

RESEARCH CENTRE

Grenoble - Rhône-Alpes

IN PARTNERSHIP WITH:

CNRS, Ecole normale supérieure de Lyon,
Université Claude Bernard (Lyon 1)

2020

ACTIVITY REPORT

Project-Team

AVALON

**Algorithms and Software Architectures for
Distributed and HPC Platforms**

IN COLLABORATION WITH: Laboratoire de l'Informatique du
Parallélisme (LIP)

DOMAIN

**Networks, Systems and Services,
Distributed Computing**

THEME

**Distributed and High Performance
Computing**

Contents

Project-Team AVALON	1
1 Team members, visitors, external collaborators	3
2 Overall objectives	5
2.1 Presentation	5
2.2 Objectives	5
3 Research program	6
3.1 Energy Application Profiling and Modeling	6
3.2 Data-intensive Application Profiling, Modeling, and Management	6
3.3 Resource-Agnostic Application Description Model	7
3.4 Virtualization	7
3.5 Application Mapping and Scheduling	8
3.5.1 Application Mapping and Software Deployment	8
3.5.2 Non-Deterministic Workflow Scheduling	8
3.5.3 Software Asset Management	8
4 Application domains	8
4.1 Overview	8
4.2 Climatology	9
4.3 Astrophysics	9
4.4 Bioinformatics	9
5 New software and platforms	10
5.1 New software	10
5.1.1 DIET	10
5.1.2 SimGrid	10
5.1.3 libkomp	11
5.1.4 XKBLAS	12
5.1.5 Concerto	12
5.1.6 Qirinus-Orchestra	13
5.1.7 execo	13
5.1.8 Kwollect	13
5.2 New platforms	14
5.2.1 Platform: Grid'5000	14
5.2.2 Platform: Leco	14
5.2.3 Platform: SILECS	14
5.2.4 Platform: SLICES	14
6 New results	15
6.1 Energy Efficiency in HPC and Large Scale Distributed Systems	15
6.1.1 Predicting the Energy Consumption of CUDA Kernels using SimGrid	15
6.1.2 Energy Consumption and Energy Efficiency in a Precision Beekeeping System	15
6.1.3 Ecodesign of Digital Services	16
6.2 Modeling and Simulation of Parallel Applications and Distributed Infrastructures	16
6.3 Edge and Cloud Ressouce Management	17
6.4 HPC Applications and Runtimes	18
6.5 Improving Virtualized Systems	19
7 Bilateral contracts and grants with industry	20
7.1 Bilateral contracts with industry	20
7.2 Bilateral grants with industry	20

8 Partnerships and cooperations	21
8.1 International initiatives	21
8.1.1 Inria International Labs	21
8.1.2 Inria international partners	21
8.2 European initiatives	21
8.2.1 FP7 & H2020 Projects	21
8.3 National initiatives	25
8.3.1 Inria Large Scale Initiative	25
8.4 Regional initiatives	25
8.4.1 CPER	25
8.4.2 Action Exploratoire Inria	25
9 Dissemination	26
9.1 Promoting scientific activities	26
9.1.1 Scientific events: organisation	26
9.1.2 Scientific events: selection	26
9.1.3 Journal	27
9.1.4 Scientific expertise	27
9.1.5 Research administration	27
9.2 Teaching - Supervision - Juries	27
9.2.1 Teaching	27
9.2.2 Supervision	28
9.2.3 Juries	29
9.3 Popularization	30
9.3.1 Internal or external Inria responsibilities	30
9.3.2 Articles and contents	30
9.3.3 Interventions	30
10 Scientific production	30
10.1 Publications of the year	30
10.2 Cited publications	32

Project-Team AVALON

Creation of the Team: 2012 February 01, updated into Project-Team: 2014 July 01

Keywords

Computer sciences and digital sciences

- A1.1.1. – Multicore, Manycore
- A1.1.2. – Hardware accelerators (GPGPU, FPGA, etc.)
- A1.1.4. – High performance computing
- A1.1.5. – Exascale
- A1.1.13. – Virtualization
- A1.3.5. – Cloud
- A1.3.6. – Fog, Edge
- A1.6. – Green Computing
- A2.1.6. – Concurrent programming
- A2.1.7. – Distributed programming
- A2.1.10. – Domain-specific languages
- A2.5.2. – Component-based Design
- A2.6. – Infrastructure software
- A2.6.1. – Operating systems
- A2.6.2. – Middleware
- A2.6.3. – Virtual machines
- A2.6.4. – Ressource management
- A4.4. – Security of equipment and software
- A6.2.7. – High performance computing
- A7.1. – Algorithms
- A7.1.1. – Distributed algorithms
- A7.1.2. – Parallel algorithms
- A8.2.1. – Operations research
- A8.9. – Performance evaluation

Other research topics and application domains

- B1.1.7. – Bioinformatics
- B3.2. – Climate and meteorology
- B4.1. – Fossile energy production (oil, gas)
- B4.2.2. – Fusion
- B4.5. – Energy consumption
- B4.5.1. – Green computing
- B6.1.1. – Software engineering

B8.1.1. – Energy for smart buildings

B9.5.1. – Computer science

B9.7. – Knowledge dissemination

B9.7.1. – Open access

B9.7.2. – Open data

B9.8. – Reproducibility

1 Team members, visitors, external collaborators

Research Scientists

- Christian Perez [Team leader, Inria, Senior Researcher, HDR]
- Marcos Dias de Assunção [Inria, Starting Research Position, until Apr 2020]
- Thierry Gautier [Inria, Researcher, HDR]
- Laurent Lefevre [Inria, Researcher, HDR]

Faculty Members

- Yves Caniou [Univ Claude Bernard, Associate Professor]
- Eddy Caron [École Normale Supérieure de Lyon, Associate Professor, HDR]
- Olivier Glück [Univ Claude Bernard, Associate Professor]
- Elise Jeanneau [Univ Claude Bernard, Associate Professor, from Sep 2020]
- Etienne Mauffret [École Normale Supérieure de Lyon, ATER, from Sep 2020]
- Alain Tchana [École Normale Supérieure de Lyon, Professor, HDR]

Post-Doctoral Fellow

- Mohamed Karaoui [École Normale Supérieure de Lyon, from Mar 2020]

PhD Students

- Ghoshana Bista [Orange Labs, CIFRE]
- Dorra Boughzala [Inria, until Nov 2020]
- Vo Quoc Bao Bui [INP Toulouse, until Sep 2020]
- Arthur Chevalier [Orange Labs, CIFRE]
- Felipe Rodrigo De Souza [École Normale Supérieure de Lyon]
- Idriss Doudadi [INRIA, work at INRIA Bordeaux - co-advisor S. Thibault]
- Hugo Hadjur [aivancity School for Technology, Business & Society Paris-Cachan, from Sep 2020]
- Zeina Houmani [École Normale Supérieure de Lyon]
- Barbe Thystere Mvondo Djob [Univ Grenoble Alpes]
- Lucien Ndjie Ngale [Univ Jules Vernes Picardie, from Nov 2020]
- Celestine Stella Ndonga Bitchebe [CNRS]
- Kevin Nguetchouang [Univ de Lyon, from Dec 2020]
- Romain Pereira [CEA, from Nov 2020]
- Pierre Etienne Polet [Thales, CIFRE, from Jul 2020]
- Laurent Turpin [Inria]
- Patrick Lavoisier Wapet [INP Toulouse]

Technical Staff

- Simon Delamare [CNRS, Engineer]
- Marie Durand [Inria, Engineer, until Aug 2020]
- Zakaria Fraoui [Inria, Engineer]
- Matthieu Imbert [Inria, Engineer]
- Patrice Kirmizigul [Inria, Engineer, from Dec 2020]
- Vincent Lanore [Inria, Engineer, until Sep 2020]
- David Loup [Inria, Engineer]
- Jean-Christophe Mignot [CNRS, Engineer]
- Olivier Mornard [Inria, Engineer, from Jun 2020 until Oct 2020]
- Cyril Seguin [Inria, Engineer, until Jun 2020]

Interns and Apprentices

- Mohamed Bassiouny [Inria, from May 2020 until Jul 2020]
- Adrien Berthelot [Inria, until Sep 2020]
- Jerome Boillot [École Normale Supérieure de Lyon, from Apr 2020 until Jul 2020]
- Matteo Delabre [Inria, from Feb 2020 until Jun 2020]
- Oregane Desrentes [École Normale Supérieure de Lyon, from Apr 2020 until Jul 2020]
- Theophile Dubuc [École Normale Supérieure de Lyon]
- Meriem Ghali [Inria, from Jun 2020 until Jul 2020]
- Haci Yusuf Gundogan [Inria, from Jun 2020 until Jul 2020]
- Yves Kone [École Normale Supérieure de Lyon, from Mar 2020 until Aug 2020]
- Nils Moynac [Inria, from Jun 2020 until Jul 2020]
- Lucien Ndjie Ngale [École Normale Supérieure de Lyon, from Apr 2020 until Oct 2020]
- Kevin Nguetchouang [École Normale Supérieure de Lyon, from Apr 2020 until Oct 2020]
- Stephane Pouget [École Normale Supérieure de Lyon, from Mar 2020 until Aug 2020]
- Raoufdine Said [Inria, from Jun 2020 until Jul 2020]
- Gaspard Thevenon [École Normale Supérieure de Lyon, from Jun 2020 until Jul 2020]

Administrative Assistant

- Evelyne Blesle [Inria]

Visiting Scientists

- Jean-Philippe Aboumou [Saham Life Insurance - Cameroun, until Feb 2020]
- Daniel Ndjodo Bessala [Université de Yaoundé - Cameroun, from Oct 2020 until Nov 2020]

External Collaborators

- Doreid Ammar [aivancity School for Technology, Business & Society Paris-Cachan, from Sep 2020, Professor]
- Frédéric Suter [CNRS, Researcher, HDR]

2 Overall objectives

2.1 Presentation

The fast evolution of hardware capabilities in terms of wide area communication, computation and machine virtualization leads to the requirement of another step in the abstraction of resources with respect to parallel and distributed applications. These large scale platforms based on the aggregation of large clusters (Grids), huge datacenters (Clouds) with IoT (Edge/Fog), collections of volunteer PCs (Desktop computing platforms), or high performance machines (Supercomputers) are now available to researchers of different fields of science as well as to private companies. This variety of platforms and the way they are accessed also have an important impact on how applications are designed (*i.e.*, the programming model used) as well as how applications are executed (*i.e.*, the runtime/middleware system used). The access to these platforms is driven through the use of multiple services providing mandatory features such as security, resource discovery, virtualization, load-balancing, monitoring, *etc.*

The goal of the AVALON team is to execute parallel and/or distributed applications on parallel and/or distributed resources while ensuring user and system objectives with respect to performance, cost, energy, security, *etc.* Users are generally not interested in the resources used during the execution. Instead, they are interested in how their application is going to be executed: the duration, its cost, the environmental footprint involved, *etc.* This vision of utility computing has been strengthened by the cloud concepts and by the short lifespan of supercomputers (around three years) compared to application lifespan (tens of years). Therefore a major issue is to design models, systems, and algorithms to execute applications on resources while ensuring user constraints (price, performance, *etc.*) as well as system administrator constraints (maximizing resource usage, minimizing energy consumption, *etc.*).

2.2 Objectives

To achieve the vision proposed in Section 2.1, the AVALON project aims at making progress on four complementary research axes: energy, data, programming models and runtimes, application scheduling and virtualization.

Energy Application Profiling and Modeling AVALON will improve the profiling and modeling of scientific applications with respect to energy consumption. In particular, it will require to improve the tools that measure the energy consumption of applications, virtualized or not, at large scale, so as to build energy consumption models of applications.

