

RESEARCH CENTRE

**Inria Centre at Rennes
University**

IN PARTNERSHIP WITH:

**Institut national des sciences
appliquées de Rennes**

2024
**ACTIVITY
REPORT**

Team
KERDATA

Enabling the Edge-Cloud-HPC Data Continuum

Inria teams are typically groups of researchers working on the definition of a common project, and objectives, with the goal to arrive at the creation of a project-team. Such project-teams may include other partners (universities or research institutions)

**IN COLLABORATION WITH: Institut de recherche en informatique
et systèmes aléatoires (IRISA)**

DOMAIN

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

THEME

**Distributed and High Performance
Computing**

Inria

Contents

Team KERDATA	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	3
3 Research program	4
3.1 Research Axis 1: Convergence of Extreme-Scale Computing and Big Data infrastructures	5
3.2 Research Axis 2: Advanced data processing, analytics and AI in a reproducible way on the Edge-to-Cloud Continuum	6
3.3 Research Axis 3: I/O management, in situ visualization and analysis on HPC systems at extreme scales	7
4 Application domains	8
4.1 Radio astronomy	8
4.2 Climate and meteorology	9
4.3 Earth science	9
4.4 Sustainable development through precision agriculture	10
4.5 Smart cities	10
4.6 Botanical Science	10
5 Social and environmental responsibility	11
5.1 Footprint of research activities	11
5.2 Impact of research results	11
6 Highlights of the year	12
6.1 Awards	12
6.2 SC'24 Program Chair	12
6.3 HiPC 2025 Program Chair	12
6.4 ISPDC 2025 General Co-Chairs	12
7 New software, platforms, open data	12
7.1 New software	12
7.1.1 Damaris	12
7.1.2 E2Clab	13
7.1.3 StorAlloc	14
7.1.4 Fives	15
7.1.5 MOSAIC	15
7.1.6 Neomem	16
7.1.7 FLAdversary	17
7.1.8 FLDrift	17

8	New results	18
8.1	Convergence of HPC and Big Data infrastructures for supporting workflows in the Computing Continuum	18
8.1.1	Provisioning storage resources for HPC and Cloud systems	18
8.1.2	Data profiling and benchmarking of computational workflows beyond the classical Computing Continuum	18
8.2	Advanced data processing support for Artificial Intelligence across the Computing Continuum	19
8.2.1	Investigating the Gap between Simulation, Emulation and Real-World Deployments for Reproducible Federated Learning Experiments	19
8.2.2	Formalisation of workflow provenance for trustworthy and explainable AI	19
8.2.3	Efficient Resource-Constrained Federated Learning Clustering with Local Data Compression on the Edge-to-Cloud Continuum	20
8.2.4	Efficient Data-Parallel Continual Learning with Asynchronous Distributed Rehearsal Buffers	21
8.2.5	Comparative Analysis of Federated Learning: Simulations Versus Real-World Testbeds	21
8.3	Scalable I/O, in-situ Visualization and Resource Management at Large Scale	22
8.3.1	Multi-level analysis of the I/O pattern of HPC applications	22
8.3.2	Topology and affinity-aware data aggregation	23
8.3.3	Scalable asynchronous I/O and in-situ processing with Damaris	23
8.3.4	Qualitatively Analyzing Optimization Objectives in the Design of HPC Resource Manager	24
8.3.5	Allocation Strategies for Disaggregated Memory in HPC Systems	24
8.3.6	Scheduling distributed I/O resources in HPC systems	24
9	Partnerships and cooperations	25
9.1	International initiatives	25
9.1.1	Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program	25
9.2	International research visitors	26
9.2.1	Visits of international scientists	26
9.2.2	Visits to international teams	26
9.3	European initiatives	26
9.3.1	H2020 projects	26
9.3.2	Collaborations with Major European Organizations	29
9.4	National initiatives	30
10	Dissemination	32
10.1	Promoting scientific activities	32
10.1.1	Scientific events: organisation	32
10.1.2	Scientific events: selection	33
10.1.3	Journal	34
10.1.4	Invited talks	34
10.1.5	Leadership within the scientific community	35

10.1.6 Scientific expertise	35
10.1.7 Research administration	35
10.2 Teaching - Supervision - Juries	36
10.2.1 Teaching	36
10.2.2 Supervision	37
10.2.3 Juries	38
10.3 Popularization	38
10.3.1 Productions (articles, videos, podcasts, serious games, ...)	38
11 Scientific production	38
11.1 Major publications	38
11.2 Publications of the year	40
11.3 Cited publications	42

Team KERDATA

Creation of the Team: 2024 July 01

Keywords

Computer sciences and digital sciences

- A1.1.1. – Multicore, Manycore
- A1.1.4. – High performance computing
- A1.1.5. – Exascale
- A1.1.9. – Fault tolerant systems
- A1.3. – Distributed Systems
- A1.3.5. – Cloud
- A1.3.6. – Fog, Edge
- A2.6.2. – Middleware
- A3.1.2. – Data management, quering and storage
- A3.1.3. – Distributed data
- A3.1.8. – Big data (production, storage, transfer)
- A6.2.7. – High performance computing
- A6.3. – Computation-data interaction
- A7.1.1. – Distributed algorithms
- A9.2. – Machine learning
- A9.7. – AI algorithmics

Other research topics and application domains

- B3.2. – Climate and meteorology
- B3.3.1. – Earth and subsoil
- B8.2. – Connected city
- B9.5.6. – Data science
- B9.8. – Reproducibility
- B9.11.1. – Environmental risks

1 Team members, visitors, external collaborators

Research Scientists

- Gabriel Antoniu [Team leader, INRIA, Senior Researcher]
- Silvina Caino Lores [INRIA, ISFP]
- Jakob Luettgau [INRIA, Starting Research Position]
- Guillaume Pallez [INRIA, Researcher]
- François Tessier [INRIA, ISFP]

Faculty Member

- Alexandru Costan [INSA RENNES, Associate Professor]

PhD Students

- Robin Boezennec [INRIA]
- Thomas Bouvier [INRIA, until Jun 2024]
- Arthur Jaquard [INRIA, from Oct 2024]
- Theo Jolivel [INRIA, from Oct 2024]
- Julien Monniot [INRIA]
- Cedric Prigent [INRIA]
- Mathis Valli [INRIA]

Technical Staff

- Thomas Badts [INRIA, Engineer, from Apr 2024]
- Joshua Bowden [INRIA, Engineer, until Aug 2024]
- Jean Etienne Ndamlabin Mboula [INRIA, Engineer, from Aug 2024]

Interns and Apprentices

- Theo Jolivel [INRIA, Intern, from Mar 2024 until Jul 2024]
- Ugo Thay [INRIA, Intern, from Jun 2024 until Aug 2024]
- Alix Tremodeux [ENS DE LYON, Intern, from Oct 2024]

Administrative Assistant

- Laurence Dinh [INRIA]

2 Overall objectives

Note: This version of the team’s objectives correspond to the 2020–2024 period. Renewed objectives have been defined starting with 2025.

