

## Education

2017 – 2021 Grenoble (FR)	<b>PhD in Computer Science</b> Under the supervision of Arnaud Legrand. High Performance Computing: Towards Better Performance Predictions and Experiments <ul style="list-style-type: none"><li>I present a methodology for predicting the performance of parallel non-trivial applications through simulation. I describe several models for communications and computations, with an increasing complexity. I compare these models through an extensive validation by matching the predictions with real experiments. This validation shows that modeling the spatial and temporal variability of the platform is essential for faithful predictions. As a consequence, predictions require careful sensibility analysis accounting for the uncertainty on the resource models, which I illustrate through several case studies.</li><li>I present the lessons learned while making the numerous experiments required in the first part and how I improved the methodology. I show that measurements can suffer from multiple experimental biases and I explain how some of these biases can be overcome. I also present how I implemented systematic performance non-regression testing, which allowed to detect many significant changes of the platform throughout this thesis.</li><li>Main results published at Cluster'19 conference (CORE A) [CLH19] and currently under review for publication in JPDC journal (CORE A) [CL21].</li></ul>	Université Grenoble Alpes
2015 – 2017 Grenoble (FR)	<b>Master's &amp; Engineering Degrees in Computer Science (M1-M2)</b> Graduate specialization in parallel and distributed systems. Obtained a Master of Science, with highest honor, ranked 2 <sup>nd</sup> /88.	Ensimag
2013 – 2015 Lyon (FR)	<b>Bachelor's Degree in Theoretical Computer Science (L3-M1)</b> Undergraduate and postgraduate intensive program in theoretical computer science. Obtained a Bachelor of Science, with great honor.	ENS Lyon
2011 – 2013 Grenoble (FR)	<b>Undergraduate program (L1-L2)</b> Undergraduate program in computer science and mathematics.	Université Grenoble Alpes

## Internships & Mobility

Oct/17 – Dec/17 Chicago (US)	<b>Performance Variability in Supercomputers</b> Under the supervision of Swann Perarnau. <ul style="list-style-type: none"><li>Carried experiments and statistical analyses to characterize performance variability.</li></ul>	Argonne Laboratory
Feb/17 – Jul/17 Grenoble (FR)	<b>Efficient Simulation of Large-Scale MPI Applications</b> Under the supervision of Arnaud Legrand. <ul style="list-style-type: none"><li>Modified the simulator and the simulated application to enable large scale simulations.</li><li>Outcome: simulate executions several orders of magnitude larger.</li><li>Obtained results for the large scale experiment of a conference article [Hei+17].</li></ul>	Inria
May/16 – Aug/16 Walldorf (DE)	<b>Multicast Communication in SAP HANA</b> Under the supervision of Norman May. <ul style="list-style-type: none"><li>Analyzed several multicast algorithms.</li><li>Implemented these algorithms in C++, using HANA codebase.</li></ul>	SAP
May/15 – Aug/15 Grenoble (FR)	<b>Job Isolation in Fat Tree Topologies</b> Under the supervision of Matthieu Perotin. <ul style="list-style-type: none"><li>Designed several algorithms to prevent the leak of sensitive information in a cluster.</li><li>Results published in a conference article [PC17].</li></ul>	Bull
Jun/14 – Jul/14 Sophia-Antipolis (FR)	<b>Modeling and Verification of Concurrent Systems</b> Under the supervision of Robert de Simone.	Inria
Jun/13 – Jul/13 Grenoble (FR)	<b>Monitoring of Distributed Systems</b> Under the supervision of Yliès Falcone. <ul style="list-style-type: none"><li>Designed an algorithm for decentralized monitoring of distributed systems.</li><li>Results published in a conference article [FCF14].</li></ul>	Université Grenoble Alpes
Jun/12 Grenoble (FR)	<b>Monitoring of Distributed Systems</b> Under the supervision of Yliès Falcone.	Université Grenoble Alpes

# Teaching

Jan/20 – Mar/20 Grenoble (FR)	<b>Statistical Modeling and Literate Programming</b> 20 hours – 3 <sup>rd</sup> year students (L3) in business informatics. Data analysis and visualization in R.	Université Grenoble Alpes
Sep/19 – Dec/19 Grenoble (FR)	<b>Algorithmics and Imperative Programming</b> 30 hours – 3 <sup>rd</sup> year students (L3) in software engineer school. Computational complexity, correctness proof, data structures.	Polytech Grenoble
Sep/18 – Dec/18 Grenoble (FR)	<b>Principles of Operating Systems</b> 30 hours – 4 <sup>th</sup> year students (M1) in computer science. Memory allocation, multithreading, synchronization, buffered I/O, performance evaluation.	Université Grenoble Alpes
Sep/18 – Dec/18 Grenoble (FR)	<b>Software Development Basis</b> 30 hours – 2 <sup>nd</sup> year students (L2) in computer science. Functional & robustness testing, modularization, type abstraction.	Université Grenoble Alpes
Jan/18 – May/18 Grenoble (FR)	<b>Introduction to Python</b> 64 hours – 1 <sup>st</sup> year students (L1) in earth science. Variables and types, control flow statements, data structures, files.	Université Grenoble Alpes

