MaX and POP

The case story of Quantum ESPRESSO

F. Affinito, C. Cavazzoni, G. Erbacci

CINECA

MAX DRIVING THE EXASCALE TRANSITION

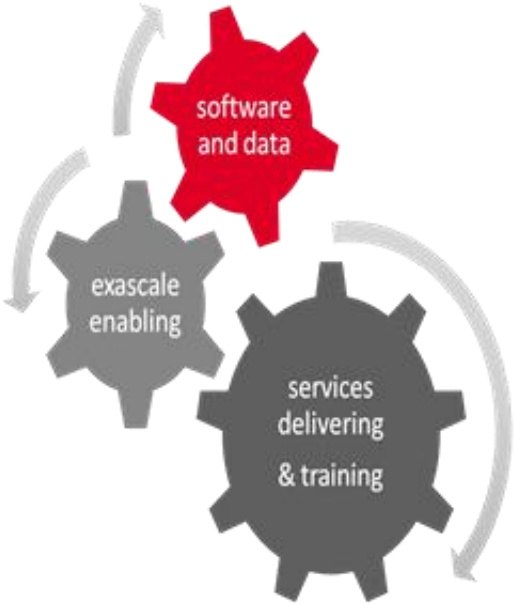

The mission:

building, managing and evolving a centre of excellence aimed to enlarge the EU leadership in materials modelling, simulations, discovery and design.

5 research institutions
5 supercomputer centers
1 global education institution
2 business partners
4 flagship codes
1 platform for workflows and data management
> 10.000 users

CINECA

QUANTUM ESPRESSO

siesta

www.flapw.de
fleur

Yambo

AiiDA
Automated Interactive Infrastructure and Database for Computational Science

Quantum ESPRESSO

Quantum ESPRESSO is an integrated suite of Open-Source computer codes for electronic-structure calculations and materials modeling at the nanoscale. It is based on density-functional theory, plane waves, and pseudopotentials.

It is adopted and developed by a large community of users spread worldwide.

Quantum ESPRESSO implements a coarse grained MPI parallelization on top of which OpenMP threads are used.

Quantum ESPRESSO showed good scalability on a wide range of supercomputing machines on different architectures

www.quantumespresso.org

Quantum ESPRESSO algorithmic schema (SCF)