Context: the need for scalable data management. For several years now we have been witnessing a rapidly increasing number of application areas generating and processing very large volumes of data on a regular basis. Such applications, called *data-intensive*, range from traditional large-scale simulation-based scientific domains such as climate modeling, cosmology, and bioinformatics to more recent industrial applications triggered by the Big Data phenomenon: governmental and commercial data analytics, financial transaction analytics, etc. Recently, the data-intensive application spectrum further broadened by the emergence of IoT applications that need to process data coming from large numbers of distributed sensors.

Our objective. The KerData project-team is focusing on designing innovative architectures and systems for *scalable data storage and processing*. We target three types of infrastructures: *pre-Exascale high-performance supercomputers*, *cloud-based* and *edge-based* infrastructures, according to the current needs and requirements of data-intensive applications. In addition, as emphasized by the latest Strategic Research Agenda of ETP4HPC [7], new complex applications have started to emerge: they combine simulation, analytics and learning and require hybrid execution infrastructures combining supercomputers, cloud-based and edge-based systems. Our most recent research aims to address the data-related requirements (storage, processing) for such complex workflows. They are structured in three research axes summarized below.

Challenges and goals related to the HPC-Big Data convergence. Traditionally, HPC and Big Data analytics have evolved separately, using different approaches for data storage and processing as well as for leveraging their respective underlying infrastructures. The KerData team has been tackling the convergence challenge from a data storage and processing perspective, trying to provide answers to questions like: what common storage abstractions and data management techniques could fuel storage convergence, to support seamless execution of hybrid simulation/analytics workflows on potentially hybrid supercomputer/cloud infrastructures? From a broader perspective, additional challenges are posed by the question: how does the emergence of the computing continuum impact the data storage and processing infrastructure on HPC systems? The team’s activities in this area are grouped in Research Axis 1 (see 3.1).

Challenges and goals related to cloud-based and edge-based storage and processing. The growth of the Internet of Things is resulting in an explosion of data volumes at the edge of the Internet. To reduce costs incurred due to data movement and centralized cloud-based processing, cloud workflows have evolved from single-datacenter deployment towards multiple-datacenter deployments, and further from cloud deployments towards distributed, edge-based infrastructures.

This allows applications to distribute analytics while preserving low latency, high availability, and privacy. Jointly exploiting edge and cloud computing capabilities for stream-based processing leads however to multiple challenges.

In particular, understanding the dependencies between the application workflows to best leverage the underlying infrastructure is crucial for the end-to-end performance. We are currently missing models enabling this adequate mapping of distributed analytics pipelines on the Edge-to-Cloud Continuum. The community needs tools that can facilitate the modeling of this complexity and can integrate the various components involved. In particular, the need for such tools is increasing when considering AI-enabled data analytics pipelines (e.g., based on Federated Learning or Continual Learning). This is the challenge we address in Research Axis 2 (described in 3.2).

Challenges and goals related to storage and I/O for data-intensive HPC applications. Key research fields such as climate modeling, solid Earth sciences and astrophysics rely on very large-scale simulations running on post-Petascale supercomputers. Such applications exhibit requirements clearly identified by international panels of experts like IESP [59], EESI [57], ETP4HPC [58]. A jump of one order of magnitude in the size of numerical simulations is required to address some of the fundamental questions in several communities in this context. In particular, the lack of data-intensive infrastructures and methodologies to analyze the huge results of such simulations is a major limiting factor. The high-level challenge we have been addressing in Research Axis 3 (see 3.3) is to find scalable ways to store, visualize and analyze massive outputs of data during and after the simulations through asynchronous I/O and *in-situ processing*.

Approach, methodology, platforms. KerData’s global approach consists in studying, designing, implementing and evaluating distributed algorithms and software architectures for scalable data storage and I/O management for efficient, large-scale data processing. We target three main execution infrastructures: edge and cloud platforms and pre-Exascale HPC supercomputers.

The highly experimental nature of our research validation methodology should be emphasized. To validate our proposed algorithms and architectures, we build software prototypes, then validate them at large scale on real testbeds and experimental platforms.

We strongly rely on the Grid’5000 platform. Moreover, thanks to our projects and partnerships, we have access to reference software and physical infrastructures.

In the cloud area, we use the Microsoft Azure and Amazon cloud platforms, as well as the Chameleon [53] experimental cloud testbed. In the post-Petascale HPC area, we are running our experiments on systems including some top-ranked supercomputers, such as Titan, Jaguar, Kraken, Theta, Pangea and Hawk. This provides us with excellent opportunities to validate our results on advanced realistic platforms.

Collaboration strategy. Our collaboration portfolio includes international teams that are active in the areas of data management for edge, clouds and HPC systems, both in Academia and Industry. Our academic collaborating partners include Argonne National Laboratory, University of Illinois at Urbana-Champaign, Universidad Politécnica de Madrid, Barcelona Supercomputing Center. In industry, through bilateral or multilateral projects, we have been collaborating with Microsoft, IBM, Total, Huawei, ATOS/Eviden.

Moreover, the consortia of our collaborative projects include application partners in multiple application domains from the areas of climate modeling, precision agriculture, earth sciences, smart cities or botanical science. This multidisciplinary approach is an additional asset, which enables us to take into account application requirements in the early design phase of our proposed approaches to data storage and processing, and to validate those solutions with real applications and real users.

Alignment with Inria’s scientific strategy. We have engaged in collaborative projects with some of Inria’s main strategic partners: DFKI (the main German research center in artificial intelligence) through the ENGAGE Inria-DFKI project started in 2022; and with ATOS, through the ACROSS and EUPEX H2020 EuroHPC projects, started in March 2021 and January 2022, respectively. Gabriel Antoniu, Head of the KerData team, serves as a scientific lead for Inria in these three projects. The ENGAGE project is carried out in collaboration with the DataMove and HiePACS teams, while the EUPEX project also involves the TADaaM and HiePACS teams.

3 Research program

Note: This version of the team’s research program corresponds to the 2020–2024 period. A renewed scientific program has been defined starting with 2025.

The scientific landscape in the areas of High-Performance Computing and Cloud Computing has changed significantly over the last few years. Two evolutions strongly impacted this landscape.

First, while High-Performance Computing and Big Data analytics had already started their convergence movement before 2015, this phenomenon was further enforced by the increased usage of machine learning for data analytics. This led to a triple convergence in the end: HPC, Big Data *and AI* (where the term "AI" is actually mainly referring in practice to machine learning). This convergence was driven by the emergence of new, complex application workflows. Modern use cases such as autonomous vehicles, digital twins, smart buildings and precision agriculture, are contexts where such application workflows are useful. They typically combine physics-based simulations, analysis of large data volumes and machine learning.

Second, the execution of such workflows requires a hybrid infrastructure: edge devices create streams of input data, which are processed by data analytics and machine learning applications in the Cloud, and simulations on large, specialized HPC systems provide insights into and prediction of future system state. From these results, additional steps create and communicate output data across the infrastructure levels, and, for some use cases, devices or cyber-physical systems in the real world are controlled (as in the case of smart factories). Thus, such workflows need to create different requirements for every step of their execution and require a hybrid combination of interconnected underlying infrastructure subsystems: supercomputers, cloud data centers and edge-processing systems connected to sensors (emergence of the *computing continuum*).

