

RESEARCH CENTRE

Inria Lyon Center

IN PARTNERSHIP WITH:

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

2022

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**

Inria

Contents

Project-Team AVALON	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	3
2.1 Presentation	3
2.2 Objectives	3
3 Research program	4
3.1 Energy Application Profiling and Modeling	4
3.2 Data-intensive Application Profiling, Modeling, and Management	4
3.3 Resource-Agnostic Application Description Model	5
3.4 Application Mapping and Scheduling	5
3.4.1 Application Mapping and Software Deployment	5
3.4.2 Non-Deterministic Workflow Scheduling	6
3.4.3 Software Asset Management	6
4 Application domains	6
4.1 Overview	6
4.2 Climatology	6
4.3 Astrophysics	7
4.4 Bioinformatics	7
5 Social and environmental responsibility	7
5.1 Footprint of research activities	7
5.2 Impact of research results	7
6 Highlights of the year	8
6.1 Awards	8
6.2 Start'up	8
7 New software and platforms	8
7.1 New software	8
7.1.1 IQ Orchestra	8
7.1.2 execo	9
7.1.3 Halley	9
7.1.4 libkomp	9
7.1.5 XKBLAS	10
7.2 New platforms	10
7.2.1 Platform: Grid'5000	10
7.2.2 Platform: SLICES-FR	10
7.2.3 Platform: SLICES	11
8 New results	11
8.1 Energy Efficiency in HPC and Large Scale Distributed Systems	11
8.1.1 Energy Consumption and Energy Efficiency in a Precision Beekeeping System	11
8.1.2 Ecodesign of large scale distributed applications	11
8.1.3 Environmental assessment of projects involving AI methods	12
8.1.4 Challenging life cycle analysis of distributed ICT services	12
8.1.5 Comparing software-based power meters dedicated on CPU and GPU	12
8.1.6 Immersion cooling system and analysis	13
8.2 Modeling and Simulation of Parallel Applications and Distributed Infrastructures	13
8.2.1 SDN-based Fog and Cloud Interplay for Stream Processing	13
8.2.2 Building dynamic urgent applications on continuum computing platforms	14
8.3 Edge and Cloud Resource Management	14

8.3.1	Total Cost Modeling for VNF based on Licenses and Resources	14
8.4	HPC Applications and Runtimes	14
8.4.1	Enhancing MPI+OpenMP task based applications for heterogenous architectures with GPU support	14
9	Bilateral contracts and grants with industry	15
9.1	Bilateral grants with industry	15
10	Partnerships and cooperations	16
10.1	International initiatives	16
10.1.1	Participation in other International Programs	16
10.2	European initiatives	16
10.2.1	Horizon Europe	16
10.2.2	H2020 projects	18
10.3	National initiatives	21
10.4	Regional initiatives	21
10.4.1	Action Exploratoire Inria	21
11	Dissemination	21
11.1	Promoting scientific activities	21
11.1.1	Scientific events: organisation	21
11.1.2	Scientific events: selection	22
11.1.3	Journal	22
11.1.4	Invited talks	22
11.1.5	Scientific expertise	23
11.1.6	Research administration	23
11.2	Teaching - Supervision - Juries	23
11.2.1	Teaching	23
11.2.2	Teaching administration	24
11.2.3	Supervision	25
11.2.4	Juries	25
11.3	Popularization	26
11.3.1	Articles and contents	26
11.3.2	Education	26
11.3.3	Interventions	26
12	Scientific production	26
12.1	Major publications	26
12.2	Publications of the year	26
12.3	Other	28
12.4	Cited publications	28

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
- 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
- B4.5. – Energy consumption
 - B4.5.1. – Green computing
- B6.1.1. – Software engineering
- 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 LYON I, Associate Professor]
- Eddy Caron [ENS DE LYON, Associate Professor, HDR]
- Olivier Glück [UNIV LYON I, Associate Professor]
- Elise Jeanneau [UNIV LYON I, Associate Professor]
- Etienne Maufret [ENS Lyon, ATER, until Aug 2022]

Post-Doctoral Fellows

- Thierry Arrabal [CNRS, until Sep 2022]
- Yasmina Bouizem [INRIA, from Sep 2022]
- Jerry Lacmou Zeutouo [INRIA, from Sep 2022]

PhD Students

- Maxime Agusti [OVHCloud, CIFRE]
- Adrien Berthelot [OCTO TECHNOLOGY, CIFRE]
- Ghoshana Bista [ORANGE LABS, CIFRE]
- Hugo Hadjur [AIVANCITY]
- Mathilde Jay [UGA]
- Simon Lambert [CIRIL GROUP, CIFRE, from Apr 2022]
- Lucien Ndjie Ngale [UPJV]
- Vladimir Ostapenco [INRIA]
- Romain Pereira [CEA]
- Pierre-Etienne Polet [THALES, CIFRE]

Technical Staff

- Brice-Edine Bellon [INRIA, Engineer, from Feb 2022]
- Arthur Chevalier [ENS DE LYON, Engineer]
- Simon Delamare [CNRS, Engineer]
- Matthieu Imbert [INRIA, Engineer]
- Pierre Jacquot [INRIA, Engineer]
- Jean Christophe Mignot [CNRS, Engineer]
- Dominique Ponsard [CNRS, Engineer]

Interns and Apprentices

- Ioan-Tudor Cebere [ENS DE LYON, from Feb 2022]

Administrative Assistants

- Evelyne Blesle [INRIA, until Nov 2022]
- Chrystelle Mouton [INRIA, from Oct 2022]

External Collaborator

- Doreid Ammar [AIVANCITY]

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 and SLIES 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 [20] 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 [23], phase detection for specific HPC applications [27], *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 [22] exposes the number of nodes but hides the complexity of network topology behind a set of collective operations; OpenMP [26] simplifies the management of threads on top of a shared memory machine while OpenACC [25] 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 [21]. 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 [28] 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 [24] for ocean modelization),

code-coupling applications (*e.g.*, the OASIS coupler [29] 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.