Data-intensive Application Profiling, Modeling, and Management AVALON will improve the profiling, modeling, and management of scientific applications with respect to CPU and data intensive applications. Challenges are to improve the performance prediction of parallel regular applications, to model and simulate (complex) intermediate storage components, and data-intensive applications, and last to deal with data management for hybrid computing infrastructures.

Programming Models AVALON will design component-based models to capture the different facets of parallel and distributed applications while being resource agnostic, so that they can be optimized for a particular execution. In particular, the proposed component models will integrate energy and data modeling results. AVALON in particular targets OpenMP runtime as a specific use case and contributes to improve it for multi-GPU nodes.

Application Mapping and Scheduling AVALON will propose multi-criteria mapping and scheduling algorithms to meet the challenge of automating the efficient utilization of resources taking into consideration criteria such as performance (CPU, network, and storage), energy consumption, and security. AVALON will in particular focus on application deployment, workflow applications, and security management in clouds.

Virtualization Virtualization is a powerful approach to abstract resources. However, many challenges remain at the hypervisor level to provide an efficient usage of resources, in particular for disaggregated data centers.

All our theoretical results will be validated with software prototypes using applications from different fields of science such as bioinformatics, physics, cosmology, *etc.* The experimental testbeds GRID'5000, Leco, and Silecs will be our platforms of choice for experiments.

3 Research program

3.1 Energy Application Profiling and Modeling

Despite recent improvements, there is still a long road to follow in order to obtain energy efficient, energy proportional and eco-responsible exascale systems. Energy efficiency is therefore a major challenge for building next generation large-scale platforms. The targeted platforms will gather hundreds of millions of cores, low power servers, or CPUs. Besides being very important, their power consumption will be dynamic and irregular.

Thus, to consume energy efficiently, we aim at investigating two research directions. First, we need to improve measurement, understanding, and analysis on how large-scale platforms consume energy. Unlike some approaches [23] that mix the usage of internal and external wattmeters on a small set of resources, we target high frequency and precise internal and external energy measurements of each physical and virtual resource on large-scale distributed systems.

Secondly, we need to find new mechanisms that consume less and better on such platforms. Combined with hardware optimizations, several works based on shutdown or slowdown approaches aim at reducing energy consumption of distributed platforms and applications. To consume less, we first plan to explore the provision of accurate estimation of the energy consumed by applications without pre-executing and knowing them while most of the works try to do it based on in-depth application knowledge (code instrumentation [26], phase detection for specific HPC applications [30], *etc.*). As a second step, we aim at designing a framework model that allows interaction, dialogue and decisions taken in cooperation among the user/application, the administrator, the resource manager, and the energy supplier. While smart grid is one of the last killer scenarios for networks, electrical provisioning of next generation large IT infrastructures remains a challenge.

3.2 Data-intensive Application Profiling, Modeling, and Management

The term “Big Data” has emerged to design data sets or collections so large that they become intractable for classical tools. This term is most of the time implicitly linked to “analytics” to refer to issues such as data curation, storage, search, sharing, analysis, and visualization. However, the Big Data challenge is not limited to data-analytics, a field that is well covered by programming languages and run-time systems such as Map-Reduce. It also encompasses data-intensive applications. These applications can be sorted into two categories. In High Performance Computing (HPC), data-intensive applications leverage post-petascale infrastructures to perform highly parallel computations on large amount of data, while in High Throughput Computing (HTC), a large amount of independent and sequential computations are performed on huge data collections.

These two types of data-intensive applications (HTC and HPC) raise challenges related to profiling and modeling that the AVALON team proposes to address. While the characteristics of data-intensive applications are very different, our work will remain coherent and focused. Indeed, a common goal will be to acquire a better understanding of both the applications and the underlying infrastructures running them to propose the best match between application requirements and infrastructure capacities. To

achieve this objective, we will extensively rely on logging and profiling in order to design sound, accurate, and validated models. Then, the proposed models will be integrated and consolidated within a single simulation framework (SIMGRID). This will allow us to explore various potential “what-if?” scenarios and offer objective indicators to select interesting infrastructure configurations that match application specificities.

Another challenge is the ability to mix several heterogeneous infrastructures that scientists have at their disposal (*e.g.*, Grids, Clouds, and Desktop Grids) to execute data-intensive applications. Leveraging the aforementioned results, we will design strategies for efficient data management service for hybrid computing infrastructures.

3.3 Resource-Agnostic Application Description Model

With parallel programming, users expect to obtain performance improvement, regardless its cost. For long, parallel machines have been simple enough to let a user program use them given a minimal abstraction of their hardware. For example, MPI [25] exposes the number of nodes but hides the complexity of network topology behind a set of collective operations; OpenMP [29] simplifies the management of threads on top of a shared memory machine while OpenACC [28] aims at simplifying the use of GPGPU.

However, machines and applications are getting more and more complex so that the cost of manually handling an application is becoming very high [24]. Hardware complexity also stems from the unclear path towards next generations of hardware coming from the frequency wall: multi-core CPU, many-core CPU, GPGPUs, deep memory hierarchy, *etc.* have a strong impact on parallel algorithms. Parallel languages (UPC, Fortress, X10, *etc.*) can be seen as a first piece of a solution. However, they will still face the challenge of supporting distinct codes corresponding to different algorithms corresponding to distinct hardware capacities.

Therefore, the challenge we aim to address is to define a model, for describing the structure of parallel and distributed applications that enables code variations but also efficient executions on parallel and distributed infrastructures. Indeed, this issue appears for HPC applications but also for cloud oriented applications. The challenge is to adapt an application to user constraints such as performance, energy, security, *etc.*

Our approach is to consider component based models [31] as they offer the ability to manipulate the software architecture of an application. To achieve our goal, we consider a “compilation” approach that transforms a resource-agnostic application description into a resource-specific description. The challenge is thus to determine a component based model that enables to efficiently compute application mapping while being tractable. In particular, it has to provide an efficient support with respect to application and resource elasticity, energy consumption and data management. OpenMP runtime is a specific use case that we target.

3.4 Virtualization

Hypervisors are a major building block of cloud computing: they are the cornerstone of Infrastructure-as-a-Service (IaaS) platforms, which are themselves used to implement other cloud models such as Containers-as-a-service (CaaS), Functions-as-a-Service (FaaS) and (some types of) Platform-as-a-Service (PaaS). Moreover, the design of data centers is expected to evolve from a monolithic server centric model to a disaggregated architecture (resource centric), in order to better address concerns such as resource utilization (energy saving), elasticity and failure management: the current set of designs and implementations is quite diverse but a common characteristic among them is that elementary resources (such as CPUs, memory banks, network cards, and disks) will be managed as a set of independent components (each of them embedding an operating system as “firmware”) interconnected by a large-scale network fabric. This paradigm change has fundamental ramifications regarding the design of hypervisors, raising additional challenges in terms of performance (for both the control plane and the data plane) and security. In particular, one of the main consequences of this new, distributed model is that the current architecture of hypervisors will necessarily become hierarchical, with resource management logic spread into multiple components.

3.5 Application Mapping and Scheduling

This research axis is at the crossroad of the AVALON team. In particular, it gathers results of the other research axis. We plan to consider application mapping and scheduling addressing the following three issues.

3.5.1 Application Mapping and Software Deployment

Application mapping and software deployment consist in the process of assigning distributed pieces of software to a set of resources. Resources can be selected according to different criteria such as performance, cost, energy consumption, security management, *etc.* A first issue is to select resources at application launch time. With the wide adoption of elastic platforms, *i.e.*, platforms that let the number of resources allocated to an application to be increased or decreased during its execution, the issue is also to handle resource selection at runtime.

The challenge in this context corresponds to the mapping of applications onto distributed resources. It will consist in designing algorithms that in particular take into consideration application profiling, modeling, and description.

A particular facet of this challenge is to propose scheduling algorithms for dynamic and elastic platforms. As the number of elements can vary, some kind of control of the platforms must be used accordingly to the scheduling.

3.5.2 Non-Deterministic Workflow Scheduling

Many scientific applications are described through workflow structures. Due to the increasing level of parallelism offered by modern computing infrastructures, workflow applications now have to be composed not only of sequential programs, but also of parallel ones. New applications are now built upon workflows with conditionals and loops (also called non-deterministic workflows).

These workflows cannot be scheduled beforehand. Moreover cloud platforms bring on-demand resource provisioning and pay-as-you-go billing models. Therefore, there is a problem of resource allocation for non-deterministic workflows under budget constraints and using such an elastic management of resources.

Another important issue is data management. We need to schedule the data movements and replications while taking job scheduling into account. If possible, data management and job scheduling should be done at the same time in a closely coupled interaction.

3.5.3 Software Asset Management

The use of software is generally regulated by licenses, whether they are free or paid and with or without access to their sources. The world of licenses is very vast and unknown (especially in the industrial world). Often only the general public version is known (a software purchase corresponds to a license). For enterprises, the reality is much more complex, especially for main publishers. We work on the OpTISAM software, a software offering tools to perform Software Asset Management (SAM) much more efficiently in order to be able to ensure the full compliance with all contracts from each software and a new type of deployment taking into account these aspects and other additional parameters like energy and performance. This work is built on an OrangeTM collaboration.

4 Application domains

4.1 Overview

The AVALON team targets applications with large computing and/or data storage needs, which are still difficult to program, deploy, and maintain. Those applications can be parallel and/or distributed applications, such as large scale simulation applications or code coupling applications. Applications can also be workflow-based as commonly found in distributed systems such as grids or clouds.

The team aims at not being restricted to a particular application field, thus avoiding any spotlight. The team targets different HPC and distributed application fields, which brings use cases with different

issues. This will be eased with our participation to the Joint Laboratory for Extreme Scale Computing (JLESC), and to BioSyL, a federative research structure about Systems Biology of the University of Lyon. Moreover, the team members have a long tradition of cooperation with application developers such as CERFACS and EDF R&D. Last but not least, the team has a privileged connection with CC-IN2P3 that opens up collaborations, in particular in the astrophysics field.

In the following, some examples of representative applications that we are targeting are presented. In addition to highlighting some application needs, they also constitute some of the use cases that will be used to validate our theoretical results.

4.2 Climatology

The world's climate is currently changing due to the increase of the greenhouse gases in the atmosphere. Climate fluctuations are forecasted for the years to come. For a proper study of the incoming changes, numerical simulations are needed, using general circulation models of a climate system. Simulations can be of different types: HPC applications (e.g., the NEMO framework [27] for ocean modelization), code-coupling applications (e.g., the OASIS coupler [32] for global climate modeling), or workflows (long term global climate modeling).

As for most applications the team is targeting, the challenge is to thoroughly analyze climate-forecasting applications to model their needs in terms of programming model, execution model, energy consumption, data access pattern, and computing needs. Once a proper model of an application has been set up, appropriate scheduling heuristics can be designed, tested, and compared. The team has a long tradition of working with CERFACS on this topic, for example in the LEGO (2006-09) and SPADES (2009-12) French ANR projects.

4.3 Astrophysics

Astrophysics is a major field to produce large volumes of data. For instance, the Vera C. Rubin Observatory (<https://www.vro.org/>) will produce 20 TB of data every night, with the goals of discovering thousands of exoplanets and of uncovering the nature of dark matter and dark energy in the universe. The Square Kilometer Array (<http://www.skatelescope.org/>) produces 9 Tbits/s of raw data. One of the scientific projects related to this instrument called Evolutionary Map of the Universe is working on more than 100 TB of images. The Euclid Imaging Consortium (<https://www.euclid-ec.org/>) will generate 1 PB data per year.

AVALON collaborates with the *Institut de Physique des deux Infinis de Lyon* (IP2I) laboratory on large scale numerical simulations in astronomy and astrophysics. Contributions of the AVALON members have been related to algorithmic skeletons to demonstrate large scale connectivity, the development of procedures for the generation of realistic mock catalogs, and the development of a web interface to launch large cosmological simulations on GRID'5000.