To leverage the computing continuum, cooperation between multiple areas (HPC, Big Data analytics, AI, cyber-security, etc.) is necessary; in Europe, this motivated the creation of the TransContinuum Initiative (TCI), whose vision is summarized in [51]. We are proud to play a leading role in TCI, where Gabriel Antoniu co-leads the use case analysis working group, in charge of "Big Data" aspects. In addition, in the framework of ETP4HPC, we have contributed to the definition of a European vision on how the HPC area is being reshaped due to the emergence of the computing continuum by co-authoring the ETP4HPC agenda in 2020 [47] and in 2022 [46]. Very recently, we have also contributed to a community white paper [40] describing the challenges of creating an integrated software/hardware ecosystem for the computing continuum.

These two evolutions are the major factors that are directly impacting the definition of our scientific program for the upcoming years. In short, we maintain our three major research axes defined five years ago, while adapting them to cope with these important evolutions.

3.1 Research Axis 1: Convergence of Extreme-Scale Computing and Big Data infrastructures

This axis keeps HPC-Big Data convergence at storage infrastructure level as a major investigation area for the team, while shifting focus from storage abstractions to the convergence of the underlying storage resources (namely, HPC storage systems and cloud storage systems). In addition, we plan to focus on I/O orchestration on hybrid HPC/cloud infrastructures as part of the computing continuum.

Dynamic provisioning of hybrid storage resources. While for years high-performance computing (HPC) systems were the predominant means of meeting the requirements expressed by large-scale scientific workflows, today some components have moved away from supercomputers to cloud-type infrastructures. This migration has been mainly motivated by the cloud's ability to perform data analysis tasks efficiently.

From an I/O and storage perspective, this means having to deal with two very different worlds: the world of cloud computing, where direct access to resources is extremely limited due to a very high level of abstraction, and the world of on-premise supercomputers offering a low level approach requiring tight user control. The abstraction layer of clouds also allows storage, network and computing resources to have a certain elasticity and to be exclusively allocated.

In this context, we propose to converge these two worlds by exploring ways to provide storage resources distributed across hybrid HPC/cloud infrastructures to complex scientific workflows combining simulation and data analysis.

To do so, we continue our recently started work on scheduling algorithms dedicated to storage resources and implemented in a storage-aware scheduler developed in the team (simulator and scheduler). We also start a new research line focused on the abstraction of storage resources in order to provide a unified interface allowing to query any type of storage on an hybrid infrastructure.

I/O Orchestration over hybrid infrastructures. On hybrid infrastructures, in the same way as the amount of generated data increases, the need for persistence has stepped up. A broad variety of large-scale scientific applications and workflows in scientific domains such as materials, high energy physics or engineering have massive I/O needs.

On a HPC system, for instance, it is typically estimated that around 10% to 20% of the wall time of this class of applications is spent in I/O. In addition, in the case of workflows running on hybrid infrastructures, these I/O are extremely varied and are no longer restricted to a single system but are spread across complex architectures.

To take advantage of the capability of current systems and hope to leverage future ones, improving the I/O performance is decisive. The complexity of both the federation and the different underlying systems implies having a strong knowledge of the workloads' I/O behavior and adopting a topology-aware approach for data movement orchestration.

We focus our effort on two research lines here. We first model the I/O behavior from the application and workflow's point of view. The parameters influencing I/O performance may be as diverse as the data size, the data model (multidimensional arrays, meshes, etc.), the data layout (array of structures, structure of arrays, etc.) or the access frequency.

The impact of each characteristic on I/O performance will be evaluated with benchmarks and real applications on the different systems of an hybrid infrastructure (HPC, cloud and edge later on) and an I/O workload model will be proposed.

Then, while using this I/O characterization, we focus our effort on data aggregation, taking into account the underlying topology, which consists of selecting a subset of intermediate resources to collect data before moving it from/to the destination (a storage system or a data processing system in case of in-transit workflows, for instance). This technique has several advantages: it increases the I/O bandwidth by reading or writing larger chunks of data, it highly reduces the number of concurrent streams to the destination and it minimizes the network contention.

3.2 Research Axis 2: Advanced data processing, analytics and AI in a reproducible way on the Edge-to-Cloud Continuum

This second axis explores challenges posed by the Computing Continuum to data processing. For the short term we will continue our current work investigating the best ways to leverage the Edge-to-Cloud continuum (using *E2Clab* as an experimental platform), and we plan to extend the infrastructure scope to also include HPC subsystems (i.e., to cover the full computing continuum), in support to application workloads where machine learning will play an increasing role.

Supporting repeatable, replicable and reproducible automatic deployments across the continuum. As communities from an increasing number of scientific domains are leveraging the Computing Continuum, a desired feature of any experimental research is that its scientific claims are verifiable by others in order to build upon them. This can be achieved through repeatability, replicability, and reproducibility (3 Rs).

E2Clab is a first step towards enabling these goals and, to the best of our knowledge, it is the first platform to support the complete analysis cycle of an application on the Computing Continuum. We plan to further consolidate *E2Clab* in order to make it a promising platform

for future performance optimization of applications on the Edge-to-Cloud Continuum through reproducible experiments.

Specifically, we plan to focus on three main directions: (1) develop new, finer grained abstractions to model the components of the entire data processing pipeline across the continuum (from data production to permanent storage) and allow researchers to trade between different costs with increased accuracy; (2) enable built-in support for other large-scale experimental testbeds, besides Grid'5000, such as Vagrant and Chameleon and ultimately provide a community driven tool for large scale experimentation; and (3) develop a benchmark for processing frameworks within the Computing Continuum atop **E2Clab**.

Many exciting research questions could then be explored leveraging such an enhanced deployment and optimization tool, especially in domains like machine and deep learning: how to improve the convergence speed of distributed algorithms (i.e., gradient descent) to reach good accuracy quickly? how to appropriately partition a model based on the capability of different cloud or edge devices?

Continual learning and inference in parallel across the Computing Continuum. As neural network architectures and their training data are getting more and more complex, so are the infrastructures that are needed to execute them sufficiently fast. Hyperparameter setting and tuning, training, inference, dataset handling are operations that are all putting a growing pressure on the underlying compute infrastructure and call for novel approaches at all levels of the workflow, including the algorithmic level, the middleware and deployment level, and the resource optimization level.

Our goal is to address the following specific research questions: how can the various possible deployment options of complex AI workflows on the available underlying infrastructure impact performance metrics? how can this infrastructure be best leveraged in practice, potentially through seamless integration of supercomputers, clouds, and fog/edge systems?

We will focus on the middleware and the deployment level. Our objective is to investigate various deployment strategies for complex AI workflows (e.g., potentially combining continual training, simulations and inference, all in parallel and in real-time) on hybrid execution infrastructures (e.g., combining supercomputers and cloud/fog/edge systems).