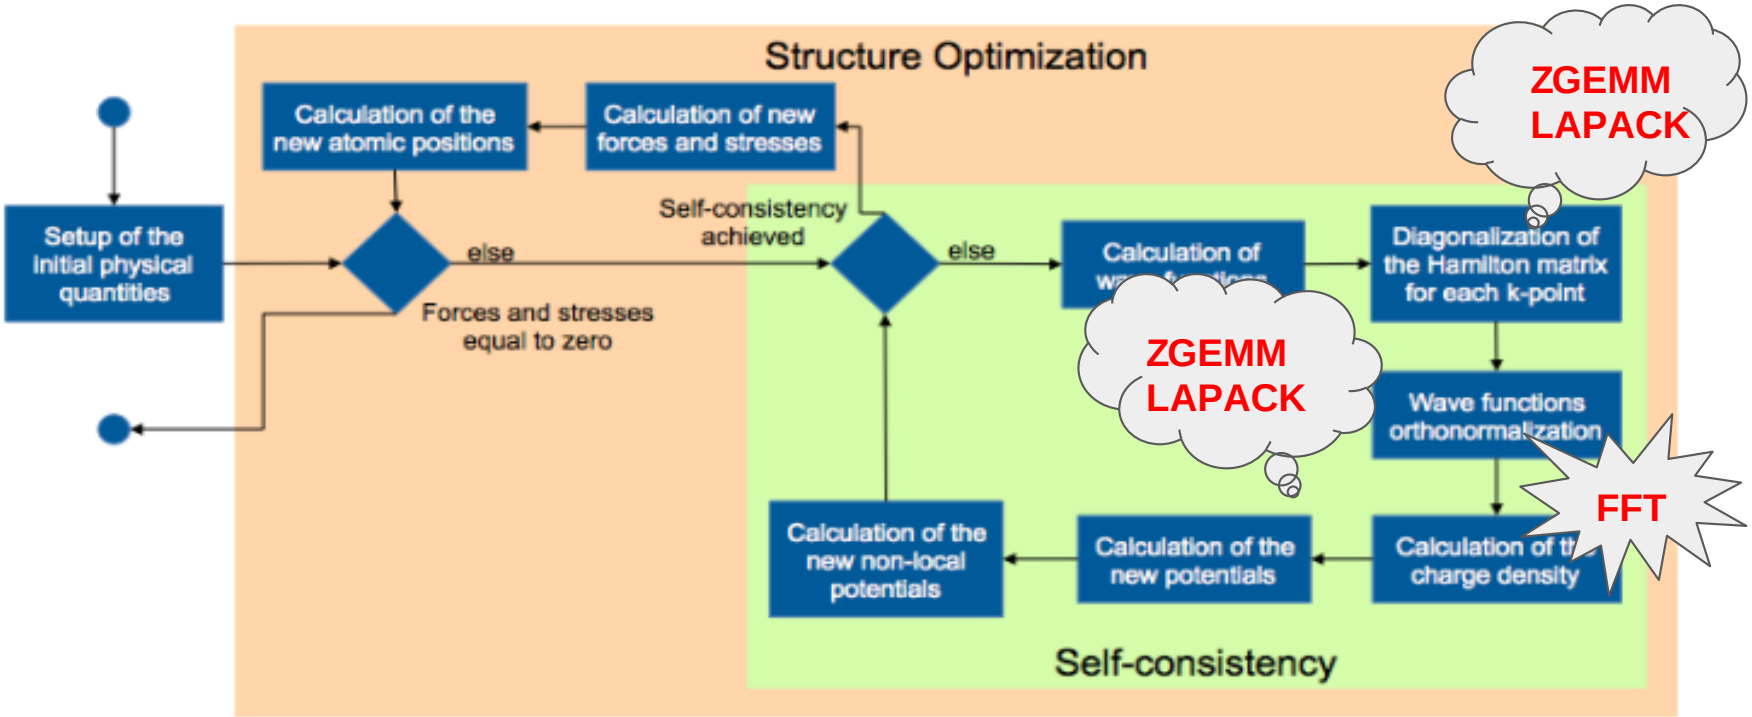

Analysis performed by POP

We provided to POP the traces from a Car Parrinello calculation for a system of 32 water molecules. The run was performed on the PICO Cineca cluster (Intel Xeon 10 Core E5-2670v2 2.50 GHz).

We ran using 16 MPI processes, each with 1 and 5 OpenMP thread per node.

The EXTRAE traces were analyzed and we got back the report from POP

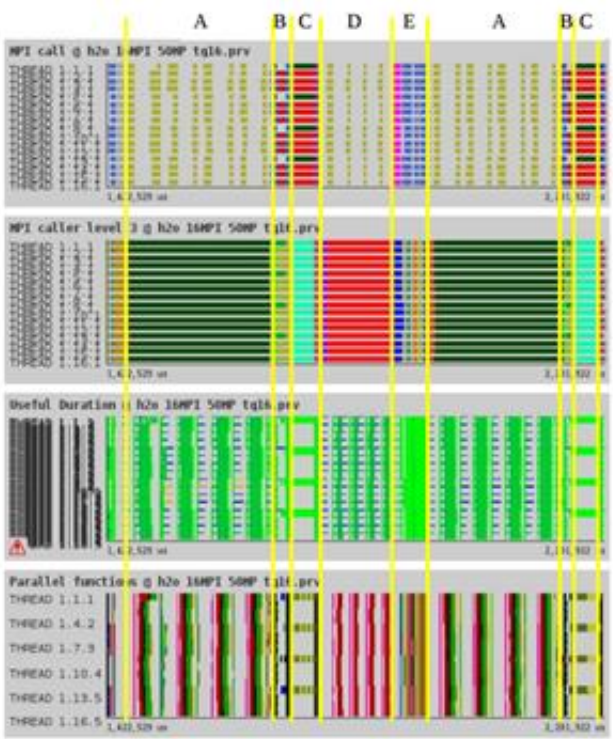

Figure 2: Timelines for the main ROI for the trace of 16 MPI x 5 OpenMP threads

Results of the analysis

Among the most important results presented in the report we focused on the FFT and on the dependence on the differenct choices for the taskgroup distribution.

