Real-time online emulation of real applications on SimGrid with Simterpose

Executive summary: The goal of this project is to design an evaluation environment for distributed applications (e.g.; P2P applications) where the instances of the real application (a standard & unmodified BitTorrent client) are executed in a virtual environment simulated by the SimGrid simulator.

Key skills required: system programming in C on Linux; deep understanding of OS principles

Research team name:	AlGorille (leader: Jens Gustedt)
Research Unit:	Nancy – Grand Est
Intern tutors:	Lucas Nussbaum <lucas.nussbaum@loria.fr>
	Martin Quinson <martin.quinson@loria.fr>
Intern level:	Master student (or PhD student)
Internship duration:	4 to 6 months
Followed by a PhD:	possible (but not mandatory)

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 applications, high-traffic web applications) 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 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 experiments using testbeds, simulation and emulation. *Real-scale* (or *in situ*) consists in executing the real application under study on an experimental platform like Grid'5000 or PlanetLab. On the opposite, with *simulation*, both the application and the environment 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.

SimGrid (developed by the AlGorille team in collaboration with other teams) 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. It is however not possible to use real applications directly on SimGrid: users have to extract the logic of their applications and rewrite them using the specific interfaces of SimGrid.

The Simterpose project, which is the core of this internship proposal, tries to alleviate this by providing a way to use SimGrid as an emulator. This would allow real applications to be executed on virtual platforms emulated by SimGrid. This project naturally relates to the *Distem* emulator also developed in the AlGorille team, but follows a completely different approach. Distem emulates the target platform by reducing the performance of the host platform running the experiment while Simterpose intercepts all computations and communications and delay them according to the computations of the simulator.

Description

Simterpose is currently under development. It currently only allows the extraction of a trace of applications' actions that could be suitable for replay in SimGrid. The simulator is not able to deal with these traces yet, neither for offline analysis once the complete trace has been captured, nor online directly to delay the application's actions according to simulation results provided by SimGrid.

The goals of this internship are:

- Develop a simple offline simulator of Simterpose traces;
- Continue the development of Simterpose to provide real-time online emulation on top of SimGrid;
- Evaluate Simterpose by running real distributed applications: P2P applications, high-performance computing applications written using MPI, . . . ;
- Distribute the execution of the user application on several cluster nodes (with a centralized SimGrid instance to provide the simulation) to achieve large-scale emulation using SimGrid and Simterpose. The developed solution will be evaluated on Grid'5000;
- Use Simterpose to understand the semantics of collective operations in MPI implementations such as OpenMPI and MPICH2. The long term goal is to integrate the resulting models in SMPI, the SimGrid-based MPI simulator;
- Use the SimGrid model-checker with Simterpose to provide an exhaustive exploration of execution paths in a real application, for bug-finding purposes.

Skills required

In addition to the skills that can reasonably be expected from Master-level students, the applicant should have a strong knowledge of system programming in C, and of modern Unix-based Operating Systems such as Linux.

Links

- SimGrid : http://simgrid.gforge.inria.fr/
- Research team (ALGORILLE): http://www.loria.fr/equipes/algorille/
- Tutors: http://www.loria.fr/~quinson/ and http://www.loria.fr/~lnussbau/
- Publication from a previous intern on Simterpose:
  M. Guthmuller, L. Nussbaum et M. Quinson. Émulation d'applications distribuées sur des plates-formes virtuelles simulées
  http://hal.archives-ouvertes.fr/inria-00565341/en/