StarPU scheduler vs. the analyst: who is right?

Vincius Garcia Pinto
Informatics Institute
Federal University of Rio Grande do Sul
Porto Alegre, Brazil
vinicius.pinto@inf.ufrgs.br

Lucas Mello Schnorr
Informatics Institute
Federal University of Rio Grande do Sul
Porto Alegre, Brazil
schnorr@inf.ufrgs.br

Arnaud Legrand
CNRS, INRIA, LIG
Univ. Grenoble Alpes
Grenoble, France
arnaud.legrand@inria.fr

***Abstract*—State-of-the-art High-Performance Computing (HPC) platforms are built with a hybrid design combining multicore processors and manycore accelerators. This design motivates the use of appropriate programming models (e.g., task parallelism) in order to reduce programming complexity, increase the performance portability and efficiently exploit such multi-level parallelism. Task-based executions are supported by dynamic runtime systems which are in charge of several critical procedures such as scheduling, load balancing, communications, and synchronizations. This way, the overall application performance depends heavily on how efficient and correct are the decisions taken by this runtime. In this work, we propose a workflow to investigate potential scheduling mistakes in the StarPU runtime system. This workflow relies on debugging and simulation strategies to enable us to rebuild the estimations computed by the scheduler at the exact moment when a new task is scheduled. We evaluate our approach by comparing the scheduler estimations and the final execution of a set of delayed tasks. Our results allow us to refute scheduling mistakes and on the other hand to identify that the real origin of the delays is related to pipeline and prefetch mechanisms.**

***Index Terms*—task parallelism, GPU, hybrid-architectures, StarPU**