Efficient federated learning in heterogeneous and volatile environments. The latest technological advances in hardware accelerators like the GPUs enable the execution of machine and deep learning tasks on large volumes of data in a time that has become reasonable. Embedded systems make it possible to deploy some inference tasks as close as possible to the operational context. One of the major challenges of these heterogeneous distributed systems lies in the ability to have relevant data in a given place and at a given time.

One approach is to rely on the recent privacy-preserving Federated Learning paradigm that leverages the edge devices for training. However, such solutions raise some major challenges related to system and statistical heterogeneity, energy footprint and security.

Our goal is to identify and adapt such emerging approaches resulting from the Computing Continuum in order to respond to the problems of distribution of computations and processing, particularly in the case of workflows involving AI. This exploratory topic has concrete application use-cases such as with the smart autonomous vehicles or military and civilian warning systems.

3.3 Research Axis 3: I/O management, in situ visualization and analysis on HPC systems at extreme scales

Our third research axis (mainly dedicated to our HPC-centered activity during the past years) will now be redefined to address challenges posed by the increasing HPC/Big Data/AI convergence at

the application level and the evolutions of the HPC infrastructures that are becoming hybrid as well, as CPU/GPU architectures become the norm for pre-Exascale/Exascale machines.

Towards unified data processing techniques for hybrid simulation/analytics workflows executed across potentially hybrid CPU/GPU infrastructures. In the high-performance computing area (HPC), the need to get fast and relevant insights from massive amounts of data generated by extreme-scale computations led to the emergence of *in situ/in transit processing*. In the Big Data area, the search for real-time, fast analysis was materialized through a different approach: *stream-based processing*. A major challenge is the joint use of these techniques in a unified data processing architecture.

Preliminary work already started within the "frameworks" work package of the HPC-Big Data Inria Challenge. It is also a core direction of our team's involvement in the ACROSS H2020 EuroHPC project. A typical scenario considered in ACROSS consists in executing hybrid workflows combining simulations and (potentially learning-based) analytics running concurrently.

The challenge is to integrate both stream and in-situ/in-transit processing tasks in the targeted workflows, leading to a decrease in execution times for data-intensive/deep learning like HPC simulations and modeling workloads. In particular, we will introduce programmatic support for on-demand data analytics on platforms that were traditionally used only for simulations. This new type of workflow (combining simulations with data analytics) could help anticipate the future behavior of the simulated systems.

Analyzing and exploiting stored data jointly with simulated data can provide a richer tool for much deeper interpretation of the targeted systems, enabling more reliable, transparent and innovative decision making. To this purpose, Damaris will be extended to support asynchronously Big Data analytics plugins, to enable in-situ and in transit analysis of simulation data, then to support hybrid (stream-based and batch-based) in transit data analysis. These new, hybrid workflows will allow on one hand to reduce the simulation time (by pre-analyzing some parts of the results locally, in-situ) and on the other hand to use simulations to train proxy models for optimization.

In the EUPEX EuroHPC Project, one goal is to introduce cross-application optimizations for data-driven stream parallel applications. This will rely on Damaris to orchestrate transfers, by leveraging various storage capabilities to provide scalable asynchronous I/O and non-intrusive in situ and in transit data processing on the data nodes. This provides another motivation to adapt Damaris to support workflows and Big Data analytics plugins by enabling in-situ and in-transit analysis of stream data.

Finally, as a piece of software considered for the Inria Exascale Software Task, in collaboration with the CEA, we plan to investigate new types of scenarios for hybrid CPU/GPU machines, where simulations could trigger on-demand analytics potentially run on GPU hardware.

4 Application domains

The KerData team investigates the design and implementation of architectures for data storage and processing across clouds, HPC and edge-based systems, which address the needs of a large spectrum of applications. The use cases we target to validate our research results come from the following domains.

4.1 Radio astronomy

The international SKA [60] project aims to create the largest telescope in the world in order to observe a part of the universe. A very large volume of data is generated at the telescope level, pre-processed on local clusters (filtering, reduction) in real time and sent to a supercomputer (SDP) at a rate of 1TB/s. This data feeds numerical simulation, generating 1PB of daily output data that

needs to be saved. At this stage, the computing power and storage resources required are such that machines capable of reaching the exascale become necessary. However, the efficient use of these systems raises new challenges, especially regarding data management.

In the context of the ExaDoST project (NumPEX PEPR), for which SKA is one of the main target demonstrators, we are working on optimizing the I/O of a data processing pipeline that is a serious candidate for the radio telescope. This work has also taken the form of active participation in the ECLAT (Extreme Computing Lab for Astronomical Telescopes) joint laboratory [55].

4.2 Climate and meteorology

The European Centre for Medium-Range Weather Forecasts (ECMWF) [56] is one of the largest weather forecasting centers in the world that provides data to national institutions and private clients. ECMWF's production workflow collects data at the edge through a large set of sensors (satellite devices, ground and ocean sensors, smart sensors). This data, approximately 80 million observations per day, is then moved to be assimilated, i.e. analyzed and sorted, before being sent to a supercomputer to feed the prediction models.

The compute and I/O intensive large-scale simulations built upon these models use ensemble forecasting methods for the refinement. To date, these simulations generate approximately 60 TB per hour, while the center predicts an annual increase of 40 % of this volume. Structured datasets called "products" are then generated from this output data and are disseminated to different clients, such as public institutions or private companies, at a rate of 1PB per month transmitted.

In the framework of the ACROSS EuroHPC Project started in 2020, our goal is to participate in the design of a hybrid software stack for the HPC, Big Data and AI domains. This software stack must be compatible with a wide range of heterogeneous hardware technologies and must meet the needs of the trans-continuum ECMWF workflow.

4.3 Earth science

Earthquakes cause substantial loss of life and damage to the built environment across areas spanning hundreds of kilometers from their origins. These large ground motions often lead to hazards such as tsunamis, fires and landslides. To mitigate the disastrous effects, a number of Earthquake Early Warning (EEW) systems have been built around the world. Those critical systems, operating 24/7, are expected to automatically detect and characterize earthquakes as they happen, and to deliver alerts before the ground motion actually reaches sensitive areas so that protective measures can be taken.

One goal of our research is to improve the accuracy of Earthquake Early Warning (EEW) systems. These systems are designed to detect and characterize medium and large earthquakes before their damaging effects reach a certain location. Traditional EEW methods based on seismometers fail to accurately identify large earthquakes due to their low sensitivity to ground motion velocity. The recently introduced high-precision GPS stations, on the other hand, are ineffective to identify medium earthquakes due to their propensity to produce noisy data. In addition, GPS stations and seismometers may be deployed in large numbers across different locations and may produce a significant volume of data consequently, affecting the response time and the robustness of EEW systems.

Integrating and processing in a timely manner high-frequency data streams from multiple sensors scattered over a large territory requires high-performance computing techniques and equipments. We therefore design distributed machine learning-based approaches [5] to earthquake detection, jointly with experts in machine learning and Earth data. Our expertise in swift processing of data on edge and cloud infrastructures allows us to learn from the data from the large number of sensors arriving at high sampling rate, without transferring all data to a single point and thus enables real-time alerts.

