

RESEARCH CENTRE

Grenoble - Rhône-Alpes

IN PARTNERSHIP WITH:

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

2021

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	4
2.1 Presentation	4
2.2 Objectives	5
3 Research program	5
3.1 Energy Application Profiling and Modeling	5
3.2 Data-intensive Application Profiling, Modeling, and Management	6
3.3 Resource-Agnostic Application Description Model	6
3.4 Application Mapping and Scheduling	7
3.4.1 Application Mapping and Software Deployment	7
3.4.2 Non-Deterministic Workflow Scheduling	7
3.4.3 Software Asset Management	7
4 Application domains	7
4.1 Overview	7
4.2 Climatology	8
4.3 Astrophysics	8
4.4 Bioinformatics	8
5 Social and environmental responsibility	9
5.1 Footprint of research activities	9
5.2 Impact of research results	9
6 Highlights of the year	9
7 New software and platforms	9
7.1 New software	9
7.1.1 SimGrid	9
7.1.2 libkomp	10
7.1.3 XKBLAS	10
7.1.4 Qirinus-Orchestra	11
7.1.5 execo	11
7.1.6 Kwollect	12
7.2 New platforms	12
7.2.1 Platform: Grid'5000	12
7.2.2 Platform: Leco	13
7.2.3 Platform: SILECS	13
7.2.4 Platform: SLICES	13
8 New results	13
8.1 Energy Efficiency in HPC and Large Scale Distributed Systems	13
8.1.1 Experimental Workflow for Energy and Temperature Profiling on HPC Systems	13
8.1.2 Energy Consumption and Energy Efficiency in a Precision Beekeeping System	14
8.1.3 Metrics Collection for Experiments at Scale	14
8.2 Modeling and Simulation of Parallel Applications and Distributed Infrastructures	14
8.2.1 SDN-based Fog and Cloud Interplay for Stream Processing	14
8.2.2 Latency-Aware Strategies for Deploying Data Stream Processing Applications on Large Cloud-Edge Infrastructure	15
8.2.3 Budget-aware static scheduling of stochastic workflows with DIET	15
8.3 Edge and Cloud Resource Management	15
8.3.1 Toward Safe and Efficient Reconfiguration with Concerto	15

8.3.2	Heuristic for license-aware, performant and energy efficient deployment of multiple software in Cloud architecture	16
8.3.3	Enabling microservices management for Deep Learning applications across the Edge-Cloud Continuum	16
8.4	HPC Applications and Runtimes	16
8.4.1	Evaluation of two topology-aware heuristics on level-3 BLAS library for multi-GPU platforms	16
8.4.2	Communication-Aware Task Scheduling Strategy in Hybrid MPI+OpenMP Applications	17
8.4.3	X-Aevol: GPU implementation of an evolutionary experimentation simulator	17
9	Bilateral contracts and grants with industry	17
9.1	Bilateral grants with industry	17
10	Partnerships and cooperations	18
10.1	International initiatives	18
10.1.1	Inria International Labs	18
10.1.2	Participation in other International Programs	18
10.2	European initiatives	19
10.2.1	Horizon Europe	19
10.3	National initiatives	22
10.4	Regional initiatives	22
10.4.1	CPER	22
10.4.2	Action Exploratoire Inria	23
11	Dissemination	23
11.1	Promoting scientific activities	23
11.1.1	Scientific events: organisation	23
11.1.2	Scientific events: selection	23
11.1.3	Journal	24
11.1.4	Invited talks	24
11.1.5	Scientific expertise	24
11.1.6	Research administration	24
11.2	Teaching - Supervision - Juries	25
11.2.1	Teaching	25
11.2.2	Supervision	26
11.2.3	Juries	27
11.3	Popularization	27
11.3.1	Education	27
11.3.2	Interventions	27
12	Scientific production	27
12.1	Major publications	27
12.2	Publications of the year	28
12.3	Cited publications	29

Project-Team AVALON

Creation of the 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
- A1.3.5. – 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
- A2.6.2. – Middleware
- 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]
- 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]
- Etienne Mauffret [École Normale Supérieure de Lyon, ATER]

PhD Students

- Maxime Agusti [OVHCloud, from Dec 2021]
- Adrien Berthelot [OCTO Technology, CIFRE, from Nov 2021]
- Ghoshana Bista [Orange Labs, CIFRE]
- Arthur Chevalier [École Normale Supérieure de Lyon, Jan 2021]
- Idriss Douadi [INRIA, co-advised with STROM [Inria Bordeaux]]
- Hugo Hadjur [Aivancity]
- Zeina Houmani [École Normale Supérieure de Lyon]
- Mathilde Jay [Université Grenoble Alpes, from Oct 2021, co-advised with DataMove [Inria Grenoble]]
- Lucien Ndjie Ngale [Univ Jules Vernes Picardie]
- Vladimir Ostapenco [Inria, from Dec 2021, co-advised with Myriads [Inria Rennes]]
- Romain Pereira [CEA]
- Pierre Etienne Polet [Thales, CIFRE]
- Laurent Turpin [Inria, co-advised with Beagle [Inria Lyon]]