This collaboration, that continues around the topics addressed by the CLUES project (<http://www.clues-project.org>), has been extended thanks to the tight links with the CC-IN2P3. Major astrophysics projects execute part of their computing, and store part of their data on the resources provided by the CC-IN2P3. Among them, we can mention SNFactory, Euclid, or VRO. These applications constitute typical use cases for the research developed in the AVALON team: they are generally structured as workflows and a huge amount of data (from TB to PB) is involved.

4.4 Bioinformatics

Large-scale data management is certainly one of the most important applications of distributed systems in the future. Bioinformatics is a field producing such kinds of applications. For example, DNA sequencing applications make use of MapReduce skeletons.

The AVALON team is a member of BioSyL (<http://www.biosyl.org>), a Federative Research Structure attached to University of Lyon. It gathers about 50 local research teams working on systems biology. Moreover, the team cooperated with the French Institute of Biology and Chemistry of Proteins (IBCP <http://www.ibcp.fr>) in particular through the ANR MapReduce project where the team focuses on a bio-chemistry application dealing with protein structure analysis. AVALON has also started working with

the Inria Beagle team (<https://team.inria.fr/beagle/>) on artificial evolution and computational biology as the challenges are around high performance computation and data management.

5 New software and platforms

5.1 New software

5.1.1 DIET

Name: Distributed Interactive Engineering Toolbox

Keywords: Scheduling, Clusters, Grid, Cloud, HPC, Middleware, Data management.

Functional Description: Middleware for grids and clouds. Toolbox for the use and porting of intensive computing applications on heterogeneous architectures.

Release Contributions: - Native Google Drive Support for the data manager - Standardization of internal integer types. - New types (see Changelog for more information)

News of the Year: We have worked on a new Budget-aware workflow scheduling with DIET We have proposed and extended MADAG for static scheduling (<https://hal.inria.fr/hal-03080468>)

URL: <http://graal.ens-lyon.fr/diet/>

Publication: [hal-03080468](https://hal.inria.fr/hal-03080468)

Contact: Eddy Caron

Participants: Joel Faubert, Hadrien Croubois, Abdelkader Amar, Arnaud Lefray, Aurélien Bouteiller, Benjamin Isnard, Daniel Balouek, Eddy Caron, Eric Bois, Frédéric Desprez, Frédéric Lombart, Gaël Le Mahec, Guillaume Verger, Huaxi Zhang, Jean-Marc Nicod, Jonathan Rouzaud-Cornabas, Lamiel Toch, Maurice Faye, Peter Frauenkron, Philippe Combes, Philippe Laurent, Raphaël Bolze, Yves Caniou, Cyril Seguin, Aurélie Kong Win Chang

Partners: CNRS, ENS Lyon, UCBL Lyon 1, Sysfera

5.1.2 SimGrid

Keywords: Large-scale Emulators, Grid Computing, Distributed Applications

Scientific Description: SimGrid is a toolkit that provides core functionalities for the simulation of distributed applications in heterogeneous distributed environments. The simulation engine uses algorithmic and implementation techniques toward the fast simulation of large systems on a single machine. The models are theoretically grounded and experimentally validated. The results are reproducible, enabling better scientific practices.

Its models of networks, cpus and disks are adapted to (Data)Grids, P2P, Clouds, Clusters and HPC, allowing multi-domain studies. It can be used either to simulate algorithms and prototypes of applications, or to emulate real MPI applications through the virtualization of their communication, or to formally assess algorithms and applications that can run in the framework.

The formal verification module explores all possible message interleavings in the application, searching for states violating the provided properties. We recently added the ability to assess liveness properties over arbitrary and legacy codes, thanks to a system-level introspection tool that provides a finely detailed view of the running application to the model checker. This can for example be leveraged to verify both safety or liveness properties, on arbitrary MPI code written in C/C++/Fortran.

Functional Description: SimGrid is a toolkit that provides core functionalities for the simulation of distributed applications in heterogeneous distributed environments. The simulation engine uses algorithmic and implementation techniques toward the fast simulation of large systems on a single machine. The models are theoretically grounded and experimentally validated. The results are reproducible, enabling better scientific practices.

Its models of networks, cpus and disks are adapted to (Data)Grids, P2P, Clouds, Clusters and HPC, allowing multi-domain studies. It can be used either to simulate algorithms and prototypes of applications, or to emulate real MPI applications through the virtualization of their communication, or to formally assess algorithms and applications that can run in the framework.

The formal verification module explores all possible message interleavings in the application, searching for states violating the provided properties. We recently added the ability to assess liveness properties over arbitrary and legacy codes, thanks to a system-level introspection tool that provides a finely detailed view of the running application to the model checker. This can for example be leveraged to verify both safety or liveness properties, on arbitrary MPI code written in C/C++/Fortran.

News of the Year: There were 2 major releases in 2020. SMPI is now regularly tested on medium scale benchmarks of the exascale suite. The Wifi support was improved, through more example and documentation, and an energy model of wifi links was proposed. Many bugs were fixed in the bindings to the ns-3 packet-level network simulator, which now allows to simulate Wifi links using ns-3 too. We enriched the API expressiveness to allow the construction of activity tasks. We also pursued our efforts to improve the documentation of the software, simplified the web site, and made a lot of bug fixing and code refactoring.

URL: <https://simgrid.org/>

Contact: Martin Quinson

Participants: Adrien Lèbre, Anne-Cécile Orgerie, Arnaud Legrand, Augustin Degomme, Emmanuelle Saillard, Frédéric Suter, Jean-Marc Vincent, Jonathan Pastor, Luka Stanisic, Martin Quinson, Samuel Thibault

Partners: CNRS, ENS Rennes

5.1.3 libkomp

Name: Runtime system libkomp

Keywords: HPC, Multicore, OpenMP

Functional Description: libKOMP is a runtime support for OpenMP compatible with différent compiler: GNU gcc/gfortran, Intel icc/ifort or clang/llvm. It is based on source code initially developed by Intel for its own OpenMP runtime, with extensions from Kaapi software (task representation, task scheduling). Moreover it contains an OMPT module for recording trace of execution.

Release Contributions: Initial version

News of the Year: libKOMP is supported by EoCoE-II project. Tikki, an OMPT monitoring tools was extracted from libKOMP to be reused outside libKOMP (<https://gitlab.inria.fr/openmp/tikki>).

URL: <http://gitlab.inria.fr/openmp/libkomp>

Contact: Thierry Gautier

5.1.4 XKBLAS

Name: XKBLAS

Keywords: BLAS, Dense linear algebra, GPU

Functional Description: XKBLAS is yet another BLAS library (Basic Linear Algebra Subroutines) that targets multi-GPUs architecture thanks to the XKaapi runtime and with block algorithms from PLASMA library. XKBLAS is able to exploit large multi-GPUs node with sustained high level of performance. The library offers a wrapper library able to capture calls to BLAS (C or Fortran). The internal API is based on asynchronous invocations in order to enable overlapping between communication by computation and also to better composed sequences of calls to BLAS.

This current version of XKBlas is the first public version and contains only BLAS level 3 algorithms, including XGEMMT:

XGEMM XGEMMT: see MKL GEMMT interface XTRSM XTRMM XSYMM XSYRK XSYR2K XHEMM XHERK XHER2K

For classical precision Z, C, D, S.

Release Contributions: 0.1 versions: calls to BLAS kernels must be initiated by the same thread that initializes the XKBlas library. 0.2 versions: better support for libblas_wrapper and improved scheduling heuristic to take into account memory hierarchy between GPUs

News of the Year: MUMPS software runs natively on top of the XKBLAS library and obtains the best performances on multi-GPUs systems with XKBLAS.

URL: <https://gitlab.inria.fr/xkblas/versions>

Contact: Thierry Gautier

Participants: Thierry Gautier, João Vicente Ferreira Lima

5.1.5 Concerto

Name: Concerto

Keywords: Reconfiguration, Distributed Software, Component models, Dynamic software architecture

Scientific Description: Concerto is a reconfiguration model which allows to describe distributed software as an evolving assembly of components.

Functional Description: Concerto is an implementation of the formal model Concerto written in Python. Concerto allows to : 1. describe the life-cycle and the dependencies of software components, 2. describe a components assembly that forms the overall life-cycle of a distributed software, 3. automatically reconfigure a Concerto assembly of components by using a set of reconfiguration instructions as well as a formal operational semantics.

News of the Year: In 2020, we added the ability to read from and write to non-data provide ports. We also updated the Madeus wrapper to the new version with only USE and PROVIDE ports and no groups.

URL: <https://gitlab.inria.fr/VerDi-project/concerto>

Publications: [hal-03103714](#), [hal-02535077](#), [hal-01897803](#)

Contact: H el ene Coullon

Partners: IMT Atlantique, LS2N, LIP

5.1.6 Qirinus-Orchestra

Keywords: Automatic deployment, Cybersecurity

Functional Description: IQ-orchestra (previously Qirinus-Orchestra) is a meta-modeling software dedicated to the securized deployment of virtualized infrastructures.

It is built around three main paradigmes:

1 - Modelization of a catalog of supported application 2 - A dynamic securized architecture 3 - An automatic virtualized environment Deployment (i.e. Cloud)

The software is strongly modular and uses advanced software engineering tools such as meta-modeling. It will be continuously improved along 3 axes:

* The catalog of supported applications (open source, legacy, internal). * The catalog of security devices (access control, network security, component reinforcement, etc.) * Intelligent functionalities (automatic firewalls configuration, detection of non-secure behaviors, dynamic adaptation, etc.)

News of the Year: - Upgrade of IQ-Orchestra/IQ-Manager - Update of all old software embedded - New workflow compilation - Bugs fix - User guide v0.1

Publications: [hal-00840734](#), [hal-01355681](#), [tel-01229874](#)

Contact: Eddy Caron

Participants: Eddy Caron, Arthur Chevalier, Patrice Kirmizigul, Arnaud Lefray

5.1.7 execo

Keywords: Toolbox, Deployment, Orchestration, Python

Functional Description: Execo offers a Python API for asynchronous control of local or remote, standalone or parallel, unix processes. It is especially well suited for quickly and easily scripting workflows of parallel/distributed operations on local or remote hosts: automate a scientific workflow, conduct computer science experiments, perform automated tests, etc. The core python package is execo. The execo_g5k package provides a set of tools and extensions for the Grid5000 testbed. The execo_engine package provides tools to ease the development of computer sciences experiments.

Release Contributions: - misc python3 support fixes - basic documentation for wheezy compatible package build - remove some debug outputs - fix crash in processes conductor in some situations - improve/fix process stdout/stderr handlers - fix get_cluster_network equipments - add a FAQ

URL: <http://execo.gforge.inria.fr>

Contact: Matthieu Imbert

Participants: Florent Chuffart, Laurent Pouilloux, Matthieu Imbert

5.1.8 Kwollect

Keywords: Monitoring, Power monitoring, Energy, Infrastructure software

Functional Description: Kwollect is a monitoring framework for IT infrastructures. It focuses on collecting environmental metrics (energy, sensors, etc.) and make them available to users.

News of the Year: Since June 2020, Kwollect is available under Grid'5000, as a testing phase. It is intended to supersede other existing monitoring solution on the infrastructure.

An article describing Kwollect is being published in CNERT 2021 workshop (Workshop on Computer and Networking Experimental Research using Testbeds, in conjunction with IEEE INFOCOM 2021).

URL: <https://gitlab.inria.fr/grid5000/kwollect>

Contact: Simon Delamare

5.2 New platforms

5.2.1 Platform: Grid'5000

Participants Simon Delamare, Laurent Lefèvre, David Loup, Olivier Mornard, Christian Perez.

