# Accurate Emulation of Processor and Memory in Wrekavoc

Master Internship

**Research team name:** AlGorille
**Research Unit:** Nancy – Grand Est
**Research theme:** NUM (AlGorille)
**Research team leader:** Jens Gustedt
**Intern tutors:** Lucas Nussbaum
**Intern level:** Master or PhD student
**Internship duration:** 4 to 6 months
**Possibility of a follow-up Ph-D:** yes

## 1. Context and Contact

Distributed computing and distributed systems is a branch of computer science that has recently gained very large attention. Grids, Clusters of clusters, Peer-to-peer systems, desktop environments, are examples of successful environments on which applications (scientific, data managements, etc.) are executed routinely.

However, such environments, are composed of different elements that make them more and more complex. The consequence is that applications (and the algorithms implemented by them) are also very complex and very hard to validate.

In general, *experimental validation* is used to study such environments. One can distinguish three different methodologies for performing experiments [GJQ09]: 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 or PlanetLab. On the opposite, with *Simulation*, a model of the application is executed against a model of the environment. *Emulation*, which is at the core of the internship proposal, can be seen as an intermediate approach, since the real application is executed, within synthetic conditions.

Wreakavoc [CJ06, DGJ08] is an emulator developed in the AlGorille team. It typically allows to execute distributed applications on a cluster of machines, emulating different features of the environment (CPU speed, memory size, network characteristics – bandwidth, latency, topology) by using system-level tools, like the network emulator included in the Linux kernel.

## 2. Description

Wrekavoc already provides CPU and memory limitation [CJ06]. But recent developments in the Linux kernel (*cgroups*, also known as *Linux Containers*) could be used to improve the CPU and memory limiters, providing a more accurate emulation of slower CPUs, and more complex usage scenarios (applications composed of several Linux processes on the same node, for example).

Additionally, *cgroups* could be used as the basis for implementing *process folding*, that is, the ability to execute several instances of the application under study on the same physical node. *Process folding* is a prerequisite for performing experiments with thousands of nodes.

The goals of this internship are:

1. To evaluate the capabilities of *cgroups* in the context of distributed systems emulation ;
2. To propose a model for CPU and memory limitation based on *cgroups* ;
3. To implement that model in Wrekavoc ;
4. To evaluate it with experiments on real distributed systems.

# References


[CJ06] L.-C. Canon and E. Jeannot. Wrekavoc: a tool for emulating heterogeneity. *Parallel and Distributed Processing Symposium, International*, 0:129, 2006.

[DGJ08] Olivier Dubuisson, Jens Gustedt, and Emmanuel Jeannot. Multi-Site Emulation using Wrekavoc: Validating Distributed Algorithms and Applications. Research Report RR-6696, INRIA, 2008.

[GJQ09] Jens Gustedt, Emmanuel Jeannot, and Martin Quinson. Experimental Validation in Large-Scale Systems: a Survey of Methodologies. *Parallel Processing Letters*, 19:399–418, 2009.