

Parallel and Distributed Simulation of Large-Scale Distributed Applications

Advisor: Martin Quinson (IRISA, team Myriads). Martin.Quinson@ens-rennes.fr

Duration: 4 to 6 months. Follow-up possible in PhD thesis.

Executive summary: The context of this project is to allow the efficient parallel and distributed simulation of large systems within the SimGrid framework. The proposed work will improve the existing parallel simulation mode, and propose a novel distributed simulation mode. We target a simulation comprising millions of heavy computational nodes on a much smaller cluster.

Key skill required: System programming and networking programming in C on Linux.

Context

Recent and foreseen technical evolution allow to build information systems of unprecedented dimensions. The potential power of the resulting distributed systems offers new possibilities in terms of applications, be them scientific such as multi-physic simulations in High Performance Computing (HPC), commercial in the Cloud with the data centers underlying the Internet, or public in very large peer-to-peer systems. For example, ExaScale systems in the HPC area are expected to aggregate millions of high end compute nodes by the end of this decade for unprecedented scientific computations.

Evaluating computer systems of this extreme scale raises severe methodological challenges. Simply executing them is not always possible as it requires to build the complete system beforehand (what is not possible for ExaScale systems for example), and it may not even be enough when uncontrolled external load prevents reproducibility. Simulation is an appealing alternative to study such systems. It may not capture the whole complexity of every phenomena, but allows to easily capture some important trends, while ensuring the controllability and reproducibility of experiments.

SimGrid¹ (developed by the Myriads team in collaboration) is a toolkit providing core functionalities for the simulation of distributed applications in heterogeneous distributed environments. The specific goal of the project is to facilitate research in the area of distributed and parallel application scheduling on distributed computing platforms ranging from simple network of workstations to Computational Grids.

This framework was shown orders of magnitude faster than concurrent simulators such as GridSim or PeerSim, and can simulate up to a few million lightweighted P2P processes on a single node [1]. This falls however short to simulate ExaScale systems, as these systems are expected to count dozen of millions of heavy processes. Both CPU and memory limitations must be overtaken to scale the simulation further. In a previous work, we shown that parallel simulation can improve the computational performance in some cases [2], but the memory limitation claim for the distribution of the simulation to leverage the memory of several computational nodes.

¹ The SimGrid Project: <http://simgrid.org>

Precise Work Description

First, we would like to improve the *parallel* execution mode described in [2]. Currently, only events occurring at the exact same timestamp are executed in parallel. This criteria must be relaxed to increase the amount of potential parallelism exhibited by the application. This require to determine the future events that are independent of the events executed in the current timestamp. Such future events can safely be added to the current timestamp. This detection must rely on an independence theorem, very similar (but still novel) to the ones used when applying Dynamic Partial Order Reduction on model checking.

Then, we would like to propose a *distributed* execution mode allowing to overcome memory limitations in very large scenarios. Several designs are possible to that extend. The intern is expected to develop several proof of concepts to understand their relative advantages. S/he will then select the best design through a careful evaluation before implementing the selected design. The ultimate goal is to run a typical HPC application (such as linpack, used for the Top 500 ranking²) using a sizable portion of the Grid'5000 experimental facility³.

References

- [1] Laurent Bobelin, Arnaud Legrand, David Marquez, Pierre Navarro, Martin Quinson, Frédéric Suter, Christophe Thiéry. *Scalable Multi-Purpose Network Representation for Large Scale Distributed System Simulation*. 12th Intl Symposium on Cluster Computing and the Grid (CCGrid'12), 2012. <http://hal.inria.fr/hal-00650233>
- [2] Martin Quinson, Cristian Rosa, Christophe Thiéry. *Parallel Simulation of Peer-to-Peer Systems*. 12th ACM/IEEE Intl Symposium on Cluster Computing and the Grid (CCGrid'12), Canada, May 2012. <http://hal.inria.fr/inria-00602216>

Skills required

In addition to the skills that can reasonably be expected from Master-level C, and of modern Unix-based Operating Systems such as Linux.

² Top 500 SuperComputing ranking: <http://www.top500.org>

³ The Grid'5000 Scientific Instrument: <https://www.grid5000.fr>