HPC Hackathon

Industrial code optimization

Principle

We aim with this hackathon to bring together Master II-level HPC students in a virtual competition around computational codes provided mostly by Electricité De France - EDF (Saturn and Telemac code) and CGG (seismic core). Other industrial partners might complete this first list.

This competition will rely on Amazon Web Services - AWS instances based on Arm technologies. Indeed, the target architectures (AWS Graviton 2 and 3 processors) offer certain approaches (software ecosystem, design) that motivate a specific effort compared to traditional Intel or AMD type architectures. Typically, on the Graviton 3 processor, we can mention the use of SIMD units (Scalable Vector Extension/SVE compared to AVX2 or AVX-512 on x86 architectures) or the availability of 8 DDR5 memory channels.

This hackathon is structured around computational codes, software environments and hardware solutions that have already been tested by industry. The compilation recipes and optimization phases have been validated before the event. The students will therefore be in a framework close to a guided practical session with the opportunity to increase their understanding of industrial issues around high-performance simulation. No specific contributions from students in terms of porting or optimization are expected (these codes are largely proven in a production context and on multiple hardware architectures). The participants will nevertheless be credited by the industrial partners (the codes have an open status) in case of significant advances.

Details

Students are encouraged to take an iterative approach to porting and optimizing these industrial codes/kernels. The latter are of different nature/complexity and will allow the deployment of the classical phases of taking control of a scientific application. The evaluation of the teams (4 students maximum per team) will be based on the following points:

- **Porting**: it will be a question of validating the application on Arm architectures (Graviton2/Graviton3) by focusing on the test case provided by the industrial partner. The validation will be done by comparing the result files and/or by comparing the results on different platforms (x86/Arm).

- **Profiling**: Students will be asked to use classical application profiling tools to identify performance locks... this will include identifying the hotspots of these applications (compiler report, dynamic code analysis...)

- **Advanced optimization**: During this phase, the students will be able to make some modifications to the codes in order to improve the performances. For small applications (e.g., CGG code, this may involve adding OpenMP directives or modifying the organization of loops ...). In the case of complex code, participants can focus on the impact of the different compilation chains and work on extracting some kernels (mini-apps). The latter could then be the subject of specific efforts.

- Arm, Nvidia and GNU compilers will be available during the hackathon. Students will also have access to various code profiling tools (e.g., Arm Allinea Forge).

- The synthesis of the obtained results (compilation recipes, validation procedure, optimization, loops extraction ...) will be made available on a Git repository in order to facilitate the evaluation and the reproducibility.

Proposed codes / Identified Kernels

CGG

seismic core: www.cgg.com

- Focus: numerical validation (x86 Intel/AMD comparison), evaluation of different compilers (Arm, Nvidia, GNU), Vectorization

EDF

Code Saturne: https://www.code-saturne.org/cms/web/

- Focus: porting and numerical validation (intel/AMD comparison), evaluation of different compilers, possibly extraction of kernels (hotspots) for advanced analysis.

Code Telemac: http://www.opentelemac.org/

- Focus: porting and numerical validation (intel/AMD comparison), evaluation of different compilers, possibly extraction of kernels (hotspots) for advanced analysis

Format

Planning:

- Announcement at the 2022 Teratec Forum: June 14-15, 2022
- Team registration: September – October 2022
- Webinar Launch: **October 7 from 4:00 to 5:30 pm**
- Open training: **October / November**
- **Hackathon from Monday, 28th of November, 2022 at 9:00 am to Monday, 5th of December, 2022 at 9:00 am**
    a. Estimated 48h work left to the choice of the institution within the week
    b. Team of 4 students
    c. ARM/AWS support

Prize

The competition will result in a ranking, and the "winning" team will be awarded a prize based on ARM processors that smell like apple

Industrial grounding

Specificities of the codes:

The HPC codes are representative of the interest and concern of TERATEC's industrial partners. Controlling the source code and dependencies (possibility of recompiling the code) is an important point, as it is for any porting of a source architecture to a new architecture.

Each code will be presented by the industrial partner to explain its interest. This presentation will take place during a webinar on the 7th of October in front of registered students during which the platform will also be presented.

ARM/AWS/UCIT

Europe has a strategy in order to build a European processor (EPI) using the Arm based instruction set (AArch64) which should arrive on the market in 2023 with SiPearl. This processor will be used in particular by ATOS as recently announced with their latest generation of supercomputer).

With Amazon graviton2 in 2019, AWS introduced the ability for communities around the world to develop, build, and test applications using the arm64 instruction set on their Cloud platform. Adoption has been impressive with many open source and commercial applications running on AWS Cloud and significant investments to improve the software development tool chain.

With AWS graviton3 available since mid-May, AWS is now adding additional mechanisms (e.g., SVE) to support advanced technical workloads such as High Performance Computing. AWS wants to provide a valuable alternative to the current solution (mainly x86 based, with a strong adoption of GPUs lately) and will give researchers and engineers access to build, test and prove their numerical simulation code using an architecture that will be very similar to the one SiPearl Rhea will provide.

UCit will administer and make available the necessary computational resources to Hackathon participants. Thanks to Cluster-in-a-Box (https://ucit.fr/index.php/cluster-in-a-box/), each team will be able to have their own customized HPC environment on the AWS Cloud and access it seamlessly

Pre-Hackathon Support:

Training on how to use the AWS platform + compilation and behavior analysis by ARM and AWS.

Support during Hackathon:

The use of slack messaging for asynchronous support is considered to provide technical assistance to the different teams. Each team will have a specific channel and shared technical channels will be available (depending on the applications for example), in order to allow the Arm and AWS technical teams to remove possible locks.

Contacts:

TERATEC: Emmanuelle Vergnaud – emmanuelle.vergnaud@teratec.fr

ARM: Conrad Hillairet - Conrad.Hillairet@arm.com

AWS: Gilles Tourpe - gtourpe@amazon.com

UCIT: Jorik Remy - jorik.remy@ucit.fr