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 SCOOP team in SKAO's SAFe framework that deals with platforms related issues such as application benchmarking and profiling, hardware-software co-design.

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. AVALON is in particular collaborating 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.

Recently, 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 with University of Sherbrooke, in November 2022, a workshop from the Entrepreneurs Jacques Cartier on the topic of "**Rebound effects in ICT. How to detect them? How to measure them? How to avoid them?**".

6 Highlights of the year

6.1 Awards

- Laurent Lefevre (co General Chair) and Eddy Caron (Local chair) have co-organized the IPDPS 2022 virtual conference from Lyon 2022 : 35th IEEE International Parallel Distributed Processing Symposium, Lyon, France (May 30 - June 3, 2022). Most of the Avalon PhD candidates (Adrien Berthelot, Ghoshana Bista, Hugo Hdjur, Mathilde Jay, Simon Lambert, Etienne Maufret, Lucien Ndjie Ngale, Vladimir Ostapenco, Pierre-Etienne Polet) were also involved as IPDPS volunteers in order to help organizing and synchronizing the multiple parallel tracks and workshops.

6.2 Start'up

- At the end of the year, Eddy Caron started a new Start up called Qirinus. This startup used the results of a joint works from Eddy Caron (Inria. Avalon), Arthur Chevalier (Inria. Avalon), Arnaud Lefray (Inria. Avalon).

7 New software and platforms

7.1 New software

7.1.1 IQ 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.)

Release Contributions: - Microservices Architecture - Multi-cloud support - Terraform export - Update of all old software embedded - Bugs fix

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

URL: <http://www.qirinus.com>

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

Contact: Eddy Caron

Participants: Eddy Caron, Arthur Chevalier, Arnaud Lefray

Partner: ENS Lyon

7.1.2 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

7.1.3 Halley

Name: Halley

Keywords: Software Components, HPC

Scientific Description: Halley is an implementation of the COMET component model that enable to efficiently compose independent parallel code using task graph for multi-core shared-memory machines.

Functional Description: Halley is an implementation of the COMET component model that enable to efficiently compose independent parallel code using task graph for multi-core shared-memory machines.

News of the Year: First operational version.

Publications: [tel-01663718](#), [hal-01518730](#), [hal-01566288](#), [hal-01901806](#)

Contact: Christian Perez

Participants: Jérôme Richard, Christian Perez

7.1.4 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

7.1.5 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

7.2 New platforms

7.2.1 Platform: Grid'5000

Participants: Simon Delamare, Pierre Jacquot, Laurent Lefèvre, 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.

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

7.2.2 Platform: SLICES-FR

Participants: Laurent Lefèvre, Simon Delamare, Christian Perez.

FUNCTIONAL DESCRIPTION The SLICES-FR infrastructure (IR ministère), that was known as SILECS, 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/>

7.2.3 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. SLICES-FR is the The French node of SLICES.

- 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 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 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 2022, we published a survey dedicated on challenges in precision beekeeping systems [3].

8.1.2 Ecodesign of large scale distributed applications

Participants: Laurent Lefèvre.

Creating energy aware with limited environmental impacts applications 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 new version of brochure available for software developers. This brochure [15] has been downloaded several thousands of times since the publication of the first version. Some joint papers have been published in 2022 [2] [11]

8.1.3 Environmental assessment of projects involving AI methods

Participants: Laurent Lefèvre.

With colleagues from the GDS EcoInfo group (Anne-Laure Ligozat, Dernis Trystram) we explore criteria for assessing the environmental impacts of responses to calls for projects involving Artificial Intelligence (AI) methods. When proposing these criteria, we take into account, in addition to the general impacts of digital services, the specificities of the AI field and in particular of machine learning: impacts of the learning and inference phases and data collection. This resulted in two available brochures [13] and [12].

8.1.4 Challenging life cycle analysis of distributed ICT services

Participants: Adrien Berthelot, Eddy Caron, Laurent Lefèvre.

The omnipresence of ICT technology in a society trying to reorient itself towards greater sustainability raises the question of the sustainability of these technologies. If many studies exist on the ecological cost of individual equipment or on the consumption of systems and software, many crucial issues remain too little explored. On the one hand, the costs are too often limited to the electrical consumption of these technologies, invisibilizing a wide range of significant, if not prominent, environmental damages. On the other hand, current approaches are too focused to assess the real footprint of human activities related to digital services, whether by omitting the footprint related to equipment manufacturing or infrastructure costs.

Our work for greener ICT, within the framework of Adrien Berthelot's PhD is mainly based on 2 methodological changes. First, an increased interest in the life cycle assessment methodology, which allows taking into account more exhaustively the environmental impacts. Second, a shift to the ICT service scale to study how a set of hardware and software elements provides a service and at what price. In the absence of a dominant standard, we seek to contribute to a better methodology for measuring the impacts of ICT services. A methodology that could be scientifically reliable, but also effectively help decision-making regarding environmental choices. In 2022, we published an article summarizing the issues and limitations of environmental assessment applied to ICT services.[10].

8.1.5 Comparing software-based power meters dedicated on CPU and GPU

Participants: Mathilde Jay, Laurent Lefèvre, Vladimir Ostapenco.

