

---

# Faithful and Efficient Simulation of High Performance Linpack

Arnaud Legrand\*<sup>1</sup>

<sup>1</sup>Laboratoire d'Informatique de Grenoble (LIG) – CNRS : UMR5217, Université Pierre-Mendès-France - Grenoble II, Institut polytechnique de Grenoble (Grenoble INP), Université Joseph Fourier - Grenoble I – UMR 5217 - Laboratoire LIG - 38041 Grenoble cedex 9 - France Tél. : +33 (0)4 76 51 43 61 - Fax : +33 (0)4 76 51 49 85, France

## Abstract

With a power consumption of several MW per hour on a TOP500 machine, running applications on supercomputers at scale solely to optimize their performance is extremely expensive. Likewise, High-Performance Linpack (HPL), the benchmark used to rank supercomputers in the TOP500, requires a careful tuning of many parameters (problem size, grid arrangement, granularity, collective operation algorithms, etc.) and supports exploration of the most common and fundamental performance issues and their solutions. In this talk, we will explain how we both extended the SimGrid's SMPI simulator and slightly modified the open-source version of HPL to allow a fast emulation on a single commodity server at the scale of a supercomputer. We explain how to model the different components (network, BLAS, ...) and show that a careful modeling of both spatial and temporal node variability allows us to obtain predictions within a few percents of real experiments.

**Keywords:** performance prediction, consequences of variability

---

\*Speaker