# Publications

## Conference Articles

- [CLH19] Fast and Faithful Performance Prediction of MPI Applications: the HPL Case Study  
Cornebize, T.; Legrand, A., and Heinrich, F. C.  
URL: <https://hal.inria.fr/hal-02096571>  
*2019 IEEE International Conference on Cluster Computing (CLUSTER), 2019*
- [Hei+17] Predicting the Energy Consumption of MPI Applications at Scale Using a Single Node  
Heinrich, F. C.; Cornebize, T.; Degomme, A.; Legrand, A.; Carpen-Amarie, A.; Hunold, S.; Orgerie, A.-C., and Quinson, M.  
URL: <https://hal.inria.fr/hal-01523608>  
*2017 IEEE International Conference on Cluster Computing (CLUSTER), 2017*
- [PC17] Isolating Jobs for Security on High-Performance Fabrics  
Perotin, M. and Cornebize, T.  
*2017 IEEE 3rd International Workshop on High-Performance Interconnection Networks in the Exascale and Big-Data Era (HiPINEB), 2017*
- [FCF14] Efficient and Generalized Decentralized Monitoring of Regular Languages  
Falcone, Y.; Cornebize, T., and Fernandez, J.-C.  
URL: <https://hal.archives-ouvertes.fr/hal-00972559>  
*34th Formal Techniques for Networked and Distributed Systems (FORTE), 2014*

## Preprints

- [CL21] Simulation-based Optimization and Sensibility Analysis of MPI Applications: Variability Matters  
Cornebize, T. and Legrand, A. (submitted at JPDC, currently under review)  
URL: <https://hal.inria.fr/hal-03141988>
- [CL19] DGEMM performance is data-dependent  
Cornebize, T. and Legrand, A.  
URL: <https://hal.inria.fr/hal-02401760>

## Thesis

- [Cor21] High Performance Computing: Towards Better Performance Predictions and Experiments  
Cornebize, T.  
URL: [http://cornebize.net/doc/thesis\\_manuscript.pdf](http://cornebize.net/doc/thesis_manuscript.pdf)  
*June 2021*
- [Cor17] Capacity Planning of Supercomputers: Simulating MPI Applications at Scale  
Cornebize, T.  
URL: <https://hal.inria.fr/hal-01544827>  
*June 2017*

# Software Projects

Jul/19 – May/21	<b>G5K tests</b> Performance tests on Grid'5000. <ul style="list-style-type: none"><li>• On a near-daily basis, run micro-benchmarks on hundreds of Grid'5000 nodes.</li><li>• Statistical tests on several observed metrics (e.g. CPU performance, frequency, power consumption, temperature).</li><li>• Detected several events that affected significantly the nodes performance and went unnoticed by both the staff and the users (e.g. cooling issue, faulty memory, power instability, BIOS upgrade).</li></ul>	<a href="http://cornebize.net/g5k_test">cornebize.net/g5k_test</a>
Aug/18 – now	<b>Peanut</b> Experiment engine for Grid'5000. <ul style="list-style-type: none"><li>• Experiments written as readable and concise Python scripts.</li><li>• Common tasks (e.g. node deployment, resource monitoring) factorized in the engine.</li><li>• Experiment data and meta-data automatically collected in an archive.</li></ul>	<a href="https://github.com/Ezibenroc/peanut">github.com/Ezibenroc/peanut</a>
Mar/18 – now	<b>Pycewise</b> Python module to compute a segmented linear regression. Implemented using a two-phase greedy algorithm. Several objective functions are available: <ul style="list-style-type: none"><li>• Ordinary Least Square,</li><li>• Weighted Least Square,</li><li>• Logarithmic Least Square.</li></ul>	<a href="https://github.com/Ezibenroc/pycewise">github.com/Ezibenroc/pycewise</a>
Feb/17 – now	<b>Contribution to Simgrid</b> Simulation of distributed computer systems. Contributed to SMPI, the simulator of MPI applications. <ul style="list-style-type: none"><li>• Implemented several features to improve the simulation efficiency.</li><li>• Greatly improved the calibration procedure for modeling an existing platform.</li></ul>	<a href="http://simgrid.org">simgrid.org</a>
May/16 – now	<b>Contribution to Roaring bitmap</b> Fast and lightweight set for unsigned 32 bits integers. <ul style="list-style-type: none"><li>• Contributed to CRoaring, the C library.<ul style="list-style-type: none"><li>– Implemented range constructor, selection and subset queries.</li><li>– Fixed several bugs.</li><li>– Repository: <a href="https://github.com/roaringBitmap/CRoaring">github.com/roaringBitmap/CRoaring</a></li></ul></li><li>• Developed PyRoaring, a Python wrapper for the C library.<ul style="list-style-type: none"><li>– Similar API than the built-in Python set, but several orders of magnitude faster.</li><li>– Used the Cython programming language.</li><li>– Extensive tests caught several bugs of the C library.</li><li>– Repository: <a href="https://github.com/Ezibenroc/PyRoaringBitMap">github.com/Ezibenroc/PyRoaringBitMap</a></li></ul></li></ul>	<a href="http://roaringbitmap.org">roaringbitmap.org</a>
Sep/14 – Dec/14	<b>Platypus</b> Modular and open source question answering framework. <ul style="list-style-type: none"><li>• Developed a question parsing module in Python, with a grammatical approach (Stanford CoreNLP and NLTK libraries).</li><li>• Framework currently used and valorized by Lexistems SAS.</li></ul>	<a href="http://askplatyp.us">askplatyp.us</a>

# Invited Talks

May/21 Online	<b>Practical Approach to Experiment Management</b> Datamove Seminar, Inria Grenoble
Oct/20 Online	<b>Measuring the Duration of a Matrix Product is Harder than you Think</b> XUG meeting, Inria Rennes
Sep/20 Online	<b>Performance Prediction of High Performance Linpack: a Long Journey</b> PC2/HIT Seminar, Paderborn University
Sep/20 Online	<b>DGEMM Performance is Data-Dependent</b> 11 <sup>th</sup> JLESC Workshop
May/20 Online	<b>Detecting Changes on Grid'5000 Machines</b> Polaris Seminar, Inria Grenoble
Nov/19 Grenoble (FR)	<b>Quality Assurance of a HPC Cluster: Testing for Performance non-Regression</b> Journée des utilisateurs GRICAD