The global energy demand for digital activities is constantly growing. Computing nodes and cloud services are at the heart of these activities. Understanding their energy consumption is an important step towards reducing it. On one hand, physical power meters are very accurate in measuring energy but they

are expensive, difficult to deploy on a large scale, and are not able to provide measurements at the service level. On the other hand, power models and vendor-specific internal interfaces are already available or can be implemented on existing systems. Plenty of tools, called software-based power meters, have been developed around the concepts of power models and internal interfaces, in order to report the power consumption at levels ranging from the whole computing node to applications and services. However, we have found that it can be difficult to choose the right tool for a specific need. In this work, we qualitatively and experimentally compare several software-based power meters able to deal with CPU or GPU-based infrastructures. For this purpose, we evaluate them against high-precision physical power meters while executing various intensive workloads. We extend this empirical study to highlight the strengths and limitations of each software-based power meter. This research is a joint work with Denis Trystram (LIG Laboratory) and Anne-Cécile Orgerie (IRISA, Rennes). Two posters have been disseminated during the Compas 2022 conference [16] [17]

8.1.6 Immersion cooling system and analysis

Participants: Thierry Arrabal, Lucas Betencourt, Eddy Caron, Laurent Lefèvre.

With the CBP (Centre de Calcul Blaise Pascal) and the TotalLinux company we started a collaboration together to launch a study based on immersion cooling prototypes dedicated to the next data center generation. For these data centers, the exponential performance evolution of IT equipment leads to an exponential increase in energy consumption, particularly for cooling. Immersion cooling using mineral oil as a heat transfer fluid appears to be a solution in the future to both meet the growing cooling needs and limit the associated energy consumption. In [5], we compared the cooling-efficiency between air cooling and immersion cooling methods by using 8 identical HPC servers. Short and long term stresses with and without overclocking were run. A statistical analysis using the Mann-Whitney method was conducted to determine whether or not, the temperature differences observed between the servers cooled by immersion and the ones cooled by air were significant. For the immersed servers, we obtained average processor temperatures down to 15% lower and RAM strip temperatures down to 32% lower. A 1.5 times higher heat transfer coefficient was obtained in oil compared to air (with natural oil convection compared to forced air convection).

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 [4] 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) [4].

8.2.2 Building dynamic urgent applications on continuum computing platforms

Participants: Eddy Caron, Laurent Lefevre.

Advanced cyberinfrastructure aims at making the use of streaming data a common practice in the scientific community. They offer an ecosystem that links data, compute, network, and users to deliver knowledge obtained from multiple data sources using large-scale computational models. However, integrating this heterogeneous data with time-sensitive systems is difficult due to a lack of programming abstractions that can allow data-driven reactive behaviors throughout the edge-to-cloud/HPC computing continuum. Here we present a methodology for incorporating contextual information into the application logic while taking into consideration the heterogeneity of the underlying platform and the unpredictability of the data. A fire science scenario that includes sensors at the network's edge for smoke detection and computational models launched in the cloud for wildfire simulation and air quality assessment serves as the inspiration for this method. This topic is joint work with University of Utah through the collaboration with Daniel Balouek and Manish Parashar and resulted in a publication in the Urgent HPC workshop in 2022 [6]

8.3 Edge and Cloud Resource Management

03660857

8.3.1 Total Cost Modeling for VNF based on Licenses and Resources

Participants: Ghoshana Bista, Eddy Caron.

Moving to NFV (Network Function Virtualization) and SDN (Software Defined Network), Telcos face four key cloud architecture challenges: interoperability, automation, reliability, and adaptability. All these challenges encompass optimization resources; whether it is to increase the utilization of hardware resources (virtualization) or to deliver shared computing resources and functions in real-time (cloudification). Softwarization of networks is a consequence of telco cloudification. Virtual Network Function (VNF) is Software and is de facto protected by IPR (Intellectual Property Right), ensured by a license that circumscribes usage rights at a given negotiated cost. Until now, only few works have dealt with the economic dimension of softwarisation. Currently, the telco industry struggles to converge and standardize licensing and cost models. At risk: the network cloudification benefits could be swept away by poor management of resources (Hardware and Software). In [8] we have focused on the NVF software cost modeling and then in [7] we introduced a preliminary model for optimizing the total cost of a VNF, based on the Resource Cost (RC) and License Cost (LC). This analysis is inspired by measurement and licensing practices previously studied in Avalon (in collaboration with Orange) and commonly observed in the Telcos industries, i.e consumption and capacity.

The study has tried to suggest several models that are relevant to the various scenarios. We compare the traditional ways of estimating total cost and our models (capacity and consumption). Results show that our model is far better than the traditional one. We also present different kinds of possible scenarios such as VNF instances and users which have a huge range of requirements to be fulfilled. We proposed the flavours methods (Simultaneous Active Users or based on Bandwidth requirement). We introduced potential metrics and we introduced a novel model.

8.4 HPC Applications and Runtimes

8.4.1 Enhancing MPI+OpenMP task based applications for heterogenous architectures with GPU support

Participants: Romain Peirera, Thierry Gautier.

Heterogeneous supercomputers are widespread over HPC systems and programming efficient applications on these architectures is a challenge. Task-based programming models are a promising way to tackle this challenge. Since OpenMP 4.0 and 4.5, the target directives enable to offload pieces of code to GPUs and to express it as tasks with dependencies. Therefore, heterogeneous machines can be programmed using MPI+OpenMP(task+target) to exhibit a very high level of concurrent asynchronous operations for which data transfers, kernel executions, communications and CPU computations can be overlapped. Hence, it is possible to suspend tasks performing these asynchronous operations on the CPUs and to overlap their completion with another task execution. Suspended tasks can resume once the associated asynchronous event is completed in an opportunistic way at every scheduling point. We have integrated [9] this feature into the MPC framework and validated it on a AXPY microbenchmark and evaluated on a MPI+OpenMP(tasks) implementation of the LULESH proxy applications. The results show that we are able to improve asynchronism and the overall HPC performance, allowing applications to benefit from asynchronous execution on heterogeneous machines.