4.4 Sustainable development through precision agriculture

Feeding the growing world's population is a on-going challenge, especially in view of climate change, which adds a certain level of uncertainty in food production. Sustainable and precision agriculture is one of the answers that can be implemented to partly overcome this issue. Precision agriculture consists in using new technologies to improve crop management by considering environmental parameters such as temperature, soil moisture or weather conditions, for example. These techniques now need to scale up to improve their accuracy. Over recent years, we have seen the emergence of precision agriculture workflows running across the digital continuum, that is to say all the computing resources from the edge to High-Performance Computing (HPC) and Cloud-type infrastructures. This move to scale is accompanied by new problems, particularly with regard to data movements.

CybeleTech [54] is a French company that aims at developing the use of numerical technologies in agriculture. The core products of CybeleTech are based on numerical simulation of plant growth through dedicated biophysical models and machine learning methods extracting knowledge through large databases. To develop its models, CybeleTech collects data from sensors installed on open agricultural plots or in crop greenhouses. Plant growth models take weather variables as input and the accuracy of agronomic indices estimation heavily rely on the accuracy of these variables.

To this purpose, CybeleTech wishes to collect precise meteorological information from large forecasting centers such as the European Center for Medium-Range Weather Forecasting (ECMWF) [56]. This data gathering is not trivial since it involves large data movements between two distant sites under severe time constraints. In the context of the EUPEX EuroHPC project, our team is exploring innovative data management techniques and data movement algorithms to accelerate the execution of these hybrid geo-distributed workflows running on large-scale systems in the area of precision agriculture.

4.5 Smart cities

The proliferation of small sensors and devices that are capable of generating valuable information in the context of the Internet of Things (IoT) has exacerbated the amount of data flowing from all connected objects to cloud infrastructures. In particular, this is true for Smart City applications. These applications raise specific challenges, as they typically have to handle small data (in the order of bytes and kilobytes), arriving at high rates, from many geographical distributed sources (sensors, citizens, public open data sources, etc.) and in heterogeneous formats, that need to be processed and acted upon with high reactivity in near real-time.

Our vision is that, by smartly and efficiently combining the data-driven analytics at the edge and in the cloud, it becomes possible to make a substantial step beyond state-of-the-art prescriptive analytics through a new, high-potential, faster approach to react to the sensed data of the smart cities. The goal is to build a data management platform that will enable comprehensive joint analytics of past (historical) and present (real-time) data, in the cloud and at the edge, respectively, allowing to quickly detect and react to special conditions and to predict how the targeted system would behave in critical situations. This vision is the driving objective of our SmartFastData associate team with Instituto Politécnico Nacional, Mexico.

In a similar context, smart homes by leveraging numerous sensors and connected devices aim at improving the quality of life, security and making better use of the energy. This is one target of the ENGAGE project.

4.6 Botanical Science

Pl@ntNet [45] is a large-scale participatory platform dedicated to the production of botanical data through AI-based plant identification. Pl@ntNet's main feature is a mobile app allowing

smartphone owners to identify plants from photos and share their observations. It is used by around 10 million users all around the world (more than 180 countries) and it processes about 400K plant images per day. One of the challenges faced by Pl@ntNet engineers is to anticipate what should be the appropriate evolution of the infrastructure to pass the next spring peak without problems and also to know what should be done the following years.

Our research aims to improve the performance of Pl@ntNet. Reproducible evaluations of Pl@ntNet on large-scale testbed (e.g., deployed on Grid'5000 [42] by E2Clab [14]), aim to optimize its software configurations in order to minimize the user response time.

5 Social and environmental responsibility

5.1 Footprint of research activities

HPC and cloud facilities are expensive in capital outlay (both monetary and human) and in energy use and it is clear that there is a related environmental impact inherent to this area. Our work on Damaris supports the efficient use of high performance computing resources. Damaris [3] can help to minimize power needed in running computationally demanding engineering applications and can reduce the amount of storage used for results, thus supporting environmental goals and improving the cost effectiveness of running HPC systems. For the future, our scientific project for the upcoming 4 years starting with 2025 will include specific research directions to address challenges posed by sustainability and climate change, including research on frugal storage and on ways to leverage second-hand HPC hardware.

Another aspect worth mentioning is that our team has strong and active international collaborations which sometimes require intercontinental travels by plane. To minimize carbon footprint, we are careful to keep a balance between a few physical meetings (necessary to maintain substantial exchanges) and remote meetings by videoconference (used in most cases, when traveling is not necessary).

5.2 Impact of research results

Social impact. One of our target applications is Early Earthquake Warning. We proposed a solution that enables earthquakes classification with an outstandingly perfect accuracy. By enabling accurate identification of strong earthquakes, it becomes possible to trigger adequate measures and save lives. For this reason, our work was distinguished with an Outstanding Paper Award — Special Track for Social Impact at AAAI-20, an A* conference in the area of Artificial Intelligence. This result has been highlighted by the *Le Monde* journal in its edition of December 28, 2020, in a section entitled: *Ces découvertes scientifiques que le Covid-19 a masquées en 2020*. This collaborative work continued beyond 2020.

Environmental impact. As presented in Section 4, we are partners with CybeleTech in the framework of the EUPEX EuroHPC project. CybeleTech is a French company specialized in precision agriculture. Within the framework of our collaboration, we propose to focus our efforts on a scale-oriented data management mechanism targeting two CybeleTech use-cases. They address irrigation scheduling for orchards and optimal harvest date for corn, and their models require the acquisition of large volumes of remote data. The overall goal is to improve the accuracy of plant growth models and improve decision making for precision agriculture. The underlying approach in favour of precision agriculture has also encountered some criticism¹.

¹"Digital and Precision Agriculture – Criticisms and Concerns", 2020

6 Highlights of the year

6.1 Awards

Two of the team's papers [27, 28] have been nominated for the Best Paper Award at the HiPC conference in December 2024 in Bangalore, India.

6.2 SC'24 Program Chair

Guillaume Pallez was the Program Chair of SC'24, the top conference in the area of HPC (around 18,000 participants this year).

6.3 HiPC 2025 Program Chair

François Tessier has been selected to serve as a Program Co-Chair of HiPC 2025, a selective international conference in our domain.

6.4 ISPDC 2025 General Co-Chairs

Alexandru Costan and François Tessier will be the General Co-Chairs of ISPDC 2025, a distributed systems conference which will be organized in Rennes in 2025.

7 New software, platforms, open data

7.1 New software

7.1.1 Damaris

Keywords: Visualization, I/O, HPC, Exascale, High performance computing

Scientific Description: Damaris is a middleware for I/O and data management targeting large-scale, MPI-based HPC simulations. It initially proposed to dedicate cores for asynchronous I/O in multicore nodes of recent HPC platforms, with an emphasis on ease of integration in existing simulations, efficient resource usage (with the use of shared memory) and simplicity of extension through plug-ins.

Over the years, Damaris has evolved into a more elaborate system, providing the possibility to use dedicated cores or dedicated nodes to in situ data processing and visualization. It proposes a seamless connection to the VisIt visualization framework to enable in situ visualization with minimum impact on run time. Damaris provides an extremely simple API and can be easily integrated into the existing large-scale simulations.

