

Target, ns/per day	NB-LIB	GROMACS
CPU-only	28.112	23.898
GPU	57.788	56.848

- NB-LIB test simulations outperform GROMACS, results for Van der Waals gas simulations obtained on Piz Daint

Project contact: Dr. Joe Jordan, KTH Center for HPC

3 PiCKeX

- The PicKeX project focuses on two important particle-in-cell codes: EPOCH, a fully relativistic, electromagnetic model and BIT1, a sophisticated PIC/Monte-Carlo model
- EPOCH heavily used in the laser-plasma community, BIT1 is prominent in the magnetic fusion community

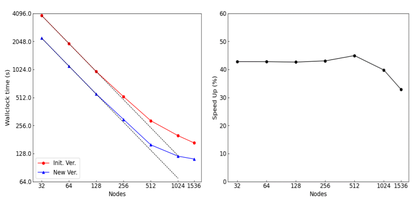

- New version of EPOCH ~30-45% faster on JUWELS than initial version for 32-1536 nodes
- In addition, new GPU version of BIT1 nearly 40% faster for particle mover algorithm on VIZ supercomputer

Project contact: Prof. Paul Gibbon, Jülich Supercomputing Centre

4 MoPHA

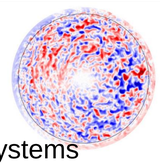

- The aim of MoPHA is to refactor and modernise the plasma simulation codes ELMFIRE, GENE, and Vlasiator to prepare for the upcoming European pre-exascale systems
- Benchmarking carried out on Puhti (similar architecture to SuperMUC-NG), Puhti-AI (similar architecture to JUWELS) and Mahti (similar architecture to Vega and MeluXina), in addition to the LUMI-C partition

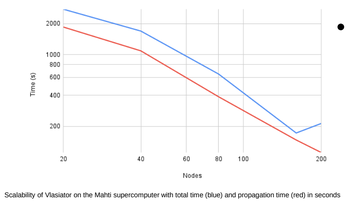

Scalability of Vlasiator on the Mahti supercomputer with total time (blue) and propagation time (red) in seconds

- Vlasiator scales well up to 160 nodes on Mahti with a 12x speedup compared to 20 nodes

Project contact: Dr. Martti Louhivuori, CSC - IT Center for Science

5 GHEX

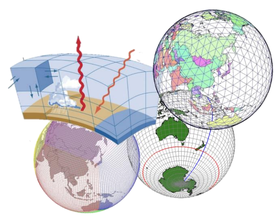

- GHEX aims to provide high-level, performance-portable communication primitives to perform halo-exchanges for massively parallel applications that use grids and meshes
- GHEX has been integrated into a number of codes, including a ~20% improvement in the performance of the atmosphere simulation Bifrost for 512 nodes on the LUMI-C partition

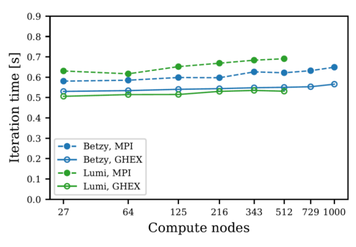

Project contact: Dr. Mauro Bianco, ETHZ-CSCS

6 QuantEx

- QuantEx aims to develop a platform for efficient simulation of quantum circuits on exascale systems with a three-layered approach, uses tensor network methods in order to simulate quantum circuits
- Benchmarks on LRZ BEAST and SuperMUC-NG systems have validated that the QuantEx framework can be run on diverse HPC architectures

	ARM A64FX (BEAST)	Intel SKL (SuperMUC-NG)	AMD ROME EPYC (BEAST)
Run-time (s) – 12 Qubits	0.14	0.061	0.030
Run-time (s) – 24 Qubits	245.4	124.92	59.06
Base frequency (MHz)	425	2300	2250
Cores/node	48	48	64

Project contact: Dr. Venkatesh Kannan, ICHEC

7 FEM/BEM domain decomposition solvers

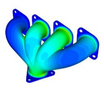

- This project aims to extend the domain decomposition library ESPRESO to enable solving of large-scale sound scattering and harmonic analysis problems

- Benchmarking performed on several European HPC systems, including Salomon, Karolina, JUWELS and JUWELS Booster

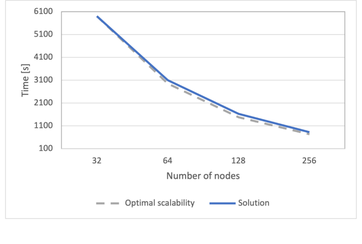

- The harmonic analysis solver obtains ~90% parallel efficiency for 256 nodes on JUWELS Booster

