Contribution ID: **14** Type: **Short talk**

Programming Heterogeneous Architectures using Hierarchical Tasks

Thursday 23 March 2023 14:40 (10 minutes)

Task-based systems have become popular due to their ability to utilize the computational power of complex heterogeneous systems. A typical programming model used is the Sequential Task Flow (STF) model, which unfortunately only supports static task graphs. This can result in submission overhead and a task graph that is not well-suited for execution on heterogeneous systems. A common approach is to find a balance between the granularity needed for accelerator devices and the granularity required by CPU cores to achieve optimal performance. To address these issues, we have extended the STF model in the STARPU runtime system by introducing the concept of hierarchical tasks. This allows for a more dynamic task graph and, when combined with an automatic data manager, it is possible to adjust granularity at runtime to best match the targeted computing resource. Additionally, submission overhead is reduced by using large-grain hierarchical tasks, as the submission process can now be done in parallel. We have shown that the hierarchical task model is correct and have conducted an early evaluation on shared memory heterogeneous systems using the CHAMELEON dense linear algebra library.

JLESC topic

Primary author: LUCAS, Gwenole (Inria Bordeaux)

Presenter: LUCAS, Gwenole (Inria Bordeaux)

Session Classification: Short Talks on Tasking

Track Classification: Programming languages and runtimes