9 Bilateral contracts and grants with industry

[ALL] A Vérifier: mettre à jour le commentaire quand c'est fait

9.1 Bilateral grants with industry

Participants: Eddy Caron, Thierry Gautier, Laurent Lefevre.

CEA We have a collaboration with CEA / DAM-Île de France. This collaboration is based on the co-advicing of a CEA PhD. The research of the PhD student (Romain Pereira) 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.

Octo technology We have a collaboration with Octo Technology (Part of Accenture). This collaboration is sealed through a CIFRE PhD grant. The research of the PhD student (Adrien Berthelot) focuses on accelerated and driven evaluation of the environmental impacts of an Information System with the full set of digital services

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).

OVHCloud We have a collaboration with OVHCloud through the FrugalCloud collaboration (Inria Défi/challenge) between Inria and OVHCloud company. This collaboration explores 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 Ostapenco.

SynAApps We have a collaboration with SynAApps (part of Cyril Group). This collaboration is sealed through a CIFRE PhD grant. The research of the PhD student (Simon Lambert) focuses on forecast and dynamic resource provisioning on a virtualization infrastructure).

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

[ALL] A Vérifier: mettre à jour le commentaire quand c'est fait

10.1 International initiatives

10.1.1 Participation in other International Programs

JLESC

Participants: Christian Perez, Thierry Gautier, Jerry Lacmou, Romain Pereira.

Title: Joint Laboratory for Extreme Scale Computing

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

Duration: 2014-2023

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.

SKA

Participants: Christian Perez, Laurent Lefevre, Thierry Gautier.

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.

10.2 European initiatives

10.2.1 Horizon Europe

SLICES-PP

Participants: Christian Perez, Laurent Lefevre.

Title: Scientific Large-scale Infrastructure for Computing/Communication Experimental Studies - Preparatory Phase

Partners:

- Institut National de Recherche en Informatique et Automatique (INRIA), France
- Sorbonne Université (SU), France
- Universiteit van Amsterdam (UvA), Netherlands
- University of Thessaly (UTH), Greece
- Consiglio Nazionale delle Ricerche (CNR), Italy
- Instytut Chemii Bioorganicznej Polskiej Nauk (PSNC), Poland
- Mandat International (MI), Switzerland
- IoT Lab (IoTLAB), Switzerland
- Universidad Carlos III de Madrid (UC3M), Spain
- Interuniversitair Micro-Electronica Centrum (IMEC), Belgium
- UCLan Cyprus (UCLAN), Cyprus
- EURECOM, France
- Számítástechnikai és Automatizálási Kutatóintézet (SZTAKI), Hungary
- Consorzio Interuniversitario Nazionale per l'Informatica (CINI), Italy
- Consorzio Nazionale Interuniversitario per le Telecomunicazioni (CNIT), Italy
- Universite du Luxembourg (Uni.Lu), Luxembourg
- Technical Universitaet Muenchen (TUM), Germany
- Euskal Herriko Unibertsitatea (EHU), Spain
- Kungliga Tekniska Hoegskolan (KTH), Sweden
- Oulun Yliopisto (UOULU), Finland
- EBOS Technologies Ltd (EBOS), Cyprus
- Simula Research Laboratory AS (SIMULA), Norway
- Centre National de la Recherche Scientifique (CNRS), France
- Institut Mines-Télécom (IMT), France
- Université de Geneve (UniGe), Switzerland

Duration: From September 1, 2022 to Decembre 31, 2025

Summary: The digital infrastructures research community continues to face numerous new challenges towards the design of the Next Generation Internet. This is an extremely complex ecosystem encompassing communication, networking, data-management and data-intelligence issues, supported by established and emerging technologies such as IoT, 5/6G, cloud-to-edge computing. Coupled with the enormous amount of data generated and exchanged over the network, this calls for incremental as well as radically new design paradigms. Experimentally-driven research is becoming worldwide a de-facto standard, which has to be supported by large-scale research infrastructures to make results trusted, repeatable and accessible to the research communities. SLICES-RI (Research Infrastructure), which was recently included in the 2021 ESFRI roadmap, aims to answer these problems by building a large infrastructure needed for the experimental research on various aspects of distributed computing, networking, IoT and 5/6G networks. It will provide the resources needed to continuously design, experiment, operate and automate the full lifecycle management of digital infrastructures, data, applications, and services. Based on the two preceding projects within SLICES-RI, SLICES-DS (Design Study) and SLICES-SC (Starting Community), the SLICES-PP (Preparatory Phase) project will validate the requirements to engage into the implementation phase of the RI lifecycle. It will set the policies and decision processes for the governance

of SLICES-RI: i.e. the legal and financial frameworks, the business model, the required human resource capacities and training programme. It will also settle the final technical architecture design for implementation. It will engage member states and stakeholders to secure commitment and funding needed for the platform to operate. It will position SLICES as an impactful instrument to support European advanced research, industrial competitiveness and societal impact in the digital era.

10.2.2 H2020 projects

SLICES-DS

Participants: Christian Perez, Laurent Lefevre.