FUNCTIONAL DESCRIPTION

The Grid'5000 experimental platform is a scientific instrument to support computer science research related to distributed systems, including parallel processing, high performance computing, cloud computing, operating systems, peer-to-peer systems and networks. It is distributed on 10 sites in France and Luxembourg, including Lyon. Grid'5000 is a unique platform as it offers to researchers many and varied hardware resources and a complete software stack to conduct complex experiments, ensure reproducibility and ease understanding of results. In 2020, a new generation of high speed wattmeters has been deployed on the Lyon site. They allow energy monitoring with up to 50 measurements per second. In parallel, a new version of kwapi (software stack for energy monitoring) called kwollect has been proposed and redesigned.

- Contact: Laurent Lefèvre
- URL: <https://www.grid5000.fr/>

5.2.2 Platform: Leco

Participants Thierry Gautier, Laurent Lefèvre, Christian Perez.

FUNCTIONAL DESCRIPTION

The LECO experimental platform is a new medium size scientific instrument funded by DRRT to investigate research related to BigData and HPC. It is located in Grenoble as part of the the HPCDA computer managed by UMS GRICAD. The platform has been deployed in 2018 and was available for experiment since the summer. All the nodes of the platform are instrumented to capture the energy consumption and data are available through the Kwapi software.

- Contact: Thierry Gautier

5.2.3 Platform: SILECS

Participants Laurent Lefèvre, Simon Delamare, Christian Perez.

FUNCTIONAL DESCRIPTION The SILECS infrastructure (IR ministère) aims at providing an experimental platform for experimental computer Science (Internet of things, clouds, HPC, big data, *etc.*). This new infrastructure is based on two existing infrastructures, Grid'5000 and FIT.

- Contact: Laurent Lefevre
- URL: <https://www.silecs.net/>

5.2.4 Platform: SLICES

Participants Laurent Lefèvre, Christian Perez.

FUNCTIONAL DESCRIPTION SLICES is an European effort that aims at providing a flexible platform designed to support large-scale, experimental research focused on networking protocols, radio technologies, services, data collection, parallel and distributed computing and in particular cloud and edge-based computing architectures and services. The French node will leverage the SILECS platform.

- Contact: Christian Perez
- URL: <https://www.slices-ri.eu>

6 New results

6.1 Energy Efficiency in HPC and Large Scale Distributed Systems

6.1.1 Predicting the Energy Consumption of CUDA Kernels using SimGrid

Participants Laurent Lefèvre, Dorra Boughzala.

Building a sustainable Exascale machine is a very promising target in High Performance Computing (HPC). To tackle the energy consumption challenge while continuing to provide tremendous performance, the HPC community have rapidly adopted GPU-based systems. Today, GPUs have become the most prevailing components in the massively parallel HPC landscape thanks to their high computational power and energy efficiency. Modeling the energy consumption of applications running on GPUs has gained a lot of attention for the last years. Alas, the HPC community lacks simple yet accurate simulators to predict the energy consumption of general purpose GPU applications. In this work, we address the prediction of the energy consumption of CUDA kernels via simulation. We propose in this paper a simple and lightweight energy model that we implemented using the open-source framework SimGrid. Our proposed model [5] is validated across a diverse set of CUDA kernels and on two different NVIDIA GPUs (Tesla M2075 and Kepler K20Xm). As our modeling approach is not based on performance counters or detailed-architecture parameters, we believe that our model can be easily approved by users who take care of the energy consumption of their GPGPU applications.

6.1.2 Energy Consumption and Energy Efficiency in a Precision Beekeeping System

Participants Laurent Lefèvre, Doreid Ammar, Hugo Hadjur.

Honey bees have been domesticated by humans for several thousand years and mainly provide honey and pollination, which is fundamental for plant reproduction. Nowadays, the work of beekeepers is constrained by external factors that stress their production (parasites and pesticides, among others). Taking care of large numbers of beehives is time-consuming, so integrating sensors to track their status can drastically simplify the work of beekeepers. Precision beekeeping complements beekeepers' work thanks to the Internet of Things (IoT) technology. If used correctly, data can help to make the right diagnosis for honey bees colony, increase honey production and decrease bee mortality. Providing enough energy for on-hive and in-hive sensors is a challenge. Some solutions rely on energy harvesting, others target usage of large batteries. Either way, it is mandatory to analyze the energy usage of embedded equipment in order to design an energy efficient and autonomous bee monitoring system. Our first work, with the beginning of the Ph.D. of Hugo Hadjur (co-advised by Doreid Ammar (Academic Dean and Professor at aivancity School for Technology, Business & Society Paris-Cachan and external member of Avalon team) and Laurent Lefevre), relies on a fully autonomous IoT framework that collects environmental and image data of a beehive[9]. It consists of a data collecting node (environmental data sensors, camera, Raspberry Pi and Arduino) and a solar energy supplying node. Supported services are analyzed task by task from an energy profiling and efficiency standpoint, in order to identify the highly pressured areas of the framework. This first step will guide our goal of designing a sustainable precision beekeeping

system, both technically and energy-wise. Some experimental parts of this work occur in the CPER LECO/GreenCube project and some parts are financially supported by aivancity School for Technology, Business & Society Paris-Cachan.

6.1.3 Ecodesign of Digital Services

Participants Laurent Lefèvre.

Creating energy aware with limited environmental impacts digital services needs a complete redesign. Laurent Lefevre with some colleagues from the GDS EcoInfo group have explored the various facets of ecodesign. This has resulted to a brochure available for software developers. This brochure has been downloaded several hundred of times since its publication in November 2020.[22]

6.2 Modeling and Simulation of Parallel Applications and Distributed Infrastructures

Developing Accurate and Scalable Simulators of Production Workflow Management Systems

Participants Frédéric Suter.

WRENCH [2] is a Workflow Management System simulation framework, whose objectives are (i) accurate and scalable simulations; and (ii) easy simulation software development. WRENCH achieves its first objective by building on the SimGrid framework. While SimGrid is recognized for the accuracy and scalability of its simulation models, it only provides low-level simulation abstractions and thus large software development efforts are required when implementing simulators of complex systems. WRENCH thus achieves its second objective by providing high-level and directly re-usable simulation abstractions on top of SimGrid.

Characterizing, Modeling, and Accurately Simulating Power and Energy Consumption of I/O-intensive Scientific Workflows

Participants Frédéric Suter.

While distributed computing infrastructures can provide infrastructure-level techniques for managing energy consumption, application-level energy consumption models have also been developed to support energy-efficient scheduling and resource provisioning algorithms. We analyzed the accuracy of a widely-used application-level model that has been developed and used in the context of scientific workflow executions. To this end, we profiled two production scientific workflows on a distributed platform instrumented with power meters. We then conduct an analysis of power and energy consumption measurements. This analysis shows that power consumption is not linearly related to CPU utilization and that I/O operations significantly impact power, and thus energy, consumption. We then propose a power consumption model that accounts for I/O operations, including the impact of waiting for these operations to complete, and for concurrent task executions on multi-socket, multi-core compute nodes. We implement our proposed model as part of a simulator that allows us to draw direct comparisons between real-world and modeled power and energy consumption. We find that our model has high accuracy when compared to real-world executions. Furthermore, our model improves accuracy by about two orders of magnitude when compared to the traditional models used in the energy-efficient workflow scheduling literature [4].

Microservice Architectures using Data-Driven Discovery and QoS Guarantees

Participants Eddy Caron, Zeina Houmani.

Microservices promise the benefits of services with an efficient granularity using dynamically allocated resources. In the current evolving architectures, data producers and consumers are created as decoupled components that support different data objects and quality of service. Actual implementations of service meshes lack support for data-driven paradigms, and focus on goal-based approaches designed to fulfill the general system goal. This diversity of available components demands the integration of users requirements and data products into the discovery mechanism. We have proposed a data-driven service discovery framework based on profile matching using data-centric service descriptions. In [10] we have designed and evaluated a microservices architecture for providing service meshes with a standalone set of components that manages data profiles and resources allocations over multiple geographical zones. Moreover, we demonstrated an adaptation scheme to provide quality of service guarantees. Evaluation of the implementation on a real life testbed shows effectiveness of this approach with stable and fluctuating request incoming rates.

6.3 Edge and Cloud Ressouce Management

Safe and Efficient Reconfiguration with Concerto

Participants Christian Perez, Maverick Chardet.

For large-scale distributed systems that need to adapt to a changing environment, conducting a reconfiguration is a challenging task. In particular, efficient reconfigurations require the coordination of multiple tasks with complex dependencies. In [3, 6, 20], we present Concerto, a model used to manage the lifecycle of software components and coordinate their reconfiguration operations. Concerto promotes efficiency with a fine-grained representation of dependencies and parallel execution of reconfiguration actions, both within components and between them. In these works, the elements of the model are described as well as their formal semantics. In addition, we outline a performance model that can be used to estimate the time required by reconfigurations, and we describe an implementation of the model. The evaluation demonstrates the accuracy of the performance estimations, and illustrates the performance gains provided by the execution model of Concerto compared to state-of-the-art systems.

An Optimal Model for Optimizing the Placement and Parallelism of Data Stream Processing Applications on Cloud-Edge Computing

Participants Eddy Caron, Felipe Rodrigo De Souza, Marcos Dias De Assunção, Laurent Lefèvre.

The Internet of Things has enabled many application scenarios where a large number of connected devices generate unbounded streams of data, often processed by data stream processing frameworks deployed in the cloud. Edge computing enables offloading processing from the cloud and placing it close to where the data is generated, thereby reducing the time to process data events and deployment costs. However, edge resources are more computationally constrained than their cloud counterparts, raising two interrelated issues, namely deciding on the parallelism of processing tasks (a.k.a. operators) and their mapping onto available resources. In [13, 19] we provided some mechanism to improve this kind of mapping. In [14], we formulate the scenario of operator placement and parallelism as an optimal mixed-integer linear programming problem. The proposed model is termed as Cloud-Edge data Stream Placement (CESP). Experimental results using discrete-event simulation demonstrate that CESP can achieve an end-to-end latency at least 80% and monetary costs at least 30% better than traditional cloud deployment.

An Operational Tool for Software Asset Management Improvement

Participants Ghoshana Bista, Eddy Caron, Arthur Chevalier.

This research takes place in the field of Software Asset Management (SAM), license management, use rights, and compliance with contractual rules. When talking about proprietary software, these rules are often misinterpreted or totally misunderstood. In exchange for the fact that we are free to license our use as we see fit, in compliance with the contract, the publishers have the right to make audits. They can check that the rules are being followed and, if they are not respected, they can impose penalties, often financial penalties. The emergence of the Cloud has greatly increased the problem because software usage rights were not originally intended for this type of architecture. We have studied the licensing methods of major publishers such as Oracle, IBM and SAP before introducing the various problems inherent in SAM. The lack of standardization in metrics, specific usage rights, and the difference in paradigm brought about by the Cloud and soon the virtualized network make the situation more complicated than it already was. In [18] our research is oriented towards modeling these licenses and metrics in order to abstract from the legal and blurry side of contracts. This abstraction allows us to develop software placement algorithms that ensure that contractual rules are respected at all times. This licensing model also allows us to introduce a deployment heuristic that optimizes several criteria at the time of software placement such as performance, energy and cost of licenses. We then introduce the problems associated with deploying multiple software at the same time by optimizing these same criteria and prove the NP-completeness of the associated decision problem. In order to meet these criteria, we present a placement algorithm that approaches the optimal and uses the above heuristic [16] to provide a Software Asset Management Compliance in Green Deployment Algorithm. In parallel, we have developed a SAM tool that uses these researches to offer an automated and totally generic software management in a Cloud architecture. All this work has been conducted in collaboration with Orange and tested in different Proof-Of-Concept before being fully integrated into the SAM tool. More recently we extended this previous work to the Virtual Network Function (VNF) [17].