Technical Staff

- Thierry Arrabal [CNRS, Engineer, from Sep 2021]
- Lucas Betencourt [École Normale Supérieure de Lyon, Engineer, from Oct 2021]
- Abderahman Cheniour [Inria, Engineer, from May 2021 until Oct 2021]
- Arthur Chevalier [Inria, Engineer, from Nov 2021]
- Simon Delamare [CNRS, Engineer]
- Matthieu Imbert [Inria, Engineer]

- Pierre Jacquot [Inria, Engineer, from Sep 2021]
- Patrice Kirmizigul [Inria, Engineer, until May 2021]
- David Loup [Inria, Engineer, until Mar 2021]
- Jean-Christophe Mignot [CNRS, Engineer]
- Mahendra Paipuri [Inria, Engineer]

Interns and Apprentices

- Imene Agusti [OVHCloud, Intern, from Dec 2021, co-advised with DataMove (Inria Grenoble)]
- Pierre Jacquot [Inria, from Feb 2021 until Jul 2021]
- Sylvere Kanapa [Inria, from May 2021 until Jul 2021]
- Binjamyn Mairesse [Inria, from May 2021 until Aug 2021]
- Thibaut Modrzyk [Inria, from May 2021 until Jul 2021]
- Catalin Moldovan [Inria, from Apr 2021 until Jul 2021]
- Augusta Mukam [École Normale Supérieure de Lyon, from Mar 2021 until Sep 2021]
- Tushita Ramkaran [Inria, from Jun 2021 until Aug 2021]

Administrative Assistant

- Evelyne Blesle [Inria]

External Collaborators

- Doreid Ammar [Aivancity]
- Frédéric Suter [CNRS, Senior 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), datacenters (Clouds) with IoT (Edge/Fog), 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, 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 the previous section, the AVALON project aims at making progress on four complementary research axes: energy, data, programming models and runtimes, application scheduling.

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 and Runtimes 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.

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 [17] 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 [20], phase detection for specific HPC applications [24], *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 [19] exposes the number of nodes but hides the complexity of network topology behind a set of collective operations; OpenMP [23] simplifies the management of threads on top of a shared memory machine while OpenACC [22] 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 [18]. 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 [25] 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 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.

3.4.1 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.

3.4.2 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.

3.4.3 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), to BioSyL, a federative research structure about Systems Biology of the University of Lyon, or to the SKA project. 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 [21] for ocean modelization), code-coupling applications (*e.g.*, the OASIS coupler [26] 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, since 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** 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** will produce 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** 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**, 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.

The SKA project is an international effort to build and operate the world's largest radiotelescopes covering all together the wide frequency range between 50 MHz and 15.4 GHz. The scale of the SKA project represents a huge leap forward in both engineering and research & development towards building and delivering a unique Observatory, whose construction has officially started on July 2021. The SKA Observatory is the second intergovernmental organisation for ground-based astronomy in the world, after the European Southern Observatory. AVALON participates to the activities of the PlaNet team of SKAO that is dedicated to platforms (benchmarking, co-design, profiling, etc) and network issues.

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**, 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) in particular through the ANR MapReduce project where the team focuses on a bio-chemistry application dealing with protein structure analysis. AVALON is also working with the Inria **Beagle** team on artificial evolution and computational biology as the challenges are around high performance computation and data management.

5 Social and environmental responsibility

5.1 Footprint of research activities

Through its research activities on energy efficiency and on energy and environmental impacts reductions, Avalon tries to reduce some impacts of distributed systems.

In May 2021, Laurent Lefevre has participated in the "Atelier Sens" proposed by some Inria colleagues which helps exchanging and discussing impact of research activities. Laurent Lefevre is also involved in the steering committee of the **EcoInfo** GDS CRNS group which deals with eco-responsibility of ICT.

5.2 Impact of research results

Rebound effects must be taken into account while proposing new approaches and solutions in ICT. This is a challenging task. Laurent Lefevre has co-organized in November 2020, a workshop from the Entretiens Jacques Cartier on the topic of "**Rebound effects in ICT. How to detect them? How to measure them? How to avoid them?**". Another event on such research challenge will be organized in May 2022.

6 Highlights of the year

The **SLICES RI** has been included in the **ESFRI roadmap**.

7 New software and platforms

7.1 New software

7.1.1 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 3 major releases in 2021. A new API was introduced to create the platform descriptions directly from the source code instead of XML, providing much more expressiveness to the experimenters. SMPI now reports memory leaks and correctly diagnoses API misuses, which makes it even more adapted to teaching settings. The documentation was thoroughly overhauled to ease the use of the framework. We also pursued our efforts to improve the overall framework, through bug fixes, code refactoring and other software quality improvement.

