



PRACE-6IP WP8: Forward-looking Software Solutions

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

Joost VandeVondele¹, Fabio Affinito², Alex Upton¹
¹ETHZ-CSCS; ²CINECA

- WP8 of PRACE-6IP 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 in two separate competitive calls, covering a wide range of scientific domains
- All projects have delivered a first production quality release of the software, including documentation and testing results, and have incorporated industry standard tools
- Projects have achieved good scaling results across a diverse range of HPC systems in Europe, including PRACE Tier-0 and EuroHPC machines, that includes a broad range of HPC architectures
- 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 are present in WP8



PRACE-6IP WP8: Forward-looking Software Solutions

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

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:
 - `nb-lib` (includes: `nb-lib`, `nb-lib`, `nb-lib`, `nb-lib`)
 - `nb-lib` (includes: `nb-lib`, `nb-lib`, `nb-lib`, `nb-lib`)
 - `nb-lib` (includes: `nb-lib`, `nb-lib`, `nb-lib`, `nb-lib`)
- NB-LIB test simulations outperform GROMACS results for "Mdr" der "Mdr" gas simulations obtained on Piz Daini

3 PICOxX

- The PICOxX project focuses on two important particle-in-cell codes: EPOCH, a fully realistic, electromagnetic model and BITI, a sophisticated PIC-Monte-Carlo model
- EPOCH heavily used in the laser-plasma community, BITI is prominent in the magnetic fusion community
- New version of EPOCH ~30-40% faster on JUWELS than initial version for 32-1536 nodes
- In addition, new GPU version of BITI nearly 40% faster for particle mover algorithm on VZ supercomputer

4 MoPhA

- The aim of MoPhA is to refactor and modernize the plasma simulation codes ELAMFIRE, CODE, and Vlasator to prepare for the upcoming European pre-exascale systems
- Benchmarking carried out on PizDaini (similar architecture to SuperMUC-NG, PizDaini) (similar architecture to JUWELS) and Mordor (similar architecture to Vega and Maelstrom), in addition to the LUM-C partition
- Vlasator scales well up to 160 nodes on Mordor with a 12x speedup compared to 20 nodes

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 libfem for 512 nodes on the LUM-C partition

6 QuantEx

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

7 FEMBEAN domain decomposition solvers

- This project aims to extend the domain decomposition library ESPRESSO to enable solving of large-scale sound scattering and harmonic analysis problems
- Benchmarking performed on several European HPC systems, including Salomon, Marlin, JUWELS and JUWELS Booster
- The harmonic analysis solver obtains ~90% parallel efficiency for 256 nodes on JUWELS Booster

8 LoSync

- LoSync aims to improve application scalability by removing unnecessary synchronization and serialization, by developing and evaluating task-aware versions of the MPI and GAPI libraries, TAMPi and TAGAPI, which are integrated with CompS-2 and OpenMP task-based runtimes
- Improved TAMPi and TAGAPI performance compared with MPI-only on Marconi2 for sparse linear systems for iterative Gauss-Seidel method for heat equation

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 HPC libraries
 - Performance improvement of the iterative eigensolver implemented in the CHASE library
- DLAF outperforms SLATE on Marconi2 (top) and SLATE and DPLASMA on Piz Daini (bottom)
- Speedup between 5x and 8x observed for optimized GPU version of CHASE on JUWELS-OC supercomputer

10 PARSEC

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

11 Lyfco

- Lyfco 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 lowest version of LIBSIB for sparse matrix computations across various HPC architectures including BEAST and SuperMUC-NG systems

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 sparse linear systems at the exascale
- 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. the PRACE HPC-RI, with a focus on transversal software, and readiness for exascale architectures