Modeling and Scheduling of Scientific Workflows with Stochastic Task Weights

Participants Yves Caniou, Eddy Caron, Aurélie Kong Win Chang.

Many applications in the form of scientific workflows are run with workflow management systems like Pegasus and Apache Airflow. Such scientific workflows are composed of tasks whose duration can only be roughly estimated (it may depend on input data length for example). Some work took place for the modelization of such workflows, and we then designed and studied the performance of budget aware static scheduling heuristics on real life workflows like LIGO, CYBERSHAKE, MONTAGE [1]. Experiments were conducted in real life using DIET and the Grid'5000 resources, with the development of a tool to automatize the generation of experiments and the DIET code [21].

6.4 HPC Applications and Runtimes

sOMP: Simulating OpenMP Task-Based Applications with NUMA Effects

Participants Thierry Gautier.

Anticipating the behavior of applications, studying, and designing algorithms are some of the most important purposes for the performance and correction studies about simulations and applications relating to intensive computing. Often studies that evaluate performance on a single-node of a simulation don't consider Non-Uniform Memory Access (NUMA) as having a critical effect. The work [7] focuses on

accurately predicting the performance of task-based OpenMP applications from traces collected through the OMPT interface. We first introduce TiKKi, a tool that records a rich high-level representation of the execution trace of a real OpenMP application. With this trace, an accurate prediction of the execution time is modeled from the architecture of the machine and sOMP, a SimGrid-based simulator for task-based applications with data dependencies. These predictions are improved when the model takes into account memory transfers. We show that good precision (10% relative error on average) can be obtained for various grains and on different numbers of cores inside different shared-memory architectures.

P-Aevol: an OpenMP Parallelization of a Biological Evolution Simulator, Through Decomposition in Multiple Loops

Participants Thierry Gautier, Christian Perez, Laurent Turpin.

In [15], we describe how we have achieved the parallelization of Aevol, a biological evolution simulator, on multi-core architecture using the OpenMP standard. While it looks like a simple for-loop problem with independent iterations, the stochastic nature of Aevol makes the duration of the iterations unpredictable and it conveys a high irregularity. Classical scheduling algorithms of OpenMP runtimes turn out to be inefficient. By analysing the origin of this irregularity, this paper presents how to transform the highly irregular Aevol for-loop to a sequence composed by a small duration irregular for-loop followed by work intensive for-loop easy to schedule using classical LPT algorithm. This method leads to a gain up to 27% from the best OpenMP loop schedule.

XKBlas: a High Performance Implementation of BLAS-3 Kernels on Multi-GPU Server

Participants Thierry Gautier, Marie Durand.

In the last ten years, GPUs have dominated the market considering the computing/power metric and numerous research works have provided Basic Linear Algebra Subprograms implementations accelerated on GPUs. Several software libraries have been developed for exploiting performance of systems with accelerators, but the real performance may be far from the platform peak performance. In the paper [8] we present XKBlas that aims to improve performance of BLAS-3 kernels on multi-GPU systems. At low level, we model computation as a set of tasks accessing data on different resources. At high level, the API design favors non-blocking calls as uniform concept to overlap latency, even by fine grain computation. Unit benchmark of BLAS-3 kernels showed that XKBlas outperformed most implementations including the overhead of dynamic task's creation and scheduling. XKBlas outperformed BLAS implementations such as cuBLAS-XT, PaRSEC, BLASX and Chameleon/StarPU.

With the work of M. Durand with the contract with MUMPS Tech., XKBlas was also a runtime for MUMPS software.

6.5 Improving Virtualized Systems

Fine-Grained Fault Tolerance For Resilient pVM-based Virtual Machine Monitors

Participants Alain Tchana, Djob Mvondo.

Virtual machine monitors (VMMs) play a crucial role in the software stack of cloud computing platforms: their design and implementation have a major impact on performance, security and fault tolerance. In this paper, we focus on the latter aspect (fault tolerance), which has received less attention, although it is now a significant concern. Our work aims at improving the resilience of the "pVM-based" VMMs, a popular design pattern for virtualization platforms. In such a design, the VMM is split into two main

components: a bare-metal hypervisor and a privileged guest virtual machine (pVM). We highlight that the pVM is the least robust component and that the existing fault-tolerance approaches provide limited resilience guarantees or prohibitive overheads. We present in [11] three design principles (disaggregation, specialization, and pro-activity), as well as optimized implementation techniques for building a resilient pVM without sacrificing end-user application performance. We validate our contribution on the mainstream Xen platform.

Fine-Grained Isolation for Scalable, Dynamic, Multi-tenant Edge Clouds

Participants Alain Tchana.

5G edge clouds promise a pervasive computational infrastructure a short network hop away, enabling a new breed of smart devices that respond in real-time to their physical surroundings. Unfortunately, today's operating system designs fail to meet the goals of scalable isolation, dense multi-tenancy, and high performance needed for such applications. In this paper we introduce EdgeOS that emphasizes system-wide isolation as fine-grained as per-client. We propose a novel memory movement accelerator architecture that employs data copying to enforce strong isolation without performance penalties. To support scalable isolation, we introduce a new protection domain implementation that offers lightweight isolation, fast startup and low latency even under high churn. We implement EdgeOS [12] in a microkernel based OS and demonstrate running high scale network middleboxes using the Click software router and endpoint applications such as memcached, a TLS proxy, and neural network inference. We reduce startup latency by 170X compared to Linux processes, and improve latency by three orders of magnitude when running 300 to 1000 edge-cloud memcached instances on one server.

7 Bilateral contracts and grants with industry

7.1 Bilateral contracts with industry

MUMPS Technologies AVALON has a collaboration with MUMPS Technologies. The funding is dedicated for Marie Durand during few months to make experimental validation of the interest of using XKBLAS library to let MUMPS software to gain in performance on multi-GPUs server.

7.2 Bilateral grants with industry

Orange We have a collaboration with Orange. This collaboration is sealed through a CIFRE Phd grant. The research of the Phd student (Arthur Chevalier) focuses on placement and compliance aspects of software licenses in a Cloud architecture. Today, the use of software is regulated by licenses, whether they are free, paid for and with or without access to its sources. The number of licenses required for specific software can be calculated with several metrics, each defined by the software vendor. Our goal is to propose a deployment algorithm that takes into account different metrics.

In 2020 we have started a new thesis to extend the previous work. With the Phd student Ghoshana Bista we focus on the software asset management dedicated to the VNF (Virtual Network Function).

Stackeo We have a collaboration with the start'up Stackeo who enhanced our work around modeling of Fog infrastructure. Eddy Caron was co-supervisor of a transfer engineer technology Inria, Zakaria Fraoui from January 2020 to December 2020.

Thales We have a collaboration with Thalès. This collaboration is sealed thanks to a CIFRE Phd grant. The research of the Phd student (Pierre-Etienne Polet) focuses executing signal processing application on GPU for embedded architecture. The problem and its solutions are at the confluence of task scheduling with memory limitation, optimization, parallel algorithm and runtime system.

CEA We have a collaboration with CEA / DAM-Île de France. This collaboration is based on the co-advising of a CEA Phd. The research of the Phd student (Romain Pereira) Polet focuses high performance OpenMP + MPI executions. MPC was developed for high performance MPI application. Recently a support for OpenMP was added. The goal of the PhD is to work on better cooperation of OpenMP and MPI thanks to the unique framework MPC.

8 Partnerships and cooperations

8.1 International initiatives

8.1.1 Inria International Labs

Title Joint Laboratory for Extreme Scale Computing ([JLESC](#))

Duration 2014-2023

Partners NCSA (US), ANL (US), Inria (FR), Jülich Supercomputing Centre (DE), BSC (SP), Riken (JP).

Summary The purpose of the Joint Laboratory for Extreme Scale Computing (JLESC) is to be an international, virtual organization whose goal is to enhance the ability of member organizations and investigators to make the bridge between Petascale and Extreme computing. The founding partners of the JLESC are Inria and UIUC. Further members are ANL, BSC, JSC and RIKEN-AICS.

JLESC involves computer scientists, engineers and scientists from other disciplines as well as from industry, to ensure that the research facilitated by the Laboratory addresses science and engineering's most critical needs and takes advantage of the continuing evolution of computing technologies.

8.1.2 Inria international partners

Declared Inria international partners

Title Square Kilometer Array ([SKA](#))

Participants Vincent Lanore, Christian Perez

Summary Through the membership of Inria to the [Maison SKA-France \(MSF\)](#), the Avalon team collaborates with [SKA Organization](#) that has been responsible for coordinating the global activities towards the SKA in the pre-construction phase.

Title Rutgers University, New-Jersey (USA)

Participants Eddy Caron, Zeina Houmani, Laurent Lefèvre

Summary From 2017 to 2019 we had a long-term collaboration between the RDI² team (Rutgers University) and our team through the associate team SUSTAM. Beyond and thanks of this collaboration we started a thesis to build a Data-driven microservices architecture for Deep Learning applications with Zeina Houmani. The funding of this Phd was 50/50 between the US and French partners.

8.2 European initiatives

8.2.1 FP7 & H2020 Projects

Title Energy oriented Centre of Excellence for computing applications – [EoCoE-II](#)

Duration: Jan 2019 - Dec 2021

Coordinator: CEA (France)

Partners: CEA, FZJ, ENEA, BSC, CNRS, INRIA, CERFACS, MPG, FRAUNHOFER, FAU, CNR, UNITN, PSNC, ULB, UBAH, CIEMAT, IFPEN, DDN, RWTH, UNITOV

Inria contact: Thierry Gautier

Summary: Europe is undergoing a major transition in its energy generation and supply infrastructure. The urgent need to halt carbon dioxide emissions and prevent dangerous global temperature rises has received renewed impetus following the unprecedented international commitment to enforcing the 2016 Paris Agreement on climate change. Rapid adoption of solar and wind power generation by several EU countries has demonstrated that renewable energy can competitively supply significant fractions of local energy needs in favourable conditions. These and other factors have combined to create a set of irresistible environmental, economic and health incentives to phase out power generation by fossil fuels in favour of decarbonized, distributed energy sources. While the potential of renewables can no longer be questioned, ensuring reliability in the absence of constant conventionally powered baseload capacity is still a major challenge.