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

Contact: Martin Quinson

Participants: Adrien Lebre, 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

7.1.2 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

7.1.3 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

7.1.4 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 paradigms:

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

7.1.5 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.

News of the Year: Many bugfixes, improvements in Python3 compatibility, and migration from Inria forge to Inria gitlab.

URL: <https://gitlab.inria.fr/mimbert/execo>

Contact: Matthieu Imbert

Participants: Florent Chuffart, Laurent Pouilloux, Matthieu Imbert

7.1.6 Kwollect

Keywords: Monitoring, Power monitoring, Energy, Infrastructure software, Sensors

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 2021, Kwollect is the new monitoring tool available under Grid'5000 and has supersede other existing monitoring solution on the infrastructure.

Kwollect is also used in Immersion 4 project where Avalon team (Inria Lyon) is involved

An article describing Kwollect has been 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>

Publication: hal-03236421

Contact: Simon Delamare

7.2 New platforms

7.2.1 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/>

7.2.2 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

7.2.3 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: Christian Perez
- URL: <https://www.silecs.net/>

7.2.4 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>

8 New results

8.1 Energy Efficiency in HPC and Large Scale Distributed Systems

8.1.1 Experimental Workflow for Energy and Temperature Profiling on HPC Systems

Participants: Laurent Lefèvre.

Despite recent advances in improving the performance of high performance computing (HPC) and distributed systems, power dissipation and thermal cooling challenges persist, impacting their total cost of ownership. Making HPC systems more energy and thermal efficient will require understanding of individual power dissipation and temperature contributions of multiple hardware system components

and their accompanying software. In this joint work with Myriads team under the IPL HAc-Specis project, we present an experimental workflow for energy and temperature profiling on systems running parallel applications. It allows full and dynamic control over the execution of applications for the entire frequency range. Through its use, we show that the energy response to frequency scaling is highly dependent on the workload characteristics and it is convex in nature with an optimal frequency point. During the course of our experimentation, we encountered a non-intuitive finding, where we observed that the tested low-power processor is consuming more power on average than the standard processor [13].

8.1.2 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 work, within the 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. 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. In 2021, we published a survey dedicated on challenges in precision beekeeping systems [5].

8.1.3 Metrics Collection for Experiments at Scale

Participants: Simon Delamare.

It has become a common requirement for testbeds to provide a service in charge of collecting and exposing metrics, thus assisting experimenters with the central task of data collection. This work [9] describes Kwollect, a service designed and developed in the context of the Grid'5000, to collect infrastructure metrics (including high-frequency wattmeters) and expose them to experimenters. Kwollect scales to high frequencies of metrics collection for hundreds of nodes. It can also be leveraged by the experimenter to collect custom metrics.

8.2 Modeling and Simulation of Parallel Applications and Distributed Infrastructures

8.2.1 SDN-based Fog and Cloud Interplay for Stream Processing

Participants: Laurent Lefèvre.

This work focuses on SDN-based approaches for deploying stream processing workloads on heterogeneous environments comprising wide-area network, cloud and fog resources. The main contribution [6] consists in dynamic workload placement algorithms operating on the stream processing request with latency constraints. Provisioning of computing infrastructure is performed by exploiting the interplay between fog and cloud under the constraint of limited network capacity. The algorithms aim at maximizing the ratio of successfully handled requests by effective utilization of available resources while meeting application latency constraints. Experiments demonstrate that the goal can be achieved by detailed analysis of requests and ensuring optimal utilization of both computing and network resources. As a result, up to 40 % improvement over the reference algorithms in terms of success ratio is observed. This research is a joint work with researchers from AGH University from Krakow, Poland (Michal Rzepka, Piotr Borylo and Artur Lason) and Ecole de Technologie Supérieure from Montreal, Canada (Marcos Dias de Assuncao).

8.2.2 Latency-Aware Strategies for Deploying Data Stream Processing Applications on Large Cloud-Edge Infrastructure

Participants: Laurent Lefèvre.

Internet of Things (IoT) applications often require the processing of data streams generated by devices dispersed over a large geographical area. Traditionally, these data streams are forwarded to a distant cloud for processing, thus resulting in high application end-to-end latency. Recent work explores the combination of resources located in clouds and at the edges of the Internet, called cloud-edge infrastructure, for deploying Data Stream Processing (DSP) applications. Most previous work, however, fails to scale to very large IoT settings. This work [4] introduces deployment strategies for the placement of DSP applications on to cloud-edge infrastructure. The strategies split an application graph into regions and consider regions with stringent time requirements for edge placement. The proposed Aggregate End-to-End Latency Strategy with Region Patterns and Latency Awareness (AELS+RP+LA) decreases the number of evaluated resources when computing an operator's placement by considering the communication overhead across computing resources. Simulation results show that, unlike the state-of-the-art, AELS+RP+LA scales to environments with more than 100k resources with negligible impact on the application end-to-end latency. This research is a joint work with researchers from Ecole de Technologie Supérieure from Montreal, Canada (Marcos Dias de Assuncao) and University of Toronto, Canada (Alexandre Veith).