Title: Scientific Large-scale Infrastructure for Computing/Communication Experimental Studies - Design Study

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- ECOLE NORMALE SUPERIEURE DE LYON (ENS DE LYON), France
- INTERUNIVERSITAIR MICRO-ELECTRONICA CENTRUM (IMEC), Belgium
- UCLAN CYPRUS LIMITED (UCLan Cyprus), Cyprus
- INSTYTUT CHEMII BIOORGANICZNEJ POLSKIEJ AKADEMII NAUK, Poland
- INSTITUT MINES-TELECOM, France
- MANDAT INTERNATIONAL ALIAS FONDATION POUR LA COOPERATION INTERNATIONALE (MI), Switzerland
- CONSIGLIO NAZIONALE DELLE RICERCHE (CNR), Italy
- PANEPISTIMIO THESSALIAS (UNIVERSITY OF THESSALY - UTH), Greece
- CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE CNRS (CNRS), France
- UNIVERSIDAD CARLOS III DE MADRID (UC3M), Spain
- UNIVERSITEIT VAN AMSTERDAM (UvA), Netherlands
- SORBONNE UNIVERSITE, France

Duration: From September 1, 2020 to August 31, 2022

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.

PRACE-6IP

Participants: Christian Perez, Thierry Gautier.

Title: PRACE 6th Implementation Phase Project

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- CENTRUM SPOLOCNYCH CINNOSTI SLOVENSKEJ AKADEMIE VIED (CENTRE OF OPERATIONS OF THE SLOVAK ACADEMY OF SCIENCES), Slovakia
- GRAND EQUIPEMENT NATIONAL DE CALCUL INTENSIF (GENCI), France
- UNIVERSIDADE DO MINHO (UMINHO), Portugal
- LINKOPINGS UNIVERSITET (LIU), Sweden
- VSB - TECHNICAL UNIVERSITY OF OSTRAVA (VSB - TU Ostrava), Czechia
- MACHBA - INTERUNIVERSITY COMPUTATION CENTER (IUCC), Israel
- TECHNISCHE UNIVERSITAET WIEN (TU WIEN), Austria
- Gauss Centre for Supercomputing (GCS) e.V. (GCS), Germany
- FUNDACION PUBLICA GALLEGA CENTRO TECNOLOGICO DE SUPERCOMPUTACION DE GALICIA (CESGA), Spain
- UNIVERSITEIT ANTWERPEN (UANTWERPEN), Belgium
- NATIONAL UNIVERSITY OF IRELAND GALWAY (NUI GALWAY), Ireland
- AKADEMIA GORNICZO-HUTNICZA IM. STANISLAWA STASZICA W KRAKOWIE (AGH / AGH-UST), Poland
- KUNGLIGA TEKNISKA HOEGSKOLAN (KTH), Sweden
- FORSCHUNGSZENTRUM JULICH GMBH (FZJ), Germany
- EUDAT OY (EUDAT), Finland
- KORMANYZATI INFORMATIKAI FEJLESZTESI UGYNOKSEG (GOVERNMENTAL INFORMATION TECHNOLOGY DEVELOPMENT AGENCY), Hungary
- COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES (CEA), France
- NATIONAL INFRASTRUCTURES FOR RESEARCH AND TECHNOLOGY (GRNET S.A.), Greece
- GEANT VERENIGING (GEANT VERENIGING), Netherlands
- UNIVERSIDADE DE EVORA (UNIVERSIDADE DE EVORA), Portugal
- KOBENHAVNS UNIVERSITET (UCPH), Denmark
- UPPSALA UNIVERSITET (UU), Sweden
- INSTYTUT CHEMII BIOORGANICZNEJ POLSKIEJ AKADEMII NAUK, Poland
- ISTANBUL TEKNIK UNIVERSITESI (ITU), Türkiye
- BAYERISCHE AKADEMIE DER WISSENSCHAFTEN (BADW), Germany
- SURF BV, Netherlands

- ASSOCIACAO DO INSTITUTO SUPERIOR TECNICO PARA A INVESTIGACAO E DESENVOLVIMENTO (IST ID), Portugal
- PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE AISBL (PRACE), Belgium
- UMEA UNIVERSITET, Sweden
- UNIVERSIDADE DE COIMBRA (UNIVERSIDADE DE COIMBRA), Portugal
- UNITED KINGDOM RESEARCH AND INNOVATION (UKRI), United Kingdom
- UNIVERSITE DU LUXEMBOURG (uni.lu), Luxembourg
- EIDGENOESSISCHE TECHNISCHE HOCHSCHULE ZUERICH (ETH Zürich), Switzerland
- SYDDANSK UNIVERSITET (SDU), Denmark
- "ASSOCIATION " "NATIONAL CENTRE FOR SUPERCOMPUTING APPLICATIONS" (NCSA), Bulgaria
- BILKENT UNIVERSITESI VAKIF (BILKENTUNIVERSITY BILIM KENTI), Türkiye
- UNIVERSITETET I OSLO (UNIVERSITY OF OSLO), Norway
- DANMARKS TEKNISKE UNIVERSITET (DTU), Denmark
- UNIVERSIDADE DO PORTO (U.PORTO), Portugal
- SIGMA2 AS (SIGMA2), Norway
- UNIVERSITY OF STUTTGART (USTUTT), Germany
- MAX-PLANCK-GESELLSCHAFT ZUR FORDERUNG DER WISSENSCHAFTEN EV (MPG), Germany
- UNIVERSITAET INNSBRUCK (UIBK), Austria
- CINECA CONSORZIO INTERUNIVERSITARIO (CINECA), Italy
- CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE CNRS (CNRS), France
- CENTRE INFORMATIQUE NATIONAL DE L'ENSEIGNEMENT SUPERIEUR (CINES), France
- POLITECHNIKA WROCLAWSKA (PWR), Poland
- POLITECHNIKA GDANSKA (GDANSK TECH), Poland
- NORGES TEKNISK-NATURVITENSKAPELIGE UNIVERSITET NTNU (NTNU), Norway
- THE CYPRUS INSTITUTE (THE CYPRUS INSTITUTE), Cyprus
- THE UNIVERSITY OF EDINBURGH (UEDIN), United Kingdom
- UNIVERZA V LJUBLJANI (UL), Slovenia
- BARCELONA SUPERCOMPUTING CENTER CENTRO NACIONAL DE SUPERCOMPUTACION (BSC CNS), Spain
- CSC-TIETEEN TIETOTEKNIKAN KESKUS OY (CSC-IT CENTER FOR SCIENCE LTD), Finland

Duration: From May 1, 2019 to December 31, 2022

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.