The EoCoE-II project will build on its unique, established role at the crossroads of HPC and renewable energy to accelerate the adoption of production, storage and distribution of clean electricity. How will we achieve this? In its proof-of-principle phase, the EoCoE consortium developed a comprehensive, structured support pathway for enhancing the HPC capability of energy-oriented numerical models, from simple entry-level parallelism to fully-fledged exascale readiness. At the top end of this scale, promising applications from each energy domain have been selected to form the basis of 5 new Energy Science Challenges in the present successor project EoCoE-II that will be supported by 4 Technical Challenges

Title: PRACE 6th Implementation Phase Project (PRACE 6IP)

Duration: May 2019 - Dec 2021

Coordinator: FORSCHUNGSZENTRUM JULICH GMBH (Germany)

Partners:

- KADEMIA GORNICZO-HUTNICZA IM. STANISŁAWA STASZICA W KRAKOWIE (Poland)
- EATEASSOCIACAO DO INSTITUTO SUPERIOR TECNICO PARA A INVESTIGACAO E DESENVOLVIMENTO (Portugal)
- TEASSOCIATION "NATIONAL CENTRE FOR SUPERCOMPUTING APPLICATIONS (Bulgaria)
- BARCELONA SUPERCOMPUTING CENTER - CENTRO NACIONAL DE SUPERCOMPUTACION (Spain)
- BAYERISCHE AKADEMIE DER WISSENSCHAFTEN (Germany)
- BILKENT UNIVERSITESI VAKIF (Turkey)
- CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE CNRS (France)
- AMCENTRUM SPOLOCNYCH CINNOSTI SLOVENSKEJ AKADEMIE VIED (Slovakia)
- CINECA CONSORZIO INTERUNIVERSITARIO (Italy)
- COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES (France)
- DANMARKS TEKNISKE UNIVERSITET (Denmark)
- EIDGENOESSISCHE TECHNISCHE HOCHSCHULE ZUERICH (Switzerland)
- FORSCHUNGSZENTRUM JULICH GMBH (Germany)
- FUNDACION PUBLICA GALLEGA CENTRO TECNOLOGICO DE SUPERCOMPUTACION DE GALICIA (Spain)
- GEANT VERENIGING (Netherlands)
- GRAND EQUIPEMENT NATIONAL DE CALCUL INTENSIF (France)
- Gauss Centre for Supercomputing (GCS) e.V. (Germany)
- ISTANBUL TEKNIK UNIVERSITESI (Turkey)

- KOBENHAVNS UNIVERSITET (Denmark)
- KORMANYZATI INFORMATIKAI FEJLESZTESI UGYNOKSEG (Hungary)
- KUNGLIGA TEKNISKA HOEGSKOLAN (Sweden)
- LINKOPINGS UNIVERSITET (Sweden)
- MACHBA - INTERUNIVERSITY COMPUTATION CENTER (Israel)
- MAX-PLANCK-GESELLSCHAFT ZUR FORDERUNG DER WISSENSCHAFTEN EV (Germany)
- NATIONAL INFRASTRUCTURES FOR RESEARCH AND TECHNOLOGY (Greece)
- NATIONAL UNIVERSITY OF IRELAND GALWAY (Ireland)
- NORGES TEKNISK-NATURVITENSKAPELIGE UNIVERSITET NTNU (Norway)
- PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE AISBL (Belgium)
- POLITECHNIKA GDANSKA (Poland)
- POLITECHNIKA WROCLAWSKA (Poland)
- SYDDANSK UNIVERSITET (Denmark)
- TECHNISCHE UNIVERSITAET WIEN (Austria)
- THE CYPRUS INSTITUTE (Cyprus)
- THE UNIVERSITY OF EDINBURGH (UK)
- UNINETT SIGMA2 AS (Norway)
- UNITED KINGDOM RESEARCH AND INNOVATION (UK)
- UNIVERSIDADE DE COIMBRA (Portugal)
- UNIVERSIDADE DE EVORA (Portugal)
- UNIVERSIDADE DO MINHO (Portugal)
- UNIVERSIDADE DO PORTO (Portugal)
- UNIVERSITAET INNSBRUCK (Austria)
- UNIVERSITAET STUTTGART (Germany)
- UNIVERSITE DU LUXEMBOURG (Luxembourg)
- UNIVERSITEIT ANTWERPEN (Belgium)
- UNIVERSITETET I OSLO (Norway)
- UNIVERZA V LJUBLJANI (Slovenia)
- UPPSALA UNIVERSITET (Sweden)
- VSB - Technical University of Ostrava (Czech Republic)

Inria contact: Christian Perez

Summary: PRACE, the Partnership for Advanced Computing is the permanent pan-European High Performance Computing service providing world-class systems for world-class science. Systems at the highest performance level (Tier-0) are deployed by Germany, France, Italy, Spain and Switzerland, providing researchers with more than 17 billion core hours of compute time. HPC experts from 25 member states enabled users from academia and industry to ascertain leadership and remain competitive in the Global Race. Currently PRACE is finalizing the transition to PRACE 2, the successor of the initial five year period. The objectives of PRACE-6IP are to build on and seamlessly continue the successes of PRACE and start new innovative and collaborative activities proposed by the consortium. These include: assisting the development of PRACE 2; strengthening the internationally recognised PRACE brand; continuing and extend advanced training which so far provided more than 36 400 person-training days; preparing strategies and best practices towards Exascale computing, work on forward-looking SW solutions; coordinating and enhancing the operation of the multi-tier HPC systems and services; and supporting users to exploit massively parallel systems and novel architectures. A high level Service Catalogue is provided. The proven

project structure will be used to achieve each of the objectives in 7 dedicated work packages. The activities are designed to increase Europe's research and innovation potential especially through: seamless and efficient Tier-0 services and a pan-European HPC ecosystem including national capabilities; promoting take-up by industry and new communities and special offers to SMEs; assistance to PRACE 2 development; proposing strategies for deployment of leadership systems; collaborating with the ETP4HPC, CoEs and other European and international organisations on future architectures, training, application support and policies. This will be monitored through a set of KPIs.

Title: Scientific Large-scale Infrastructure for Computing/Communication Experimental Studies (SLICES-DS)

Duration: Sep 2020 – Aug 2022

Coordinator: Sorbonne Université (France)

Partners:

- CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE CNRS (France)
- CONSIGLIO NAZIONALE DELLE RICERCHE (Italy)
- ECOLE NORMALE SUPERIEURE DE LYON (France)
- INTERUNIVERSITAIR MICRO-ELECTRONICA CENTRUM (Belgium)
- MANDAT INTERNATIONAL ALIAS FONDATION POUR LA COOPERATION INTERNATIONALE (Switzerland)
- PANEPISTIMIO THESSALIAS ((Greece))
- UCLAN CYPRUS LIMITED (Cyprus)
- UNIVERSIDAD CARLOS III DE MADRID (Spain)
- UNIVERSITEIT VAN AMSTERDAM (Netherlands)

Inria contact: Christian Perez

Summary: Digital Infrastructures as the future Internet, constitutes the cornerstone of the digital transformation of our society. As such, Innovation in this domain represents an industrial need, a sovereignty concern and a security threat. Without Digital Infrastructure, none of the advanced services envisaged for our society is feasible. They are both highly sophisticated and diverse physical systems but at the same time, they form even more complex, evolving and massive virtual systems. Their design, deployment and operation are critical. In order to research and master Digital infrastructures, the research community needs to address significant challenges regarding their efficiency, trust, availability, reliability, range, end-to-end latency, security and privacy. Although some important work has been done on these topics, the stringent need for a scientific instrument, a test platform to support the research in this domain is an urgent concern. SLICES ambitions to provide a European-wide test-platform, providing advanced compute, storage and network components, interconnected by dedicated high-speed links. This will be the main experimental collaborative instrument for researchers at the European level, to explore and push further, the envelope of the future Internet. A strong, although fragmented expertise, exists in Europe and could be leveraged to build it. SLICES is our answer to this need. It is ambitious, practical but overall timely and necessary. The main objective of SLICES-DS is to adequately design SLICES in order to strengthen research excellence and innovation capacity of European researchers and scientists in the design and operation of Digital Infrastructures. The SLICES Design study will build upon the experience of the existing core group of partners, to prepare in details the conceptual and technical design of the new leading edge SLICES-RI for the next phases of the RI's lifecycle.

8.3 National initiatives

8.3.1 Inria Large Scale Initiative

HAC SPECIS, *High-performance Application and Computers, Studying PErformance and Correctness In Simulation*, 4 years, 2016-2020

Participants Dorra Boughzala, Idriss Daoudi, Thierry Gautier, Laurent Lefèvre, Frédéric Suter.

Over the last decades, both hardware and software of modern computers have become increasingly complex. Multi-core architectures comprising several accelerators (GPUs or the Intel Xeon Phi) and interconnected by high-speed networks have become mainstream in HPC. Obtaining the maximum performance of such heterogeneous machines requires to break the traditional uniform programming paradigm. To scale, application developers have to make their code as adaptive as possible and to release synchronizations as much as possible. They also have to resort to sophisticated and dynamic data management, load balancing, and scheduling strategies. This evolution has several consequences:

First, this increasing complexity and the release of synchronizations are even more error-prone than before. The resulting bugs may almost never occur at small scale but systematically occur at large scale and in a non deterministic way, which makes them particularly difficult to identify and eliminate.

Second, the dozen of software stacks and their interactions have become so complex that predicting the performance (in terms of time, resource usage, and energy) of the system as a whole is extremely difficult. Understanding and configuring such systems therefore becomes a key challenge.

These two challenges related to correctness and performance can be answered by gathering the skills from experts of formal verification, performance evaluation and high performance computing. The goal of the HAC SPECIS Inria Project Laboratory is to answer the methodological needs raised by the recent evolution of HPC architectures by allowing application and runtime developers to study such systems both from the correctness and performance point of view.

8.4 Regional initiatives

8.4.1 CPER

LECO

Participants Thierry Gautier, , Laurent Lefèvre, , Christian Perez.

In the continuation of the LECO platform funding in 2019, the GreenCube project funded by the DRRT 2019-2021 aims at installing a research platform to studying application with small computer with limited energy budget. Due to the COVID-19 crisis, the project was re-oriented to install a full simulation environment.

8.4.2 Action Exploratoire Inria

EXODE

Participants Thierry Gautier, .

In biology, the vast majority of systems can be modeled as ordinary differential equations (ODEs). Modeling more finely biological objects leads to increase the number of equations. Simulating ever larger systems also leads to increasing the number of equations. Therefore, we observe a large increase in the size of the ODE systems to be solved. A major lock is the limitation of ODE numerical resolution software (ODE solver) to a few thousand equations due to prohibitive calculation time. The AEx ExODE tackles this

lock via 1) the introduction of new numerical methods that will take advantage of the mixed precision that mixes several floating number precisions within numerical methods, 2) the adaptation of these new methods for next generation highly hierarchical and heterogeneous computers composed of a large number of CPUs and GPUs. For the past year, a new approach to Deep Learning has been proposed to replace the Recurrent Neural Network (RNN) with ODE systems. The numerical and parallel methods of ExODE will be evaluated and adapted in this framework in order to improve the performance and accuracy of these new approaches.

9 Dissemination

9.1 Promoting scientific activities

9.1.1 Scientific events: organisation

General chair, scientific chair

- Frédéric was the co-organizer of the *The Many Faces of Simulation for HPC Minisymposium* at the SIAM Conference on Parallel Processing for Scientific Computing
- Laurent Lefevre was the co-organizer of the online colloquium : "Effets rebonds dans le numérique. Comment les détecter ? Comment les mesurer ? Comment les éviter ?", with Centre Jacques Cartier, Université de Sherbrooke, Inria, GDS CNRS EcoInfo, November 2, 2020

Member of the organizing committees

- Christian Perez was member of the Organizing Committee of the French Journées Calcul Données (Dijon, 2-4 Dec 2020).

9.1.2 Scientific events: selection

Chair of conference program committees

- Christian Perez was vice-chair of the program committee for the track *Industry and Experimentation* in 40th IEEE International Conference on Distributed Computing Systems 2020.
- Laurent Lefevre was :
 - Co Program Chair of SBAC-PAD2020 Conference : IEEE 32nd International Symposium on Computer Architecture and High Performance Computing , Porto, Portugal, September 8-11, 2020
 - Co Program Chair of CCGrid 2020 conference: The 20th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, Melbourne, Australia, May 11-14, 2020

Member of the conference program committees

- Yves Caniou was a program committee member for the *International Conference on Computational Science and its Applications* 2020.
- Eddy Caron is an organizing committee member for the French Compas 2020 conference.
- Christian Perez was a program committee member for the *Programming Models and Runtime Systems* Track in IEEE/ACM CCGrid 2020, for the *Programming and System Software* area in IEEE Cluster 2020, for the *Parallel and Distributed Programming, Interfaces and Languages* topic in Euro-Par2020 conference. He was a project poster committee member in ISC High Performance 2020 and a program committee member in the French Compas 2020 conference.
- Frédéric Suter was a program committee member for the *International Conference on Parallel Processing*, for the *IEEE 32nd International Symposium on Computer Architecture and High Performance Computing*, and for the *15th Workflows in Support of Large-Scale Science Workshop*.

