Extended Software Development Workshop: Scaling Electronic Structure Applications

Location: CECAM-IRL, UCD, Ireland
Webpage: https://www.e-cam2020.eu/legacy_event/extended-software-development-workshop-scaling-electronic-structure-applications/
Dates: January 7, 2019 to January 18, 2019
Organizers: Nick R. Papior, Yann Pouillon, Micael Oliveira, Fabiano Corsetti, Volker Blum, Emilio Artacho, Alan O'Cais

1 State of the art

The evolutionary pressure on electronic structure software development is greatly increasing, due to the emergence of new paradigms, new kinds of users, new processes, and new tools. The large feature-full codes that were once developed within one field are now undergoing a heavy restructuring to reach much broader communities, including companies and non-scientific users. More and more use cases and workflows are performed by highly-automated frameworks: high-throughput calculations and computational materials design, large data repositories, and multiscale modeling, for instance. At the same time, HPC Centers are paving the way to exascale, with a cascade of effects on how to operate, from computer architectures to application design.

These trends are highly challenging for the electronic structure community. As a result, long-ignored, well-established software engineering good practices are now adopted at an accelerating pace by more and more software projects. With time, this kind of migration is becoming a question of survival, the key for a successful transformation being to allow and preserve an enhanced collaboration between the increasing number of disciplines involved. Significant efforts of integration from code developers are also necessary, since both hardware and software paradigms have to change at once.

Addressing these issues requires coordinated efforts at multiple levels:

- through the creation of open standards and the use of co-design, both for programming and for data
- with a significant leap in documentation policies, helped by tools like Doxygen, Sphinx and ReadTheDocs
- by introducing test-driven development concepts and systematically publishing test suites together with software
- by creating synergies with popular software distribution systems (e.g. EasyBuild, Spack, or MacPorts)

- by disseminating the relevant knowledge and training the community, through the release of demonstrators and giving all stakeholders the opportunity to meet regularly.

2 Main outcomes and training provided

The first two days of the workshop consisted in presentations and discussions that were roughly divided into four main topics:

- current state and future plans for the Electronic Structure Library;
- methodological and software developments in Electronic Structure;
- current and future trends in high-performance computing;
- performance and scalability of ESL libraries/components.

The participants identified the following aspects as the main challenges that the ESL must overcome:

- The ESL needs more involvement from the electronic structure community to become sustainable in the long term.
- Wider integration by community codes, as code developers are still reticent in using code written by others.
- Increase project visibility by better publicizing the work done within the ESL and increasing the outreach activities.
- Improve the scalability, optimization, and hardware awareness of ESL components in order to ease the efforts of electronic structure codes in addressing the various issues related to the current hardware race.
- Engage more with other projects and communities that share similar objectives.

Concerning the methodological and software developments in Electronic Structure, participants presented how the different projects they are working on are handling the changes in HPC architectures and their use of external libraries. This overview was nicely complemented by a dedicate session on HPC, were it was highlighted that hybrid architectures (CPUs+GPUs) are beginning to dominate the market and that there is still quite some uncertainty regarding what technology will come next. As a consequence, developers of electronic structure codes will have to adapt to these changes in computer architectures, but without necessarily being able to drive them. Different strategies to deal with these changes were also presented and discussed.

A full session was dedicated to the performance and scalability of ESL libraries. Developers were given the opportunity to present the libraries, with a focus on the parallelization strategies used, and the corresponding performance and scalability. They were also encouraged to discuss existing bottlenecks and possible ways to overcome them. This information was essential to plan part of the work to be done during the coding sessions.

A hands-on tutorial on profiling parallel applications was organized before the start of the coding sessions. This tutorial was found to be very useful by the participants and will hopefully lead to more collaborations between ESL developers and the MaX and Pop CoEs.

Lectures from this workshop were recorded and stored on E-CAM's online training portal at https://training.e-cam2020.eu/collection/5c3329f4e4b0f6515d0b995d

3 List of software development projects

During the workshop the participants gathered in smaller groups to work on specific projects.

ESL Bundle and Demonstrator
A procedure and a detailed timeline to release new versions of the Bundle with updated components was decided and tested. This lead to a new release (0.3.1). It was also decided to do two new major/minor releases every year. New Docker images were generated for this new release to be used in the ESL Continuous Integration system and a new repository of ESL Easyconfig files for the EasyBuild build and installation framework was created. Work was also done in the ESL Demonstrator to complete some of its features.

ELSI
Work on GPU support for ELSI and its solvers was started and benchmarks of the MAGMA eigensolvers and for ELPA 1-stage and 2-stage eigensolvers were performed. The SIESTA-ELSI interface was updated. Several updates to the build system and to the solvers included in ELSI.

GridXC and Libvdwxc
Benchmarks for the two libraries were performed, followed by some more in-depth analysis of the FFT performance.

ESCDF
The extension of the API for optimal parallel scalability and performance was thoroughly discussed and work on its implementation was started.

4 Participant list

Organizers

Papior, Nick R.
Technical University of Denmark, Denmark

Pouillon, Yann
Universidad de Cantabria, Spain
Oliveira, Micael
Max Planck Institute for the Structure and Dynamics of Matter, Hamburg, Germany
Corsetti, Fabiano
Synopsys QuantumWise, Denmark
Blum, Volker
Duke University, Durham, NC, USA, USA
Artacho, Emilio
Cavendish Laboratory, University of Cambridge, United Kingdom
O'Cais, Alan
Jülich Supercomputing Centre, Germany

Bosoni, Emanuele - CRANN, School of Physics, Trinity College Dublin, Ireland
Caliste, Damien - Alternative Energies and Atomic Energy Commission (CEA), France
de Jong, Wibe - Lawrence Berkeley National Laboratory, USA
Elena, Alin Marin - Daresbury Laboratory, United Kingdom
Garcia, Alberto - Institute of Materials Science, Barcelona, Spain
Garcia Suarez, Victor Manuel - University of Oviedo, Spain
Gimenez, Judit - Barcelone Supercomputing Centre (BSC), Spain
Gonze, Xavier - Université Catholique de Louvain, Belgium
Hourahine, Benjamin - University of Strathclyde, United Kingdom
Keal, Thomas - STFC Daresbury Laboratory, United Kingdom
Larsen, Ask Hjorth - Nano-bio Spectroscopy Group and ETSF Scientific Development Centre, Departamento de Fisica de Materiales, Spain
Lueders, Martin - Daresbury Laboratory, United Kingdom
Martin-Samos Colomer, Layla - University of Nova Gorica, Slovenia
Ohlmann, Sebastian - Max Planck Computing and Data Facility, Germany
Payne, Mike - University of Cambridge, United Kingdom
Rampp, Markus - MPCDF, Germany
Saxe, Paul - MolSSI - The Molecular Sciences Software Institute , USA
Wylie, Brian - Juelich Supercomputing Centre (JSC), Germany
Yu, Victor - Duke University, USA
Sanvito, Stefano
Trinity College Dublin, Ireland
Patterson, Charles
Trinity College Dublin, Ireland
O'Regan, David D.
Trinity College Dublin, Ireland
MacKernan, Donal
University College Dublin, Ireland