Leveraging multiple experimentation methodologies to study P2P broadcast

Lucas Nussbaum & Martin Quinson
LORIA – AlGorille team

Keywords:
Distributed systems (Grid, P2P, HPC, Cloud); experimentation methods (simulation, emulation, real-scale experiments); SimGrid; Grid'5000

General information

Encadrants	Lucas Nussbaum	Martin Quinson
Email	Lucas.Nussbaum@loria.fr	Martin.Quinson@loria.fr
Téléphone	03 54 95 86 19	03 83 59 20 98
Bureau	LORIA, B 118	LORIA, B 122

Context

Distributed systems such as grids, clusters, peer-to-peer systems, high-performance supercomputers, cloud computing infrastructures or desktop computing environments, benefit of an ever increasing popularity nowadays. Distributed applications (such as decentralized data sharing solutions, games, scientific application, high-traffic web applications or scientific computations) are executed routinely on these systems.

By nature, the resulting environments and applications are extremely complex and dynamic because they aggregate thousands of elements that are heterogeneous and shared among several users. This make these systems very challenging to study, test, and evaluate. Computer scientists traditionally study their systems *a priori* by reasoning theoretically on the constituents and their interactions. But the complexity of these systems make this methodology is near to impossible, explaining that most of the studies are done *a posteriori* through experiments.

Three main methodologies exist to experiment with computer systems: real-scale, simulation and emulation. *Real-scale* (or *in situ*) consists in executing the real application under study on an experimental platform like Grid'5000 (a large scale experimental platform in France, composed of more than 1600 machines). On the opposite, with *simulation*, both the application and the environments are replaced by models, and the interactions between both models are computed by a simulator. *Emulation* can be seen as an intermediate approach where the real application is executed within a synthetic environment. Typically, one will use a homogeneous cluster of machines as an execution environment, and use an emulation layer to reproduce the complex conditions found on the real Internet.

The AlGorille team is deeply involved in all those methodologies. It has a leadership role in the world leading *Grid'5000* testbed and *SimGrid* simulator, and develops entirely the *distem* emulator. We see those different methodologies as complementary approaches to work on the different steps of the scientific workflow: ideas are first maturated into algorithms in the simulator, before they are implemented in as prototypes tested on a real-scale testbed. Once the prototype is known to work, an emulator is used to evaluate the behavior of the prototype under various experimental conditions.

Description

The goal of this internship is to combine all those methodologies to answer a given question. While answering the question is a valuable contribution by itself, the main contribution is expected to be lessons learned while trying to combine the various methodologies. The proposed question is:

Are P2P broadcast algorithms such as the ones used in BitTorrent suitable for high-performance data broadcast on high-bandwidth networks such as the ones found in Grids or Clouds?

The data replication problem that BitTorrent solves elegantly for the end-users of the Internet can be found in many other contexts. For example, the replication of the data produced by the LHC scientific instrument among the interested parties poses severe challenges, as the data sets last hundred of gigabytes each. Also, rebooting the hundred of machines of a cluster may require to send them a complete image of the operating system to start. Although tempting at the first glance, using BitTorrent as is in these contexts lead to poor performance, for some reasons that we propose to discover in this internship.

Answering that question will require the student to perform experiments with *SimGrid*, *Grid'5000* and *distem*, under a wide variety of configurations, to understand the root causes of the encountered performance limitations, be them in the models and algorithms used or in their implementation. Beyond these technical elements, the student will have to think deeply about the methodology used, to combine the tools to stack their respective advantages in each situations.

Skills required

In addition to the skills that can reasonably be expected, the applicant must be able to work efficiently in a Linux environment, and should have a good understanding of systems and networks.

1 Links

- SimGrid: http://simgrid.gforge.inria.fr/
- Grid'5000 testbed: https://www.grid5000.fr/
- distem: http://distem.gforge.inria.fr/
- Research team (ALGORILLE): http://www.loria.fr/equipes/algorille/
- Tutors: http://www.loria.fr/~quinson/ and http://www.loria.fr/~lnussbau/

References

[1] Jens Gustedt, Emmanuel Jeannot and Martin Quinson, "Experimental Validation in Large-Scale Systems: a Survey of Methodologies", In *Parallel Processing Letters*, 19:399-418, 2009. http://hal.inria.fr/inria-00364180/en/