9.1.3 Journal

- Eddy Caron was reviewer for Concurrency and Computation: Practice and experience, International Transactions in Operational Research and Cluster Computing.
- Laurent Lefevre is associated editor of the IEEE Transactions on Sustainable Computing Journal.

9.1.4 Scientific expertise

- Christian Perez evaluated 8 projects for the French Direction générale de la Recherche et de l'Innovation.
- Frédéric Suter evaluated a funding proposal from the Belgian Fonds de la Recherche Scientifique.

9.1.5 Research administration

- Eddy Caron is Deputy Director in charge of call for projects, research transfert and international affairs from September 2017 to December 2020 for the LIP. He is co-leader of the Distributed system and HPC team of the FIL (Fédération Informatique de Lyon).
- Olivier Glück is member of the "Conseil Académique" of Lyon 1 University and Lyon University.
- Laurent Lefevre is a member of the executive board and the sites committee of the Grid'5000 Scientific Interest Group. He is the scientific leader of the Grid'5000 Lyon site. He is animator and co-chair of the transversal action on "Energy" of the French GDR RSD ("Réseaux et Systèmes Distribués"). He is member of the scientific advisory board of the Digital League cluster (Région Rhone Alpes). He is elected member in the LIP laboratory council (ENS Lyon). He is co-director of the CNRS GDS EcoInfo group. He is the responsible of M2 training period (ENS Lyon) and leader of the PhD commission of the LIP laboratory.
- Christian Perez represents INRIA in the overview board of the France Grilles Scientific Interest Group. He is a member of the executive board and the sites committee of the Grid'5000 Scientific Interest Group and member of the executive board of the Silecs testbed. He is a member of the Inria Grenoble Rhône-Alpes Strategic Orientation Committee. He is in charge of organizing scientific collaborations between INRIA and SKA France.

9.2 Teaching - Supervision - Juries

9.2.1 Teaching

- Licence: Eddy Caron, Programmation, 48h, L3, ENS de Lyon. France.
- Master: Eddy Caron, Integrated Project, 42h, M1, ENS de Lyon. France.
- Master: Eddy Caron, Distributed System, 30h, M1, ENS de Lyon. France.
- Licence: Yves Caniou, Algorithmique programmation impérative initiation, 49h, niveau L1, Université Claude Bernard Lyon 1, France.
- Licence: Yves Caniou, Programmation Concurrente, 49h and Responsible of UE, niveau L3, Université Claude Bernard Lyon 1, France.
- Licence: Yves Caniou, supervision 2 month internship, niveau L3, Université Claude Bernard Lyon 1, France.
- Licence: Yves Caniou, Réseaux, 36h, niveau L3, Université Claude Bernard Lyon 1, France.
- Licence: Yves Caniou, Systèmes d'information documentaire, 20h, niveau L3, Université Claude Bernard Lyon 1, France.

- Licence, Yves Caniou, Responsable mission pédagogique particulière, 4h, L3, Université Claude Bernard Lyon 1, France.
- Master: Yves Caniou, Sécurité, 25.5h and Responsible of UE, niveau M2, Université Claude Bernard Lyon 1, France.
- Master: Yves Caniou, Systèmes Avancés, 4.5h, niveau M2, Université Claude Bernard Lyon 1, France.
- Master: Yves Caniou, Responsible of alternance students, 6h, niveau M1, Université Claude Bernard Lyon 1, France.
- Master: Yves Caniou, Responsible of alternance students, 12h, niveau M2, Université Claude Bernard Lyon 1, France.
- Master: Laurent Lefèvre, Parallélisme, 12h, niveau M1, Université Lyon 1, France.
- CAPES Informatique : Laurent Lefèvre, Numérique responsable, 3h, Université Lyon1, France
- Licence: Olivier Glück, Licence pedagogical advisor, 30h, niveaux L1, L2, L3, Université Lyon 1, France.
- Licence: Olivier Glück, Introduction Réseaux et Web, 54h, niveau L1, Université Lyon 1, France.
- Licence: Olivier Glück, Bases de l'architecture pour la programmation, 23h, niveau L1, Université Lyon 1, France.
- Licence: Olivier Glück, Algorithmique programmation impérative initiation, 56h, niveau L1, Université Lyon 1, France.
- Licence: Olivier Glück, Réseaux, 2x70h, niveau L3, Université Lyon 1, France.
- Master: Olivier Glück, Réseaux par la pratique, 10h, niveau M1, Université Lyon 1, France.
- Master: Olivier Glück, Responsible of Master SRIV (Systèmes, Réseaux et Infrastructures Virtuelles) located at IGA Casablanca, 20h, niveau M2, IGA Casablanca, Maroc.
- Master: Olivier Glück, Applications systèmes et réseaux, 30h, niveau M2, Université Lyon 1, France.
- Master: Olivier Glück, Applications systèmes et réseaux, 24h, niveau M2, IGA Casablanca, Maroc.
- Master: Olivier Glück, Administration des Systèmes et des Réseaux, 16h, niveau M2, Université Lyon 1, France.
- Master: Olivier Glück, DIU Enseigner l'Informatique au Lycée, 50h, Formation continue, Université Lyon 1, France.
- Licence : Frédéric Suter, Programmation Concurrente, 32.33, L3, Université Claude Bernard Lyon 1, France

9.2.2 Supervision

- PhD: Vo Quoc Bao Bui, *Extended Para-Virtualization*, Sep 29th, 2020, Alain Tchana (dir), Daniel Hagimont (INPT, co-dir).
- PhD: Arthur Chevalier, *Optimisation du placement des licences logiciel des fonctions réseau dans le Cloud pour un déploiement économique et efficace*, Nov. 24th, 2020, Eddy Caron (dir), Noëlle Baillon (co-dir, Orange).
- PhD: Maverick Chardet, *Reconfiguration of large scale distributed systems and software in a fog computing context*, Dec 3rd, 2020, Christian Perez (dir), Hélène Coullon (Stack, Nantes, co-dir), Adrien Lèbre (Stack, Nantes, co-dir).

- PhD: Felipe Rodrigo De Souza, *Networking Provisioning Algorithms for Highly Distributed Data Stream Processing*, Dec 10th, 2020, Eddy Caron (dir), Marcos Dias de Assunção (co-dir).
- PhD: Barbe Thystere Mvondo Djob, *Improvement of the privileged domain in virtualized systems*, Dec 18th, 2020, Alain Tchana (dir), Noel De Palma (UGA, co-dir)
- PhD in progress: Celestine Stella Ndonga Bitchebe, *Hardware features for virtualization*, March 1st, 2019, Alain Tchana (dir).
- PhD in progress: Dorra Boughzala, *Simulating Energy Consumption of Continuum Computing between Heterogeneous Numerical Infrastructures in HPC*, IPL Hac-Specis Inria, Laurent Lefèvre (dir), Martin Quinson and Anne-Cécile Orgerie (Myriads, Rennes, co-dir) (December 2017-November 2020).
- PhD in progress: Idriss Daoudi, *Simulating OpenMP program*, October 2018, Samuel Thibault (Univ-Bordeaux, Storm team, Bordeaux, dir) and Thierry Gautier (INRIA, AVALON team, co-dir).
- PhD in progress: Zeina Houmani, *A Data-driven microservices architecture for Deep Learning applications*, Eddy Caron (dir), Daniel Balouek-Thomert (Rutgers University) (since oct. 2018).
- PhD in progress: Patrick Lavoisier Wapet, *Illegitimate app detection in mobile phones*, 1 oct. 2017, Alain Tchana (dir), Daniel Hagimont (INPT, co-dir).
- PhD in progress: Laurent Turpin, *Mastering Code Variation and Architecture Evolution for HPC application*, October 2019, Christian Perez (INRIA, AVALON team, dir), Jonathan Rouzaud-Cornabas (INSA, Beagle team, co-dir) and Thierry Gautier (INRIA, AVALON team, co-dir).
- PhD In progress: Ghoshana Bista, *VNF and Software Asset Management*, Feb 2020, Eddy Caron (dir), Anne-Lucie Vion (Orange).
- PhD In progress: Pierre-Etienne Polet, *GPU-ification of signal processing application for embedded HPC*, July 2020, Thierry Gautier (dir), Ramy Fantar (Thalès DMS, co-dir).
- PhD In progress: Hugo Hadjur, *Designing sustainable autonomous connected systems with low energy consumption*, Sept. 2020, Laurent Lefevre (dir), Doreid Ammar (Aivancity group, co-dir)
- PhD In progress: Romain Pereira, *High Performance OpenMP + MPI on top of MPC*, Nov 2020, Thierry Gautier (dir), Patrick Carribault (CEA, co-dir), Adrien Roussel (CEA, co-dir)
- PhD In progress: Kevin Nguetchouang Ngongan, *Out-of-Hypervisor*, Dec. 2020, Alain Tchana (dir)
- PhD In progress: Lucien Ndjie Ngale, *Proposition et mise en œuvre d'une architecture pour la robotique supportant des ordonnancements efficaces et asynchrone dans un contexte d'architectures virtualisées*, Nov 2020, Eddy Caron (dir), Yulin Zhang (Université de Picardie Jules Verne, co-dir).

9.2.3 Juries

- Eddy Caron was reviewer of the HDR of Flavien Vernier (université Savoie Mont Blanc), defense is postponed to 2021.
- Christian Perez was reviewer and member of the PhD defense committee of Grégoire Todeschi, Université de Toulouse, France, June 8th, 2020.
- Frédéric Suter was reviewer and member of the PhD defense committees of Adrien Faure, Université de Grenoble Alpes, France, Alexandre Honorat, Insa de Rennes, France, and Valentin Honoré, Université de Bordeaux, France.

9.3 Popularization

- Laurent Lefevre has done interviews for the :
 - radio show "Nouveau monde", France Info Radio : "L'innovation technologique ne suffit pas pour limiter l'impact environnemental du numérique", October 10, 2020
 - radio show "Sauvons la planète!", France Bleu Périgord Radio : "La pollution numérique, des enjeux majeurs pour la planète", October 9, 2020
 - radio show "La terre au carré", France Inter Radio : "La pollution numérique a-t-elle augmenté durant le confinement ?", June 24, 2020

9.3.1 Internal or external Inria responsibilities

- Eddy Caron is member of the CDT Inria.

9.3.2 Articles and contents

- Article on "Le vrai coût énergétique du numérique", Anne-Cécile Orgerie et Laurent Lefèvre, Pour la Science, November 25, 2020

9.3.3 Interventions

- Laurent Lefevre has performed the following interventions :
 - "Pollution numérique, tous responsables ?", Laurent Lefevre, Fête de la Science, Inria Channel, October 6, 2020
 - "Pollution numérique, tous responsables ?", Laurent Lefevre, Fête de la Science, MJC Rive de Gier, October 5, 2020
 - "Numérique et environnement... un curseur personnel, des impacts, des dérapages et des ripostes venant de la recherche", Laurent Lefevre, Talk in front of Occitanie Region, Green New Deal Occitanie, June 23, 2020
 - "Numérique et environnement... un curseur personnel, des impacts, des dérapages, des ripostes, des services...", Laurent Lefevre, Congrès Société Informatique de France (SIF) Transitions numériques et écologiques, Lyon, February, 4, 2020
 - Panelist in the "La tragédie électronique", Festival des bonnes résolutions, MJC Rive de Gier, France, January 11, 2020

10 Scientific production

10.1 Publications of the year

International journals

- [1] Y. Caniou, E. Caron, A. Kong Win Chang and Y. Robert. 'Budget-aware scheduling algorithms for scientific workflows with stochastic task weights on IaaS Cloud platforms *'. In: *Concurrency and Computation: Practice and Experience* (2020). URL: <https://hal.inria.fr/hal-03029318>.
- [2] H. Casanova, R. Ferreira Da Silva, R. Tanaka, S. Pandey, G. Jethwani, W. Koch, S. Albrecht, J. Oeth and F. Suter. 'Developing Accurate and Scalable Simulators of Production Workflow Management Systems with WRENCH'. In: *Future Generation Computer Systems* 112 (Nov. 2020), pp. 162–175. DOI: [10.1016/j.future.2020.05.030](https://doi.org/10.1016/j.future.2020.05.030). URL: <https://hal.inria.fr/hal-02878322>.
- [3] M. Chardet, H. Coullon and S. Robillard. 'Toward Safe and Efficient Reconfiguration with Concerto'. In: *Science of Computer Programming* 203 (1st Mar. 2021). DOI: [10.1016/j.scico.2020.102582](https://doi.org/10.1016/j.scico.2020.102582). URL: <https://hal.inria.fr/hal-03103714>.