8.2.3 Budget-aware static scheduling of stochastic workflows with DIET

Participants: Yves Caniou, Eddy Caron.

First we proposed a new Cloud platform model and designed some budget-aware static algorithms to schedule stochastic workflows. On the strength of this result, in [2, 7] we have introduced a new scheduling functionality for DIET, and provided the user with a set of tools to implement, and experiment with, static scheduling algorithms for workflows. We then used this new functionality to compare the executions of ten static algorithms for scientific applications from the Pegasus benchmark, using both a simulator and the testbed Grid'5000. Both types of experiments gave similar results, validating our models and DIET improvements.

8.3 Edge and Cloud Resource Management

8.3.1 Toward Safe and Efficient Reconfiguration with Concerto

Participants: 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], 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 this paper, 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.

8.3.2 Heuristic for license-aware, performant and energy efficient deployment of multiple software in Cloud architecture

Participants: Eddy Caron, Arthur Chevalier.

In [8], we introduced several new advances. First, we presented a deployment problem with Software Asset Management considerations but with revisited license consumption computation to adjust it to real-world usage. Then, we proposed representation and a structure for this deployment problem in a tree-shaped manner. We gave a proof of NP-completeness of the decision problem representing our deployment. We then introduced the main contribution: a heuristic that optimizes several criteria (e.g. we tackle the problem to reduce the energy consumption) and manages to get good results with the deployment of one product on the Cloud compared to other heuristics. We can see from the evaluation that the GreenSAM heuristic needs a very little amount of memory to obtain near-optimal results and that it does it quickly. Besides, compared to other heuristics, GreenSAM ensures compliance at deployment time by removing servers that will put us in a non-compliant state.

8.3.3 Enabling microservices management for Deep Learning applications across the Edge-Cloud Continuum

Participants: Eddy Caron, Zeina Houmani.

Current data analytics systems tend to approach resource and data management solutions independently and without focusing on the entire analytics pipeline. With the increase in the quality and amount of data, it became challenging to meet the application's objectives when dealing with multiple data sources and limited resources of heterogeneous capabilities. We proposed a system that combines data and resource management solutions to support Deep Learning workflows on an Edge-to-Cloud environment. By adapting the resolution of incoming load and distributing automatically the analysis across the continuum, the system manages latency-accuracy trade-offs to meet application performance. We illustrated the system viability through the implementation and deployment of an object detection use case on Grid'5000. As we can see in [11], evaluation results showed a gain in average system makespan reaching up to 54.4% compared to a cloud-only pipeline configuration in a multi-users scenario.

8.4 HPC Applications and Runtimes

8.4.1 Evaluation of two topology-aware heuristics on level-3 BLAS library for multi-GPU platforms

Participants: Thierry Gautier.

Nowadays 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 with multiple GPUs. Our paper [10] presents two runtime heuristics to gain in performance when task based programs are performed on heterogeneous architecture such as multi-GPU systems. The first is a topology-aware policy to take into account the heterogeneity of the high speed links that interconnect GPUs. The second is an optimistic heuristic that favors communication between devices. These have been implemented in the XKBLAS library BLAS-3 library. We made experiments on a NVIDIA DGX-1 with up to 8 GPUs V100 on a set of Basic Linear Algebra Subroutines. Experimental results on kernels showed that XKBlas outperformed most implementations including the overhead of creation and scheduling of dynamic tasks.

8.4.2 Communication-Aware Task Scheduling Strategy in Hybrid MPI+OpenMP Applications

Participants: Romain Peirera, Thierry Gautier.

While task-based programming, such as OpenMP, is a promising solution to exploit large HPC compute nodes, it has to be mixed with data communications like MPI. However, performance or even more thread progression may depend on the underlying runtime implementations. In the paper [12], we focus on enhancing the application performance when an OpenMP task blocks inside MPI communications. This technique requires no additional effort on the application developers. It relies on an online task reordering strategy that aims at running first tasks that are sending data to other processes. We evaluate our approach on a Cholesky factorization and show that we gain around 19% of execution time on an Intel Skylake compute nodes machine-each node having two 24-core processors.

8.4.3 X-Aevol: GPU implementation of an evolutionary experimentation simulator

Participants: Laurent Turpin, Thierry Gautier.