Damaris was at the core of the PhD thesis of Matthieu Dorier, who received an Accessit to the Gilles Kahn Ph.D. Thesis Award of the SIF and the Academy of Science in 2015. Developed in the framework of our collaboration with the JLESC – Joint Laboratory for Extreme-Scale Computing, Damaris was the first software resulted from this joint lab validated in 2011 for integration to the Blue Waters supercomputer project. It scaled up to 16,000 cores on Oak Ridge's leadership supercomputer Titan (first in the Top500 supercomputer list in 2013) before being validated on other top supercomputers. Active development is currently continuing within the KerData team at Inria, where it is at the center of several collaborations with industry as well as with national and international academic partners.

In 2023, in the context of the ACROSS EuroHPC project, we added an interface for Damaris to enable asynchronous analytics, in particular to support Dask (www.dask.org), a Python-based library for scalable analytics. Dask offers a suite of useful distributed analytic methods using familiar Python-like interfaces, similar to NumPy and Pandas. Our proposed Python interface has enabled access to the suite of Python based visualization libraries and Damaris has been successfully tested with new options for in situ visualization.

Damaris has been selected to be one of the key software pieces of software for the NumPEX PEPR project, which aims to provide the software infrastructure for the future Exascale machine to be hosted in France in 2025 (Jules Vernes project). The capabilities within Damaris will further studied in collaboration with CEA within the NumPEX exploratory PEPR project.

Functional Description: Damaris is a middleware for data management and in-situ visualization targeting large-scale HPC simulations. Damaris enables: - In-situ data analysis by using selected dedicated cores/nodes of the simulation platform. - Asynchronous and fast data transfer from HPC simulations to Damaris. - Semantic-aware dataset processing through Damaris plug-ins, - Writing aggregated data (by hdf5 format) or visualizing them either by VisIt or ParaView. - Dask analytics supports.

Release Contributions: v1.10 Improved Dask data sharing in damaris4py library. v1.11 adds better support for variables not written in every iteration and a C API function for introspection on what Damaris plugins are available. v1.12 adds Debian / Redhat compatible packages generation for distribution, enables the initialization of Damaris with an XML object containing the configuration, and adds the ability of Damaris configuration creation programmatically to ease integration with other libraries (PDI for instance).

News of the Year: In 2024, the support for Dask data sharing was enhanced. In addition, in the context of the NumPEX project, we started the development of the PDI / Damaris plugin to enable a running integration of the Damaris library with PDI (<https://pdi.dev/main/>), an interface for data access, read and write data from HDF5 files, in-situ visualization, and more.

URL: <https://project.inria.fr/damaris/>

Contact: Gabriel Antoniu

Participants: Jean Etienne Ndamlabin Mboula, Gabriel Antoniu, Lokman Rahmani, Luc Bouge, Matthieu Dorier, Orçun Yildiz, Hadi Salimi, Joshua Bowden

Partner: ENS Rennes

7.1.2 E2Clab

Name: Edge-to-Cloud lab

Keywords: Distributed systems, Cloud, Reproducibility, Experimentation, Computing Continuum, Evaluation, Large scale, Provenance

Scientific Description: E2Clab is a framework that implements a rigorous methodology that provides guidelines to move from real-life application workflows to representative settings of the physical infrastructure underlying this application in order to accurately reproduce its relevant behaviors and therefore understand and optimize end-to-end performance.

E2Clab allows a rigorous analysis of possible application configurations in a controlled testbed environment to understand their behavior and related performance trade-offs. E2Clab can be generalized to other applications in the Edge-to-Cloud Continuum. E2Clab is currently used by the PI@ntNet team to understand and optimize the performance of the application. It is also used by our partners from Instituto Politécnico Nacional for automatic experiment deployments in the context of the SmartFastData associate team.

In an effort to enhance the reproducibility capabilities of E2Clab, we extended it to enable efficient provenance data capture across the Edge-to-Cloud Continuum. Specifically, we leverage simplified data models, data compression and grouping, and lightweight transmission protocols to reduce overheads for collecting such data on the IoT/Edge. This integration makes E2Clab a promising platform for the performance optimization of applications through reproducible experiments.

Functional Description: E2Clab is a framework that implements a rigorous methodology that provides guidelines to move from real-life application workflows to representative settings of the physical infrastructure underlying this application in order to accurately reproduce its relevant behaviors and therefore understand end-to-end performance. Understanding end-to-end performance means rigorously mapping the scenario characteristics to the experimental environment, identifying and controlling the relevant configuration parameters of applications and system components, and defining the relevant performance metrics.

Release Contributions: Changelog: https://gitlab.inria.fr/E2Clab/e2clab/-/blob/master/CHANGELOG.rst?ref_type=heads

Features (release 1.0.0):

(i) the configuration of the experimental environment, libraries and frameworks, (ii) the mapping between the application parts and machines on the Edge, Fog and Cloud, (iii) the deployment of the application on the infrastructure, (iv) Edge-to-Cloud network emulation, (v) the automated execution and monitoring, (vi) the application optimization, and (vii) the gathering of experiment metrics.

News of the Year: In an effort to make E2Clab the most complete framework for edge-to-cloud experiments, a support for energy monitoring (Kwollect) on the Grid'5000 platform has been added, paving the way for energy-aware infrastructure optimization and other experiments.

Additional contributions include: - Ongoing experiments on AI workflows on the computing-continuum and data provenance for explainable AI models. - Adapting existing provenance data capture to work with Flowcept, a runtime provenance data exporter for ML workflows developed at ORNL. - An improved documentation and CI pipeline has been implemented to ensure better reliability and stable behaviour across versions with improved code quality and throughout testing. - Improved logging files in experiment artifacts to facilitate the software's user-friendliness for members of the scientific community.

Latest release archive: <https://gitlab.inria.fr/E2Clab/e2clab/-/releases/v3.3.1>

URL: <https://e2clab.gitlabpages.inria.fr/e2clab/>

Publications: [hal-04208787](#), [hal-04779813](#), [hal-04698619](#), [hal-02916032](#), [hal-03310540](#), [hal-03269852](#), [hal-03332524](#), [hal-03270129](#), [hal-03338520](#), [hal-03324177](#), [hal-03259975](#), [hal-03409405](#), [hal-03510012](#), [hal-04659211](#)

Contact: Gabriel Antoniu

Participants: Thomas Badts, Daniel Rosendo, Gabriel Antoniu, Alexandru Costan, Mathieu Simonin

7.1.3 StorAlloc

Keywords: Simulation, HPC, Distributed Storage Systems

Functional Description: StorAlloc is a simulator of a job scheduler dedicated to heterogeneous storage resources. It allows to model storage infrastructures, to simulate their partitioning and allocation, and to evaluate various scheduling algorithms.