10.3 National initiatives

Inria Large Scale Initiative

FrugalCloud: 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 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 General Chair of IPDPS 2022: 35th IEEE International Parallel & Distributed Processing Symposium, Lyon, France, May 30 - June 3, 2022.
 - Co Poster Chair of CCGrid 2022 conference: The 22st IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing, Taormina, Italy, May 16-19, 2022.
 - Co-organizer of the 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 29, 2022.

Member of the organizing committees

- Eddy Caron was Local Organizer of IPDPS 2022 (Lyon, 1-4 June 2022).
- Christian Perez was member of the Organizing Committee of the French Journées Calcul Données (Dijon, 10-12 Oct 2022).

11.1.2 Scientific events: selection

Chair of conference program committees

- Christian Perez was Project Posters Chair in ISC High Performance 2022.

Member of the conference program committees

- Yves Caniou was a program committee member for the *International Conference on Computational Science and its Applications* 2022.
- Eddy Caron was member of the programm committees of CLOSER 2022 and SETCAC'22.
- Thierry Gautier was member of the programm committees of IPDPS 2022.

Reviewer

- Eddy Caron was reviewer for CLOSER 2022, CNRIA 2022

11.1.3 Journal

- Eddy Caron was reviewer for Cluster Computing, Journal of Parallel and Distributed Computing and Transactions on Parallel and Distributed Systems.

Reviewer - reviewing activities

- Eddy Caron was reviewer for TPDS and Cluster Computing.

11.1.4 Invited talks

Laurent Lefevre was invited for :

- “Numérique responsable & durable, et si nous avions tout faux ?”, The Green IT Day, Round table, Montpellier, October 6, 2022
- “Améliorer l’efficacité énergétique et réduire les impacts environnementaux des grands systèmes numériques”, Round table, GDR RSD Days, April 27, 2022
- “Tour d’horizon des impacts environnementaux et humains du numérique : information et formation !”, Emannuelle Frenoux et Laurent Lefevre, Eidos64 : 14 ème édition du forum des des pratiques numériques pour l’éducation, January 19, 2022.

Élise Jeanneau was invited for :

- “Multi agent federated learning in autonomous data system“, GDR Federated Learning Day, 16 June 2022.
- “SkyData, a new paradigm for autonomous data management“, Federation d’Informatique de Lyon (FIL) seminar, 24 November 2022.
- C. Perez is co-leader of the pole Distributed Systems of the French GDR RSD ("Réseaux et Systèmes Distribués") since 2015.

11.1.5 Scientific expertise

- Christian Perez evaluated 4 projects for the French Direction générale de la Recherche et de l’Innovation.

11.1.6 Research administration

- Eddy Caron is Deputy Director in charge of research transfert since 2017 for the LIP. He is co-leader of the Distributed system and HPC team of the FIL (Fédération Informatique de Lyon).
- 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 SLICES-FR testbed. He is a member of the Inria Lyon Strategic Orientation Committee. He is in charge of organizing scientific collaborations between Inria and SKA France. He was a member of the jury for recruiting CRCN candidates in Inria Lyon center.

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

- Licence: Eddy Caron, Programmation, 48h, L3, ENS de Lyon. France.
- Agreg Info (FEADÉP): Eddy Caron, Operating System and Network, 15h, Agreg, ENS de Lyon. France.
- Agreg Info (FEADÉP): Eddy Caron, 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 réursive, 36h, niveau L1, Université Claude Bernard Lyon 1, France.
- Licence: Yves Caniou, Programmation Concurrente, 49h and Responsable 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, 12h, niveau M1, Université Claude Bernard Lyon 1, France.

- Master: Yves Caniou, Responsable of alternance students, 7h, niveau M1, Université Claude Bernard Lyon 1, France.
- Master: Yves Caniou, Unix, 15h, niveau M1, IGA Casablanca, Maroc.
- Master: Yves Caniou, Sécurité, 30h and Responsable 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: Laurent Lefèvre, Parallélisme, 12h, niveau M1, Université Lyon 1, France.
- CAPES Informatique : Laurent Lefèvre, Numérique responsable, 3h, Université Lyon1, France
- Agreg Info : Laurent Lefèvre, Impacts environnementaux du numérique 3h, Agreg, ENS de Lyon. France.
- Master : Laurent Lefèvre, Impacts environnementaux du numérique 3h, Master BioInfo, Université Lyon1. France.
- Licence: Laurent Lefèvre, TP Programmation Concurrente, 10h, niveau L3, Université Lyon1, France
- Master: Thierry Gautier, Introduction to HPC, 20h, niveau M2, INSA Lyon, France.
- Licence: Olivier Glück, Introduction Réseaux et Web, 54h, 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 SRS (Systèmes, Réseaux et Infrastructures Virtuelles) located at IGA Casablanca, 20h, niveau M2, IGA Casablanca, Maroc.
- Master: Olivier Glück, Administration systèmes et réseaux, 30h, niveau M2, Université Lyon 1, France.
- Master: Olivier Glück, Administration systèmes et réseaux, 24h, niveau M2, IGA Casablanca, Maroc.
- Licence : Frédéric Suter, Programmation Concurrente, 32.33, L3, Université Claude Bernard Lyon 1, France
- Licence: Elise Jeanneau, Introduction Réseaux et Web, 24h, niveau L1, Université Lyon 1, France.
- Licence: Elise Jeanneau, Réseaux, 27h, niveau L3, Université Lyon 1, France.
- Master: Elise Jeanneau, Algorithmes distribués, 42h, 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, 14h, niveau M2, ENS de Lyon, France.