X-Aevol is the GPU port of the Aevol model, a bio-inspired genetic algorithm designed to study the evolution of microorganisms and its effects on their genome structure. This model is used for in-silico experimental evolution that requires the computation of populations of thousands of individuals during tens of millions of generations. As the model is extended with new features and experiments are conducted with larger populations, computational time becomes prohibitive. In [14] we present X-Aevol as a response to the need of more computational power. It was designed to leverage the massive parallelization capabilities of GPU. As Aevol exposes an irregular and dynamic computational pattern, it was not a straightforward process to adapt it for massively parallel architectures. We present how we have adapted the Aevol underlying algorithms to GPU architectures and we implement our new algorithms with CUDA programming language and test them on a representative benchmark of Aevol workloads. To conclude, we present our performance evaluation on NVIDIA Tesla V100 and A100. We show how we reach a speed-up of 1,000 over a sequential execution on a CPU and the speed-up gain up to 50% from using the newer Ampere micro-architecture in comparison with Volta one.

9 Bilateral contracts and grants with industry

9.1 Bilateral grants with industry

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.

Orange We have a collaboration with Orange. This collaboration is sealed through a CIFRE PhD grant. The research of the PhD student (Ghoshana Bista) focuses on the software asset management dedicated to the VNF (Virtual Network Function).

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 on 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.

TotalLinux We have a collaboration with TotalLinux around the data center project **Itrium**. More specially we study the impact, the energy consumption, the behavior and the performances of new architectures based on immersion cooling.

10 Partnerships and cooperations

10.1 International initiatives

10.1.1 Inria International Labs

JLESC

Title: Joint Laboratory for Extreme Scale Computing

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. 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.

10.1.2 Participation in other International Programs

SKA

Participants: Mahendra Paipuri, Christian Perez.

Title: Square Kilometer Array (SKA)

Summary: The Avalon team collaborates with SKA Organization that has been responsible for coordinating the global activities towards the SKA in the pre-construction phase.

Rutgers

Participants: Eddy Caron, Zeina Houmani, Laurent Lefèvre.

Title: Rutgers University, New-Jersey (USA)

Summary: In the context of our collaboration with the RDI² team (Rutgers University), we co-advise the PhD of Zeina Houmani to build a Data-driven microservices architecture for Deep Learning applications. The funding of this PhD was 50/50 between the US and French partners.

10.2 European initiatives

10.2.1 Horizon Europe

EoCoE-II

Participants: Thierry Gautier, Christian Perez.

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.

PRACE-6IP

Participants: Christian Perez.

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

Duration: May 2019 - Dec 2021

Coordinator: FORSCHUNGSZENTRUM JULICH GMBH (Germany)

Partners:

- KADEMIA GORNICZO-HUTNICZA IM. STANISLAWA 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.

SLICES-DS

Participants: Laurent Lefèvre, Christian Perez.

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.

10.3 National initiatives

Inria Large Scale Initiative

Défi Inria OVHCloud

Participants: Eddy Caron, Laurent Lefèvre, Christian Perez.

A joint collaboration between Inria and OVH Cloud company on the topic challenge of frugal cloud has been launched in October 2021. It addresses several scientific challenge on the eco-design of cloud frameworks and services for large scale energy and environmental impact reduction. Laurent Lefèvre is the scientific animator of this project. Some Avalon PhD students are involved in this Inria Large Scale Initiative (Défi) : Maxime Agusti and Vladimir Ostanpenco.

10.4 Regional initiatives

10.4.1 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.

10.4.2 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.

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

General chair, scientific chair

- Laurent Lefevre was co Program Chair of CCGrid 2021 conference: The 21st IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing, Melbourne, Australia, May 10-13, 2021

Member of the organizing committees

- Christian Perez was member of the Organizing Committee of the French Journées Calcul Données (Dijon, 13-15 Dec 2021).

11.1.2 Scientific events: selection

Chair of conference program committees

- Christian Perez was Project Posters Deputy Chair in ISC High Performance 2021. He is Project Posters Chair in ISC High Performance 2022.

Member of the conference program committees

- Yves Caniou was a program committee member for 2021 International Conference on Computational Science and Its Applications (ICCSA'21).
- Eddy Caron was Vice Chairs for CCGRID2021 track Scheduling and Resource Management. He was a program committee member for CLOSER'2021.
- Thierry Gautier was a program committee member for 2022 36th IEEE International Parallel & Distributed Processing Symposium (IPDPS'21).
- Christian Perez was a technical program committee member for 2021 Annual Modeling and Simulation Conference (ANNSIM'21).

11.1.3 Journal

Reviewer - reviewing activities

- Eddy Caron was reviewer for Concurrency and Computation: Practice and Experience (2 manuscripts), International Transactions in Operational Research, IEEE Transactions on computers, Journal of Parallel and Distributed Computing, Cluster Computing
- Christian Perez was reviewer for Future Generation Computer Systems.