In practice, StorAlloc takes a storage request as input, which represents the presumed storage requirements of a job executed on a HPC system. It then offers to select some fitting storage

resources, to be used by the client job. Storage resources are defined by the users, thanks to a YAML format with storage nodes and disks. Their selection happens by mean of an algorithm, also chosen by user (either from predefined algorithms, or user-developed). During simulation, various metrics are stored by StorAlloc all along the processing of storage requests, and eventually written to file when the simulation ends. Components of StorAlloc are independent and communicate through messages. They are easily extensible and new components may also be added.

URL: <https://github.com/heptaicie/storalloc>

Publications: [hal-03922866](#), [hal-03878252](#), [hal-03683568](#)

Contact: François Tessier

Participants: Julien Monniot, François Tessier, Gabriel Antoniu

7.1.4 Fives

Name: Simulator for Scheduling on Storage System at Scale

Keywords: Simulation, HPC, Distributed Storage Systems

Scientific Description: Development of Fives began in 2023, given the limitations of our previous StorAlloc simulator. At the end of 2023, Fives is still in active development, while its design and initial results are being submitted to a conference in the field.

Functional Description: Fives is a storage resource scheduling simulator for supercomputers based on WRENCH and SimGrid, two state-of-the-art simulation frameworks. In particular, Fives can model a parallel file system such as Lustre, a computing partition, and simulate a set of jobs performing I/O on the resulting HPC system.

Fives is based on several components. Firstly, as part of the development of this simulator, an abstraction called "Compound Storage Service" was proposed to represent a distributed storage system, and integrated into WRENCH. Within Fives, a job model was designed to represent a history of jobs and submit them to the scheduler present in WRENCH. Finally, a model of an existing supercomputer, Theta at Argonne National Laboratory, and a reverse-engineered version of its Lustre file system were developed in our simulator.

Experiments are underway to calibrate and validate Fives.

Publication: [hal-04784808](#)

Contact: François Tessier

7.1.5 MOSAIC

Name: Merging Operations and SegmentAtion for I/O Categorization

Keywords: Categorization, HPC, I/O

Scientific Description: MOSAIC is a Python categorizer that takes I/O traces as input and assigns classes to describe the patterns found inside.

Those classes form a general description of applications' I/O activity, giving information about the temporality of I/O, whether periodic operations occur, and an estimation of the impact on the metadata servers.

One of MOSAIC's building blocks is the automatic detection of recurring operations. This is achieved with a clustering algorithm that groups operations sharing the same characteristics (duration, I/O amount, etc.) into one single recurring operation.

MOSAIC automatically finds the traces that were generated by the same program to reduce the number of files to be processed and speed up a system-scale categorization.

MOSAIC works for now with traces from the Darshan monitoring tool but can be easily extended to fit other trace formats.

MOSAIC was used to process the 2019 traces from the BlueWaters supercomputer trace dataset (National Center for Supercomputing Applications - University of Illinois).

Functional Description: MOSAIC is a tool for categorizing HPC application storage activity. It processes traces containing all application storage operations and assigns classes to describe how they are performed.

MOSAIC can describe when the activity is performed (when the application starts, at the end, throughout the execution, etc.), find if some operations are recurring (e.g., saving data to a file every 10 minutes), and estimate the overhead caused by the metadata operations.

It can analyze large datasets of I/O traces coming from a supercomputer to find the general behavior of the applications that were carried out on the machine.

News of the Year: All features of the first version of MOSAIC were developed in 2024.

Publication: [hal-04808300](#)

Contact: François Tessier

7.1.6 Neomem

Keywords: Machine learning, Continual Learning, High performance computing

Scientific Description: Deep learning has emerged as a powerful method for extracting valuable information from large volumes of data. However, when new training data arrives continuously (i.e., is not fully available from the beginning), incremental training suffers from catastrophic forgetting (i.e., new patterns are reinforced at the expense of previously acquired knowledge). Training from scratch each time new training data becomes available would result in extremely long training times and massive data accumulation.

Rehearsal-based continual learning has shown promise for addressing the catastrophic forgetting challenge, but research to date has not addressed performance and scalability. To fill this gap, we propose an approach based on a distributed rehearsal buffer that efficiently complements data-parallel training on multiple GPUs to achieve high accuracy, short runtime, and scalability. It leverages a set of buffers (local to each GPU) and uses several asynchronous techniques for updating these local buffers in an embarrassingly parallel fashion, all while handling the communication overheads necessary to augment input mini-batches (groups of training samples fed to the model) using unbiased, global sampling.

We further propose a generalization of rehearsal buffers to support both classification and generative learning tasks, as well as more advanced rehearsal strategies (notably dark experience replay, leveraging knowledge distillation). We illustrate this approach with a real-life HPC streaming application from the domain of ptychographic image reconstruction. We run extensive experiments on up to 128 GPUs of the ThetaGPU supercomputer to compare our approach with baselines representative of training-from-scratch (the upper bound in terms of accuracy) and incremental training (the lower bound). Results show that rehearsal-based continual learning achieves a top-5 validation accuracy close to the upper bound, while simultaneously exhibiting a runtime close to the lower bound.

Functional Description: Training neural networks with continuously generated data poses challenges related to forgetting previously acquired knowledge, a phenomenon known as "catastrophic forgetting." An effective solution involves replaying certain previously observed data to maintain associated knowledge. Neomem implements this approach, aiming to achieve excellent predictive performance at the cost of a slight increase in training time. Our

approach allows for the utilization of dozens of GPUs, making it applicable to scientific simulations within the high-performance computing (HPC) community.

News of the Year: We introduced a set of abstractions that do not impose any particular constraints on the data shape of training samples retained in the rehearsal buffer, allowing to accommodate a large range of deep learning workloads and rehearsal strategies. Besides, we introduce the concept of annotated tuples of tensors to serve representative training samples and their associated states conveniently to the AI runtime. Such an approach (1) addresses the need to support more rehearsal strategies (notably strategies leveraging knowledge distillation) while (2) transparently augmenting minibatches produced by data pipelines of data-parallel CL training instances.

Publication: [hal-04600107](#)

Contact: Alexandru Costan

Participants: Thomas Bouvier, Alexandru Costan, Gabriel Antoniu

7.1.7 FLAdversary

Name: Emulation of Federated Learning Scenarios with Adversarial Clients

Keywords: Federated learning, Emulation, Adversarial attack

Functional Description: Federated Learning (FL) is subject to diverse threats from the Edge of the network where local training runs on widely distributed, heterogeneous and volatile resources.

FLAdversary provides tools to dynamically introduce adversarial attacks into the FL training phase. Different (model and data) poisoning attacks can be introduced at the client level to emulate adversaries in the FL training. Several defensive strategies are provided as baselines.

Publication: [hal-04208787](#)

Contact: Gabriel Antoniu

Partner: DFKI (German Research Center for Artificial Intelligence)

7.1.8 FLDrift

Name: Emulation of Federated Learning Scenarios with Client Drift

Keywords: Federated learning, Emulation, Heterogeneous Data

Functional Description: When deploying Federated Learning (FL) on the Computing Continuum, devices are subject to high variations in local data distributions. This limits the capacity of the system to generate a single model optimized for the entire federation of devices.