- [4] R. Ferreira Da Silva, H. Casanova, A.-C. Orgerie, R. Tanaka, E. Deelman and F. Suter. ‘Characterizing, Modeling, and Accurately Simulating Power and Energy Consumption of I/O-intensive Scientific Workflows’. In: *Journal of computational science* 44 (June 2020), p. 101157. DOI: [10.1016/j.jocs.2020.101157](https://doi.org/10.1016/j.jocs.2020.101157). URL: <https://hal.archives-ouvertes.fr/hal-02876736>.

International peer-reviewed conferences

- [5] D. Boughzala, L. Lefèvre and A.-C. Orgerie. ‘Predicting the energy consumption of CUDA kernels using SimGrid’. In: SBAC-PAD 2020 - 32nd IEEE International Symposium on Computer Architecture and High Performance Computing. SBAC-PAD 2020 - 32nd IEEE International Symposium on Computer Architecture and High Performance Computing. Porto, Portugal, Sept. 2020, pp. 191–198. URL: <https://hal.archives-ouvertes.fr/hal-02924028>.
- [6] M. Chardet, H. Coullon and C. Pérez. ‘Predictable Efficiency for Reconfiguration of Service-Oriented Systems with Concerto’. In: CCGrid 2020 : 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing. Melbourne, Australia, 4th June 2020. DOI: [10.1109/CCGrid49817.2020.00-59](https://doi.org/10.1109/CCGrid49817.2020.00-59). URL: <https://hal.inria.fr/hal-02535077>.
- [7] I. Daoudi, P. Virouleau, T. Gautier, S. Thibault and O. Aumage. ‘sOMP: Simulating OpenMP Task-Based Applications with NUMA Effects’. In: *OpenMP: Portable Multi-Level Parallelism on Modern Systems (IWOMP 2020)*. IWOMP 2020 - 16th International Workshop on OpenMP. Vol. 12295. LNCS. Austin / Virtual, United States, 1st Sept. 2020. DOI: [10.1007/978-3-030-58144-2_13](https://doi.org/10.1007/978-3-030-58144-2_13). URL: <https://hal.inria.fr/hal-02933803>.
- [8] T. Gautier and J. V. F. Lima. ‘XKBlas: a High Performance Implementation of BLAS-3 Kernels on Multi-GPU Server’. In: 2020 28th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP). Västerås, Sweden, 11th Mar. 2020, pp. 1–8. DOI: [10.1109/PDP50117.2020.00008](https://doi.org/10.1109/PDP50117.2020.00008). URL: <https://hal.inria.fr/hal-03121583>.
- [9] H. Hadjur, D. Ammar and L. Lefèvre. ‘Analysis of Energy Consumption in a Precision Beekeeping System’. In: IoT ’20 - 10th International Conference on the Internet of Things. Malmö, Sweden, 7th Oct. 2020, pp. 1–11. DOI: [10.1145/3410992.3411010](https://doi.org/10.1145/3410992.3411010). URL: <https://hal.inria.fr/hal-02973772>.
- [10] Z. Houmani, D. Balouek-Thomert, E. Caron and M. Parashar. ‘Enhancing microservices architectures using data-driven service discovery and QoS guarantees’. In: CCGrid 2020 - 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing. Melbourne, Australia, 2nd Nov. 2020, pp. 1–10. URL: <https://hal.inria.fr/hal-02523442>.
- [11] D. Mvondo, A. Tchana, R. Lachaize, D. Hagimont and N. De Palma. ‘Fine-Grained Fault Tolerance For Resilient pVM-based Virtual Machine Monitors’. In: *Proceedings of DSN 2020*. DSN 2020 - 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks. Valencia, France: <https://ieeexplore.ieee.org/document/9153354>, 29th June 2020, pp. 197–208. DOI: [10.1109/DSN48063.2020.00037](https://doi.org/10.1109/DSN48063.2020.00037). URL: <https://hal.archives-ouvertes.fr/hal-02959252>.
- [12] Y. Ren, G. Liu, V. Nitu, W. Shao, R. Kennedy, G. Parmer, T. Wood and A. Tchana. ‘Fine-Grained Isolation for Scalable, Dynamic, Multi-tenant Edge Clouds’. In: 2020 {USENIX} Annual Technical Conference ({USENIX} {ATC} 20). Virtual, France, 15th July 2020. URL: <https://hal.inria.fr/hal-03162363>.
- [13] F. Rodrigo De Souza, A. Da Silva Veith, M. Dias de Assuncao and E. Caron. ‘Scalable Joint Optimization of Placement and Parallelism of Data Stream Processing Applications on Cloud-Edge Infrastructure’. In: ICSOC 2020 - 18th International Conference on Service Oriented Computing. Dubai, United Arab Emirates, 14th Dec. 2020. URL: <https://hal.inria.fr/hal-02926464>.
- [14] F. Rodrigo De Souza, M. Dias de Assuncao, E. Caron and A. da Silva Veith. ‘An Optimal Model for Optimizing the Placement and Parallelism of Data Stream Processing Applications on Cloud-Edge Computing’. In: SBAC-PAD 2020 - IEEE 32nd International Symposium on Computer Architecture and High Performance Computing. Porto, Portugal, 8th Sept. 2020. URL: <https://hal.inria.fr/hal-02926459>.

- [15] L. Turpin, J. Rouzaud-Cornabas, T. Gautier and C. Pérez. ‘P-Aevol: an OpenMP Parallelization of a Biological Evolution Simulator, Through Decomposition in Multiple Loops’. In: *OpenMP: Portable Multi-Level Parallelism on Modern Systems*. 16th International Workshop on OpenMP. OpenMP: Portable Multi-Level Parallelism on Modern Systems. Austin, United States, 1st Sept. 2020, pp. 52–66. DOI: [10.1007/978-3-030-58144-2_4](https://doi.org/10.1007/978-3-030-58144-2_4). URL: <https://hal.inria.fr/hal-02962838>.

Conferences without proceedings

- [16] N. Baillon-Bachoc, E. Caron, A. Chevalier and A.-L. Vion. ‘Providing Software Asset Management Compliance in Green Deployment Algorithm’. In: SETCAC 2020 - Symposium on Emerging Topics in Computing and Communications. Chennai, India, 14th Oct. 2020, pp. 1–14. URL: <https://hal.archives-ouvertes.fr/hal-03034059>.
- [17] G. Bista, E. Caron and A.-L. Vion. ‘A Study On Optimizing VNF Software Cost’. In: GIIS 2020 - Global Information Infrastructure and Networking Symposium. Tunis, Tunisia, 28th Oct. 2020, pp. 1–4. DOI: [10.1109/GIIS50753.2020.9248486](https://doi.org/10.1109/GIIS50753.2020.9248486). URL: <https://hal.archives-ouvertes.fr/hal-03027945>.

Doctoral dissertations and habilitation theses

- [18] A. Chevalier. ‘Optimization of software license placement in the Cloud for economical and efficient deployment’. Université de Lyon, 24th Nov. 2020. URL: <https://tel.archives-ouvertes.fr/tel-03099617>.
- [19] F. R. De Souza. ‘Scheduling Solutions for Data Stream Processing Applications on Cloud-Edge Infrastructure’. Université de Lyon, 10th Dec. 2020. URL: <https://tel.archives-ouvertes.fr/tel-03122824>.

Reports & preprints

- [20] M. Chardet, H. Coullon, C. Pérez, D. Pertin, C. Servantie and S. Robillard. *Enhancing Separation of Concerns, Parallelism, and Formalism in Distributed Software Deployment with Madeus*. 2nd June 2020. URL: <https://hal.inria.fr/hal-02737859>.
- [21] A. Kong Win Chang, Y. Caniou, E. Caron and Y. Robert. *Budget-aware workflow scheduling with DIET*. Inria Grenoble Rhône-Alpes, Dec. 2020. URL: <https://hal.inria.fr/hal-03080468>.

Other scientific publications

- [22] C. Bonamy, C. Boudinet, L. Bourgès, K. Dassas, L. Lefèvre and F. Vivat. *Je code : les bonnes pratiques en éco-conception de service numérique à destination des développeurs de logiciels*. 17th Nov. 2020. URL: <https://hal.archives-ouvertes.fr/hal-03009741>.

10.2 Cited publications

- [23] R. Ge, X. Feng, S. Song, H.-C. Chang, D. Li and K. W. Cameron. ‘PowerPack: Energy Profiling and Analysis of High-Performance Systems and Applications’. In: *IEEE Trans. Parallel Distrib. Syst.* 21.5 (May 2010), pp. 658–671. DOI: [10.1109/TPDS.2009.76](https://doi.org/10.1109/TPDS.2009.76). URL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4906989.
- [24] A. Geist and S. Dosanjh. ‘IESP Exascale Challenge: Co-Design of Architectures and Algorithms’. In: *Int. J. High Perform. Comput. Appl.* 23.4 (Nov. 2009), pp. 401–402. DOI: [10.1177/1094342009347766](https://doi.org/10.1177/1094342009347766). URL: <http://dx.doi.org/10.1177/1094342009347766>.
- [25] W. Gropp, S. Huss-Lederman, A. Lumsdaine, E. Lusk, B. Nitzberg, W. Saphir and M. Snir. *MPI: The Complete Reference – The MPI-2 Extensions*. 2nd ed. Vol. 2. ISBN 0-262-57123-4. The MIT Press, Sept. 1998.

-
- [26] H. Kimura, T. Imada and M. Sato. 'Runtime Energy Adaptation with Low-Impact Instrumented Code in a Power-Scalable Cluster System'. In: *Proceedings of the 2010 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing*. CCGRID '10. Washington, DC, USA: IEEE Computer Society, 2010, pp. 378–387.
- [27] G. Madec. *NEMO ocean engine*. Note du Pole de modélisation 27. ISSN No 1288-1619. France: Institut Pierre-Simon Laplace (IPSL), 2008.
- [28] OpenACC. *The OpenACC Application Programming Interface*. Version 1.0. Nov. 2011. URL: <http://www.openacc-standard.org>.
- [29] OpenMP Architecture Review Board. *OpenMP Application Program Interface*. Version 3.1. July 2011. URL: <http://www.openmp.org>.
- [30] B. Rountree, D. K. Lownenthal, B. R. de Supinski, M. Schulz, V. W. Freeh and T. Bletsch. 'Adagio: Making DVS Practical for Complex HPC Applications'. In: *Proceedings of the 23rd international conference on Supercomputing*. ICS '09. New York, NY, USA: ACM, 2009, pp. 460–469.
- [31] C. Szyperski. *Component Software - Beyond Object-Oriented Programming*. 2nd ed. Addison-Wesley / ACM Press, 2002, p. 608.
- [32] S. Valcke. 'The OASIS3 coupler: a European climate modelling community software'. In: *Geoscientific Model Development* 6 (2013). doi:10.5194/gmd-6-373-2013, pp. 373–388.