11.1.4 Invited talks

- Laurent Lefevre was invited to give a talk during the following events:
 - "Des nuages, du brouillard, des impacts !", Bibliothèque Municipale de Lyon, December 7, 2021
 - "Efficiency, proportionality, complexity, frugality.. how to find a way towards sustainable computing ?", Computing Systems Week, HiPEAC, October 25, 2021
 - "S'il vous plait... dessine moi un service numérique ! Please... draw me a digital service !", Journées Scientifiques Inria (online), July 1, 2021
 - "Impacts environnementaux du numérique : phase d'usage... en route vers la sobriété numérique ? Mesurer -> Estimer -> Réduire", GreenIT school, Anglet (online), June 30, 2021
 - "Impacts environnementaux du numérique : en route vers la sobriété numérique ? Mesurer -> Estimer -> Réduire", MIAI Days, Grenoble, France, May 4, 2021
 - "Impacts environnementaux du numérique : Mesurer -> Estimer -> Riposter", 22 eme Conference de la Société Française de Recherche Opérationnelle et d'Aide à la Décision (ROADEF), Mulhouse, France (online), April 28, 2021

11.1.5 Scientific expertise

- Yves Caniou evaluated a project for the Agence Nationale de la recherche (ANR).
- Christian Perez evaluated 3 projects for the French Direction générale de la Recherche et de l'Innovation.

11.1.6 Research administration

- 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.
- Laurent was member of the jury for recruiting CRCN candidates in Inria Lyon center. Laurent Lefevre is Inria Scientific International correspondent for Inria Lyon center. He is member of Inria Lyon initiative committee. Laurent Lefevre participates in LIP Laboratory direction team. He is elected member of the LIP laboratory council (ENS Lyon). 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 co-director of the CNRS GDS EcoInfo group. He is the responsible of M2 training period (ENS Lyon) and local correspondent of the radar Inria process.

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

- Licence: Eddy Caron, Programmation, 48h, L3, ENS de Lyon. France.
- Agreg Info (FEADéP), Operating System and Network, 15h, Agreg, ENS de Lyon. France.
- Agreg Info (FEADéP), TP Programmation, 11h, Agreg, ENS de Lyon. France.
- Master: Eddy Caron, Distributed System, 30h, M1, ENS de Lyon. France.
- Master: Eddy Caron, Large scale sustainable distributed resource management, 16h, M2, ENS de Lyon. France.
- Licence: Yves Caniou, Algorithmique programmation impérative initiation, 60h, niveau L1, Université Claude Bernard Lyon 1, France.
- Licence: Yves Caniou, Algorithmique et programmation récursive, 36h, niveau L1, Université Claude Bernard Lyon 1, France.
- Licence: Yves Caniou, supervision 2 month internship, niveau L2, Université Claude Bernard Lyon 1, France.
- Licence: Yves Caniou, supervision 3 month internship, niveau L2, 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, Réseaux, 12h, niveau L3, Université Claude Bernard Lyon 1, France.
- Licence: Yves Caniou, Systèmes d'information documentaire, 20h, niveau L3, Université Claude Bernard Lyon 1, France.
- Master: Yves Caniou, Projet Orientation Master, 9h, niveau M1, Université Claude Bernard Lyon 1, France.
- Master: Yves Caniou, Sécurité, 30h 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, Sécurité, 20h, niveau M2, IGA Casablanca, Maroc.
- Master: Laurent Lefèvre, Parallélisme, 12h, niveau M1, Université Lyon 1, France.
- CAPES Informatique : Laurent Lefèvre, Numérique responsable, 3h, Université Lyon1, France
- Master: Thierry Gautier, Introduction to HPC, 20h, niveau M2, INSA Lyon, 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, Responsable 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
- Licence: Elise Jeanneau, Introduction Réseaux et Web, 36h, niveau L1, Université Lyon 1, France.
- Licence: Elise Jeanneau, Réseaux, 34h, niveau L3, Université Lyon 1, France.
- Master: Elise Jeanneau, Algorithmes distribués, 36h, niveau M1, Université Lyon 1, France.
- Master: Elise Jeanneau, Réseaux, 21h, niveau M1, Université Lyon 1, France.
- Master: Elise Jeanneau, Compilation et traduction de programmes, 22h, niveau M1, Université Lyon 1, France.
- Master: Elise Jeanneau, Algorithmes pour les systèmes distribués dynamiques, 10h, niveau M2, ENS de Lyon, France.

11.2.2 Supervision

- PhD: Idriss Daoudi, Simulating OpenMP program, October 2018, Samuel Thibault (Univ-Bordeaux, Storm team, Bordeaux, dir) and Thierry Gautier (INRIA, Avalon team, co-dir).
- PhD: 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: Laurent Turpin, Mastering Code Variation and Architecture Evolution for HPC application, October 2019, and Thierry Gautier (INRIA, Avalon team, dir) and Jonathan Rouzaud-Cornabas (INSA, Beagle 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: 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).