The report proposed an approach aimed to the overlap of computation and communication in this phase, with a coarse grained taskification (OpenMP or OmpSs based)

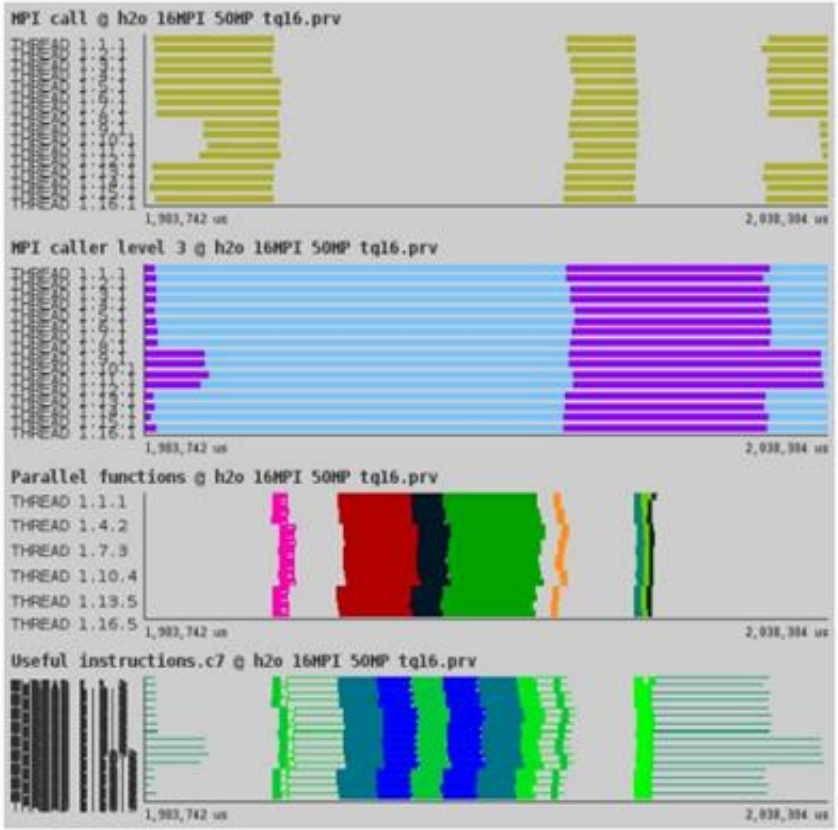

Figure 17: One FFT iteration. a) MPI call, b) MPI caller at call stack level 3, c) outlined parallel function, d) useful duration

Follow up

Standing on the results contained in the POP report, and in the framework of the MaX activity we:

- restructured the fine/coarse grain parallel distribution of the FFT
- built a mini-app containing the FFT kernel extracted from the QE distribution, used as a benchmark for algorithmic improvements and as a tool for validation and profiling of the performances
- implemented an overlap mechanism based on the double-buffer technique in the FFT cycle
- are looking forward to implement a task based approach for the execution of fine grained units of work

Conclusions

POP provided an essential **service** to MaX and Quantum ESPRESSO that permitted to:

- **highlight the weakness** in the implementation and the targets of future actions
- **identify a direction** for the development and optimization of the code

The provided suggestions have been considered for the next developments that up to now led to

- build a mini-app of the FFT for **validation and testing**
- implement **new techniques** for the improvement of the parallelization

More important, this activity started a productive **feedback and follow-up** that, in the framework of MaX is actively contributing to the improvement of the QE distribution