11.2.2 Teaching administration

- Eddy Caron is the director of the PLR (Projet Long Recherche) for the ENS student (fourth year).

11.2.3 Supervision

- Licence:
 - Yves Caniou, Programmation Concurrente, Responsable of UE, niveau L3, Université Claude Bernard Lyon 1, France.
- Master
 - Yves Caniou, Projet Orientation Master, niveau M1, Université Claude Bernard Lyon 1, France.
 - Yves Caniou, Responsable of alternance students, 7h, niveau M1, Université Claude Bernard Lyon 1, France.
 - Yves Caniou, Sécurité, Responsable of UE, niveau M2, Université Claude Bernard Lyon 1, France.
- Phd in progress:
 - Maxime Agusti. *Observation de plate-formes de co-localisation baremetal, modèles de réduction énergétique et proposition de catalogues*, Feb 2022, FrugalCloud Inria-OVHCloud collaboration, Eddy Caron (co-dir. ENS de Lyon. Inria. Avalon), Benjamin Fichel (co-dir. OVHcloud), Laurent Lefevre (dir. Inria. Avalon) et Anne-Cécile Orgerie (co-dir. Inria. Myriads).
 - Adrien Berthelot. *Évaluation accélérée et assistée des impacts environnementaux d'un Système d'Information avec l'ensemble de ses services numériques*, Jan 2022, Eddy Caron (ENS de Lyon. Inria. Avalon), Christian Fauré (Octo Technology) and Laurent Lefevre (Inria. Avalon).
 - Hugo Hadjur. *Designing sustainable autonomous connected systems with low energy consumption*, 2020, Laurent Lefevre (dir. Inria. Avalon), Doreid Ammar (co-dir. Aivancity group)
 - Simon Lambert. *Forecast and dynamic resource provisioning on a virtualization infrastructure*, 2022, Eddy Caron (dir. ENS de Lyon. Inria. Avalon), Laurent Lefevre (co-dir Inria. Avalon), Rémi Grivel (co-dir. Ciril Group).
 - Lucien Arnaud 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 (Inria. Avalon) and Yulin Zhang (CRISPI. UPJV).
 - Vladimir Ostapenco. *Modeling and design of a framework and its environmental Gantt Chart in order to manage heterogeneous energy leverages*, FrugalCloud Inria-OVHCloud collaboration, Laurent Lefevre (dir. Inria. Avalon), Anne-Cécile Orgerie (co-dir. Inria. Myriads), Benjamin Fichel (co-dir. OVHcloud).
 - Ghoshana Bista, *VNF and Software Asset Management*, Feb 2020, Eddy Caron (dir), Anne-Lucie Vion (Orange).

11.2.4 Juries

- Yves Caniou was member of the intership L3 jury, ENS de Lyon, and internship M2 jury of Université Claude Bernard Lyon 1.
- Eddy Caron was member of the CoS (*a.k.a.* Comité de Sélection) for an Assistant Professor at the ENS de Lyon and at Sorbonne University.
- Eddy Caron was member of the PhD defense committees of Safuriyawu Ahmed, Insa de Lyon. (Insa de Lyon) and Nicolas Stouls (Insa de Lyon).
- Christian Perez was member of the PhD defense committee of Sebastian Friedemann, July 2022, Université Grenoble Alpes.

11.3 Popularization

11.3.1 Articles and contents

Laurent Lefevre was interviewed for :

- " Tour d'horizon des impacts environnementaux et humains du numérique", Emmanuelle Frenoux and Laurent Lefevre, Podcast Ludomag, January 31, 2022

11.3.2 Education

- Yves Caniou co-organized the 5th Edition of Le Campus du Libre, on saturday Nov. 26 2022 at INSA Hedy Lamarr bg, Université Claude Bernard Lyon 1, La Doua.

11.3.3 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] 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>.

12.2 Publications of the year

International journals

- [2] C. Bonamy, C. Boudinet, L. Bourgès, K. Dassy, L. Lefèvre, B. Ninassi and F. Vivat. 'L'écoconception d'un service numérique : des actions pour réduire l'impact environnemental du numérique'. In: *1024 : Bulletin de la Société Informatique de France* 19 (Apr. 2022), pp. 59–68. DOI: [10.48556/SIF.1024.19.59](https://doi.org/10.48556/SIF.1024.19.59). URL: <https://hal.archives-ouvertes.fr/hal-03660118>.
- [3] 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>.
- [4] 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

- [5] T. Arrabal, L. Betencourt, E. Caron and L. Lefèvre. 'Setting up an experimental framework for analysing an immersion cooling system'. In: *IEEE SBAC-PAD 2022 : IEEE 34th International Symposium on Computer Architecture and High Performance Computing*. Bordeaux, France, 2nd Nov. 2022. URL: <https://hal.inria.fr/hal-03839381>.
- [6] D. Balouek-Thomert, L. Lefèvre, E. Caron and M. Parashar. 'Towards a methodology for building dynamic urgent applications on continuum computing platforms'. In: *CIW-IUS - International Workshop on Interactive Urgent Supercomputing Held in conjunction with SC2022: The International Conference on High Performance Computing, Networking, Storage and Analysis*. Dallas, United States, 14th Nov. 2022. URL: <https://hal.inria.fr/hal-03892286>.