- PhD In progress: Mathilde Jay, Frugal IA edge services, Sept. 2021, Laurent Lefevre (co-dir), Denis Trystram (DataMove, LIG, dir).
- PhD In progress: Adrien Berthelot, Assisted evaluation of environmental impacts of digital services, November 2021, Eddy Caron (dir), Laurent Lefevre (co-dir), Christian Faure (Octo technology, co-dir).
- PhD In progress: Vladimir Ostapenco, Dealing with energy leverages in Cloud, December 2021, Laurent Lefevre (dir), Anne-Cécile Orgerie (co-dir).

11.2.3 Juries

- Eddy Caron was reviewer and member of the HDR defense committee of Flavien Vernier, Université Savoie Mont Blanc, France, 25 June 2021. He was member of the PhD defense committee of Clement Elbaz, Université de Rennes, France, 30 March 2021. He was reviewer and member of the PhD defense committee of Medhi Belkhiria, Université de Rennes, France, 25 November 2021.
- Thierry Gautier was reviewer and member of the PhD defense committee of Andrès RUBIO PROANO, Bordeaux, France, 7 October 2021.
- Christian Perez was reviewer and member of the HDR defense committee of Jeremy Buisson, St Cyr Coëquidan, France, 2 February 2021. He was president of the HDR defense committee of Flavien Vernier, Annecy, France, 25 June 2021.
- Laurent Lefevre was reviewer and member of the PhD defense committee of Hamidreza Arkian, Université de Rennes 1, December 7, 2021. He was reviewer and member of the PhD defense committee of Matthieu Stoffel, Université Grenoble Alpes, October 1, 2021. He was member of the PhD defense committee of Simon Dembele, Ecole Nationale Supérieure de Mécanique et d'Aérotechnique, Poitiers, July 8, 2021. He was member of the PhD defense committee of Teylo Gouveia Lima, Universidade Federal Fluminense, Brasil, March 26, 2021.

11.3 Popularization

11.3.1 Education

- Yves Caniou co-organized the 4th Edition of Le Campus du Libre, on saturday Nov. 6 2021 at the Nautibus, Université Claude Bernard Lyon 1, La Doua.

11.3.2 Interventions

- Yves Caniou has been invited for a talk and the animation of a workshop by Nouvelle conférence du Disrupt'Campus sur les logiciels et services libres, in a partnership between Université de Lyon and the Métropole de Lyon, on thursday Jan. 20 2022.

12 Scientific production

12.1 Major publications

- [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* 33.17 (2021), pp. 1–25. URL: <https://hal.inria.fr/hal-03508925>.

12.2 Publications of the year

International journals

- [2] 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* 33.17 (2021), pp. 1–25. URL: <https://hal.inria.fr/hal-03508925>.
- [3] M. Chardet, H. Coullon and S. Robillard. ‘Toward Safe and Efficient Reconfiguration with Concerto’. In: *Science of Computer Programming* 203 (1st Mar. 2021), pp. 1–31. DOI: [10.1016/j.scico.2020.102582](https://doi.org/10.1016/j.scico.2020.102582). URL: <https://hal.inria.fr/hal-03103714>.
- [4] A. Da Silva Veith, M. Dias de Assuncao and L. Lefèvre. ‘Latency-Aware Strategies for Deploying Data Stream Processing Applications on Large Cloud-Edge Infrastructure’. In: *IEEE transactions on cloud computing* (20th July 2021), pp. 1–12. DOI: [10.1109/TCC.2021.3097879](https://doi.org/10.1109/TCC.2021.3097879). URL: <https://hal.inria.fr/hal-03347555>.
- [5] H. Hadjur, D. Ammar and L. Lefèvre. ‘Toward an intelligent and efficient beehive: A survey of precision beekeeping systems and services’. In: *Computers and Electronics in Agriculture* 192 (Jan. 2022), pp. 1–16. DOI: [10.1016/j.compag.2021.106604](https://doi.org/10.1016/j.compag.2021.106604). URL: <https://hal.inria.fr/hal-03483914>.
- [6] M. Rzepka, P. Boryło, M. Assunção, A. Lasoń and L. Lefèvre. ‘SDN-based fog and cloud interplay for stream processing’. In: *Future Generation Computer Systems* 131 (June 2022), pp. 1–17. DOI: [10.1016/j.future.2022.01.006](https://doi.org/10.1016/j.future.2022.01.006). URL: <https://hal.inria.fr/hal-03559874>.