Project contact: Dr Michal Merta, IT4Innovations

8 LoSync

- LoSync aims to improve application scalability by removing unnecessary synchronisation and serialisation, by developing and evaluating task-aware versions of the MPI and GASPI libraries, TAMPI and TAGASPI, which are integrated with OmpSs-2 and OpenMP task-based runtimes
- Improved TAMPI and TAGASPI performance compared with MPI-only on Marenostrum4 for iterative Gauss-Siedel method for heat equation

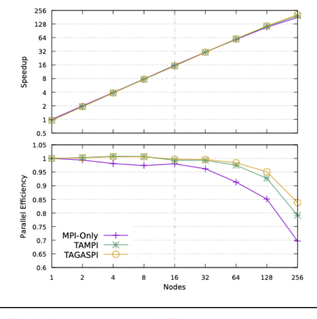

Project contact: Dr. Mark Bull, EPCC, University of Edinburgh

9 Performance Portable Linear Algebra

- This project aims to implement a distributed eigenvalue solver for dense matrices, with two main contributions:
  - Implementation of a direct eigensolver, DLA-Future (DLAF), based on MPI and HPX libraries
  - Performance improvement of the iterative eigensolver implemented in the ChASE library

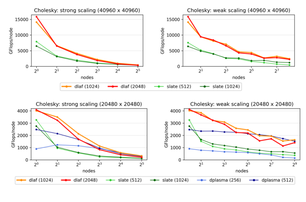

- DLAF outperforms SLATE on Marconi100 (top) and SLATE and DPLASMA on Piz Daint (bottom)

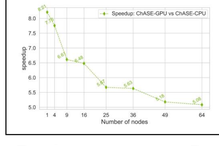

- Speedup between 5x and 8x observed for optimised GPU version of ChASE on JURECA-DC supercomputer

Project contact: Dr. Raffaele Solcà, ETHZ-CSCS

10 PARSEC

- PARSEC aims to modernise the adaptive mesh refinement (AMR) implementation of the computational fluid dynamics (CFD) codes Alya, Nek5000, and Argo to exploit (pre-)exacale machines
- Performance of GPU version of Nek5000 with turbulent pipe case shown below, scalability up to 128 nodes with good efficiency on JUWELS Booster

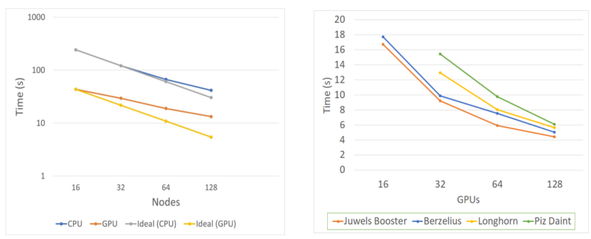

Scaling of the OpenACC version of Nek5000 on several GPU enabled machines.
(left) Results in JUWELS Booster. (right) Results including JUWELS Booster, Berzelius, Longhorn and Piz Daint

Project contact: Dr. Ricard Borrell, BSC

11 LyNcs

- LyNcs pools together software development efforts across Europe to provide research communities with the next generation of parallel libraries for solving sparse linear systems at the exascale
- It targets all levels of the software stack:
  - Sparse matrix-matrix kernels at the lowest level
  - Iterative linear solver libraries at the middle level
  - Community codes at the highest level
- Improved performance in the latest version of LIBRSB for sparse matrix computations across various HPC architectures including BEAST and SuperMUC-NG systems

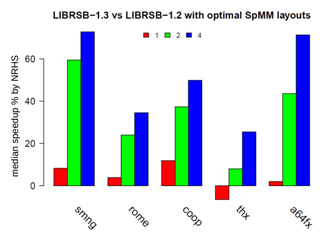

Project contact: Dr. Jacob Finkenrath, CaSToRC

12 Future directions

- The projects have achieved good scaling results across a diverse range of HPC systems in Europe, including PRACE Tier-0 and EuroHPC machines, with a view to future deployment on the upcoming European pre-Exascale systems once these are fully online
- The software solutions developed in WP8 are complementary and synergistic with the work of the Centres of Excellence (CoEs) - certain domains not represented in the CoEs, such as plasma physics, are present in WP8, whilst other projects have received strong letters of support from the CoEs
- The software engineering in WP8 supports and interacts with CoEs and other communities, e.g. via PRACE HLSTs, with a focus on transversal software, and readiness for exascale architectures

The EuroHPC Summit Week conference series receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreements 823767 (PRACE-6IP) and 824151 (HPC-GIG). The Implementation Phase of PRACE receives funding from the EU's Horizon 2020 Research and Innovation Programme (2014-2020) under grant agreement 823767

#EHPCSW