- [7] G. Bista, E. Caron and A.-L. Vion. 'Total Cost Modeling for VNF Based on Licenses and Resources'. In: CLOSER 2022 - 12th International Conference on Cloud Computing and Services Science. Virtual, Portugal, 27th Apr. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03660857>.
- [8] G. Bista, E. Caron and A.-L. Vion. 'VNF Software Cost Modeling Based On Telecom Network'. In: The 6th International workshop on Software Defined Networks and Network Function Virtualization 2022 (SDN-NFV2022). Paris, France, 22nd Nov. 2022. URL: <https://hal.inria.fr/hal-03915710>.
- [9] M. Ferat, R. Pereira, A. Roussel, P. Carribault, L.-A. Steffemel and T. Gautier. 'Enhancing MPI+OpenMP task based applications for heterogenous architectures with GPU support'. In: *OpenMP in a Modern World: From Multi-device Support to Meta Programming*. IWOMP 2022 - 18th International Workshop on OpenMP. Vol. LNCS-13527. Chattanooga, TH, United States: Springer International Publishing, 27th Sept. 2022, pp. 3–16. DOI: [10.1007/978-3-031-15922-0_1](https://doi.org/10.1007/978-3-031-15922-0_1). URL: <https://hal-cea.archives-ouvertes.fr/cea-03749364>.

National peer-reviewed Conferences

- [10] A. Berthelot. 'L'analyse de cycle de vie de services numériques sur systèmes distribués : définitions et enjeux'. In: COMPAS 2022: Conférence francophone d'informatique en Parallélisme Architecture et Système (Compas). Amiens, France, 5th July 2022. URL: <https://hal.inria.fr/hal-03941852>.

Conferences without proceedings

- [11] C. Bonamy, L. Bourgès and L. Lefèvre. 'Impact des langages de programmation sur la performance énergétique des applications de calcul scientifique : un challenge ?' In: JRES 2022 - Journées Réseau de l'Enseignement et de la Recherche. Marseille, France, 17th May 2022. URL: <https://hal.archives-ouvertes.fr/hal-03607468>.

Reports & preprints

- [12] L. Lefèvre, A.-L. Ligozat, D. Trystram, S. Bouveret, A. Bugeau, J. Combaz, E. Frenoux, G. Guennebaud, J. Lefèvre and J.-P. Nicolai. *Proposition de document de cadrage Évaluation environnementale de projets impliquant des méthodes d'IA*. EcoInfo, 2022, pp. 1–8. URL: <https://hal.science/hal-03853135>.
- [13] L. Lefèvre, A.-L. Ligozat, D. Trystram, S. Bouveret, A. Bugeau, J. Combaz, E. Frenoux, G. Guennebaud, J. Lefèvre, J.-P. Nicolai and K. Dassas. *Environmental assessment of projects involving AI methods*. 4th Jan. 2023. URL: <https://hal.science/hal-03922093>.

Other scientific publications

- [14] C. Bonamy, C. Boudinet, L. Bourgès, K. Dassas, L. Lefèvre, B. Ninassi and F. Vivat. *Good practices in digital service ecodesign for software developers*. 7th Feb. 2023. URL: <https://hal.science/hal-03977001>.
- [15] C. Bonamy, C. Boudinet, L. Bourgès, K. Dassas, L. Lefèvre, B. Ninassi and F. Vivat. *Je code : les bonnes pratiques en éco-conception de service numérique à destination des développeurs de logiciels*. 10th May 2022. URL: <https://hal.archives-ouvertes.fr/hal-03009741>.
- [16] M. Jay, D. Trystram and L. Lefèvre. 'Comparaison expérimentale de logiciels de mesure d'énergie : côté GPU'. In: Compas 2022. Amiens, France, 5th July 2022. URL: <https://hal.inria.fr/hal-03945744>.
- [17] V. Ostapenco, L. Lefevre, A.-C. Orgerie and B. Fichel. 'Une comparaison expérimentale de wattmètres logiciels : focus sur le CPU'. In: COMPAS 2022. Amiens, France, 5th July 2022. URL: <https://hal.inria.fr/hal-03946331>.

12.3 Other

Scientific popularization

- [18] A.-C. Orgerie and L. Lefèvre. *Le vrai coût énergétique du numérique*. 10th July 2022. URL: <https://hal.inria.fr/hal-03889053>.

Softwares

- [19] [SW] T. L. Group, *LinBox* version 1.7.0, 26th Sept. 2022. LIC: GNU Lesser General Public License v2.1 or later. HAL: [hal-03788347](https://hal.archives-ouvertes.fr/hal-03788347), URL: <https://hal.archives-ouvertes.fr/hal-03788347>, VCS: <https://github.com/linbox-team/linbox>, SWHID: [swh:1:dir:ab3150274e077cc4c4d18ab7bc074fab0a5c12d3;origin=https://hal.archives-ouvertes.fr/hal-03788347;visit=swh:1:snp:99e4a5b29527e7089567751cd8b75113ecd12c0b;anchor=swh:1:rel:35391e40ce3a7b02180e2dc159c095a89ed26a91;path=/](https://sw.hal.archives-ouvertes.fr/hal-03788347).

12.4 Cited publications

- [20] 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.
- [21] 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>.
- [22] 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.
- [23] 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.
- [24] G. Madec. *NEMO ocean engine*. Note du Pole de modélisation 27. ISSN No 1288-1619. France: Institut Pierre-Simon Laplace (IPSL), 2008.
- [25] OpenACC. *The OpenACC Application Programming Interface*. Version 1.0. Nov. 2011. URL: <http://www.openacc-standard.org>.
- [26] OpenMP Architecture Review Board. *OpenMP Application Program Interface*. Version 3.1. July 2011. URL: <http://www.openmp.org>.
- [27] 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.
- [28] C. Szyperski. *Component Software - Beyond Object-Oriented Programming*. 2nd ed. Addison-Wesley / ACM Press, 2002, p. 608.
- [29] 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.