International peer-reviewed conferences

- [7] Y. Caniou, E. Caron, A. Kong Win Chang and Y. Robert. ‘Budget-aware Static Scheduling of Stochastic Workflows with DIET’. In: *ADVCOMP 2021 - Fifteenth International Conference on Advanced Engineering Computing and Applications in Sciences*. Barcelona, Spain, 3rd Oct. 2021, pp. 1–8. URL: <https://hal.inria.fr/hal-03332601>.
- [8] E. Caron, A. Chevalier, N. Baillon-Bachoc and A.-L. Vion. ‘Heuristic for license-aware, performant and energy efficient deployment of multiple software in Cloud architecture’. In: *ICICS 2021 - 12th International Conference on Information and Communication Systems*. Valencia, Spain, 24th May 2021. DOI: [10.1109/ICICS52457.2021.9464578](https://doi.org/10.1109/ICICS52457.2021.9464578). URL: <https://hal.archives-ouvertes.fr/hal-03572922>.
- [9] S. Delamare and L. Nussbaum. ‘Kwollect: Metrics Collection for Experiments at Scale’. In: *CNERT 2021 - Workshop on Computer and Networking Experimental Research using Testbeds*. Virtual, United States, 10th May 2021, pp. 1–6. URL: <https://hal.inria.fr/hal-03236421>.
- [10] T. Gautier and J. V. F. Lima. ‘Evaluation of two topology-aware heuristics on level-3 BLAS library for multi-GPU platforms’. In: *PAW-ATM 2021 - 4th Annual Parallel Applications Workshop, Alternatives To MPI+X*. Saint Louis, United States, 19th Nov. 2021, pp. 1–11. URL: <https://hal.inria.fr/hal-03363275>.
- [11] Z. Houmani, D. Balouek-Thomert, E. Caron and M. Parashar. ‘Enabling microservices management for Deep Learning applications across the Edge-Cloud Continuum’. In: *SBAC-PAD 2021 - IEEE 33rd International Symposium on Computer Architecture and High Performance Computing*. Belo Horizonte, Brazil: IEEE, 26th Oct. 2021, pp. 1–10. URL: <https://hal.inria.fr/hal-03409405>.
- [12] R. Pereira, A. Roussel, P. Carribault and T. Gautier. ‘Communication-Aware Task Scheduling Strategy in Hybrid MPI+OpenMP Applications’. In: *IWOMP 2021 - 17th International Workshop on OpenMP. OpenMP : Enabling Massive Node-Level Parallelism (IWOMP 2021)*. Bristol, United Kingdom, 15th Sept. 2021, pp. 1–15. URL: <https://hal-cea.archives-ouvertes.fr/cea-03284169>.
- [13] K. Rao Vaddina, L. Lefèvre and A.-C. Orgerie. ‘Experimental Workflow for Energy and Temperature Profiling on HPC Systems’. In: *ISCC 2021 - IEEE Symposium on Computers and Communications*. Athens, Greece: IEEE, Sept. 2021, pp. 1–7. URL: <https://hal.archives-ouvertes.fr/hal-03335184>.

- [14] L. Turpin, T. Gautier and J. Rouzaud-Cornabas. ‘X-Aevol: GPU implementation of an evolutionary experimentation simulator’. In: GECCO 2021 - Genetic and Evolutionary Computation Conference. Lille France, France: ACM, 7th July 2021, pp. 1–9. DOI: [10.1145/3449726.3463195](https://doi.org/10.1145/3449726.3463195). URL: <https://hal.inria.fr/hal-03290799>.

Doctoral dissertations and habilitation theses

- [15] I. Daoudi. ‘Performance Modelling and Simulation of OpenMP Applications’. Université de Bordeaux, 21st Sept. 2021. URL: <https://tel.archives-ouvertes.fr/tel-03416335>.

Reports & preprints

- [16] I. Daoudi, S. Thibault and T. Gautier. *Draft: sOMP: NUMA and cache-aware simulations for task-based applications*. RR-9400. Inria, 22nd Mar. 2021, p. 25. URL: <https://hal.inria.fr/hal-03177026>.

12.3 Cited publications

- [17] 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.
- [18] 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>.
- [19] 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.
- [20] 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.
- [21] G. Madec. *NEMO ocean engine*. Note du Pole de modélisation 27. ISSN No 1288-1619. France: Institut Pierre-Simon Laplace (IPSL), 2008.
- [22] OpenACC. *The OpenACC Application Programming Interface*. Version 1.0. Nov. 2011. URL: <http://www.openacc-standard.org>.
- [23] OpenMP Architecture Review Board. *OpenMP Application Program Interface*. Version 3.1. July 2011. URL: <http://www.openmp.org>.
- [24] 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.
- [25] C. Szyperski. *Component Software - Beyond Object-Oriented Programming*. 2nd ed. Addison-Wesley / ACM Press, 2002, p. 608.
- [26] 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.