FLDrift provides support for various Non-IID scenarios (i.e., introducing concept-drift and label-shift between federated peers) for FL experiments. Several personalization/clustering strategies are provided as baselines.

News of the Year: We implemented several baseline clustering strategies improving personalization in Federated Learning to address client drift. FLDrift proposes 4 scenarios to evaluate the performance of clustering approaches. Each scenario introduces a different form of concept drift between client local datasets.

Publication: [hal-04779813](#)

Contact: Gabriel Antoniu

Partner: DFKI (German Research Center for Artificial Intelligence)

8 New results

8.1 Convergence of HPC and Big Data infrastructures for supporting workflows in the Computing Continuum

8.1.1 Provisioning storage resources for HPC and Cloud systems

Participants: François Tessier, Julien Monniot, Gabriel Antoniu.

Collaboration. *This work has been carried out in close co-operation with [Henri Casanova](#) (University of Manoa, HI, USA)*

One of the recent axis we are developing in the context of high-performance data access concerns the provisioning of storage resources. The way these resources are accessed on supercomputers and clouds opposes a complex low-level vision that requires tight user control (on supercomputers) and a very abstract vision that implies uncertainty for performance modeling (on clouds). Nevertheless, taking full advantage of all available resources is critical in a context where storage is central for coupling workflow components. Our goal is then to make heterogeneous storage resources distributed across HPC+Cloud infrastructures allocatable and elastic to meet the needs of I/O-intensive hybrid workloads.

This is the context of Julien Monniot's thesis (thesis started in October 2021). He explores the techniques for scheduling storage resources on large-scale systems through simulation. The modeling of storage infrastructures and the evaluation of storage-aware scheduling algorithms are the main contributions of this work.

The work started in 2022 around a first proof-of-concept called StorAlloc [49, 48, 50] evolved in 2023 and 2024 into Fives (Simulator for Scheduling on Storage Systems at Scale), a new simulator implemented with WRENCH [43], a state-of-the-art simulation framework. Fives not only reproduces the results of StorAlloc, but goes far beyond it. Thanks to this simulator, we were able to model a Lustre parallel file system (both hardware and software, which we reverse-engineered), as well as a supercomputer mimicking Theta, a 11 PFlops HPC system at Argonne National Laboratory. Using a job abstraction layer we devised, Fives enabled us to simulate several weeks of job execution on Theta from an I/O point of view (Darshan traces). The aim of this simulation was to calibrate our simulator with a view to subsequently predicting the Lustre's I/O performance on any other dataset. This work was accepted and presented in an A-rank conference in the field [27]. Our paper was a Best Paper Award nominee.

8.1.2 Data profiling and benchmarking of computational workflows beyond the classical Computing Continuum

Participants: Silvina Caino Lores, Gabriel Antoniu.

Collaboration. *This work has been carried out in cooperation with [Elaine Wong](#) and [Travis Humble](#) (Oak Ridge National Laboratory, USA)*

Quantum Computing (QC) systems are increasingly explored as the next high-impact extension to the Computing Continuum, in particular with integration into supercomputers like Joliot-Curie at GENCI (France) and cloud environments like Amazon Web Services. As part of our leadership roles in the [Workflows Community Initiative](#) and workflow-specific venues like the [WORKS workshop](#), we have identified that the successful interoperability between these systems in the

Computing Continuum will depend on middleware able to interact with heterogeneous hardware technologies (e.g., CPU, GPU, TPU, FPGA, quantum processing unit (QPU)) and their associated software stacks and data management methods [36]. In the context of QC, immediate approaches towards integrating QC into the wider Computing Continuum are hybrid and involve interaction with classical systems [41]. This leads to complex open challenges on how to combine multiple programming models in a single application with workflow steps combining quantum and classical processing in a domain-agnostic manner.

In our previous work, we explored avenues of refinement for the QPU in the context of many-tasks management in order to understand the value it can play in linking QC with HPC [25]. Notably, current works on the integration of QC into classical computing ecosystems focus on the interoperability and performance of the algorithms without considering data-oriented optimizations (e.g., data encoding, arrangement, locality, or mapping to high-level data abstractions), and workflow-specific challenges like task-resource mapping are rarely explored [52]. We are conducting exploratory work on profiling and characterising data access and transfer patterns in small scenarios of variational quantum algorithms [44]—a feasible near-term approach to QC integration— with the goal of developing new workflow benchmarks.

8.2 Advanced data processing support for Artificial Intelligence across the Computing Continuum

8.2.1 Investigating the Gap between Simulation, Emulation and Real-World Deployments for Reproducible Federated Learning Experiments

Participants: Cédric Prigent, Alexandru Costan, Gabriel Antoniu.

Collaboration. *This work has been carried out in close co-operation with Kate Keahey, University of Chicago / Argonne National Laboratory, who supervised the internship of Cédric Prigent.*

A common practice in the Federated Learning (FL) literature is to run simulations on a single compute node to assess the performance of FL algorithms. While simulation enables fast prototyping and validation of algorithmic concepts, it may face limitations in reproducing the real system’s performance in heterogeneous environments such as the Computing Continuum, and particularly on resource-constrained Edge devices. Conversely, emulation on distributed testbeds offers more effective means to accurately reproduce the performance of real-world devices.

However, to the best of our knowledge, no prior research has investigated the differences between simulation and emulation in FL experiments. In this work, we study the complementarity of these approaches and discuss their respective challenges, as a first step towards reproducibility of FL experiments. We illustrate our study with a real-life application used as a baseline: an outdoor air quality forecasting framework with real-world sensors. Our results show that simulation can be used to accurately reproduce model performance metrics, while emulation can effectively reproduce the system performance of real-world experiments. Finally, we present a set of lessons learned on the challenges of FL reproducibility and the selection of experimental infrastructures for FL experiments and applications.

This work was accepted and presented in an A-rank conference in the field [28]. This paper was also Best Paper Award nominee.

8.2.2 Formalisation of workflow provenance for trustworthy and explainable AI

Participants: Silvina Caino-Lores, Alexandru Costan.

Collaboration. *This work is part of an ongoing collaboration with Renan Souza and Rafael Ferreira da Silva (Oak Ridge National Laboratory, USA).*

Artificial Intelligence (AI) is driving scientific discovery and economic growth in all kinds of application domains while impacting from routine daily tasks to societal-level challenges. However, research communities, industry players and social actors are expressing increasing concern about the potential ethical and practical implications of the pervasive presence of AI. Of particular concern are the explainability of AI, or making AI's decision-making process understandable, and transparency of AI, ensuring clarity in AI's design, data and operation. Working towards advancing explainability and transparency of AI is currently a priority, essential for responsible and trustworthy AI applications.

In our work, we have formalised the AI lifecycle and developed a vision of workflow provenance as a tool for data-to-insights, identifying key research priorities and open challenges [29]. Provenance refers to the origin or history of data and models, capturing the lineage of their creation, modification, and usage. Provenance in the form of metadata captured at runtime during the execution of AI workflows provides a detailed record of data sources, processing steps, and model configurations, ensuring transparency and traceability throughout the AI lifecycle. A challenging aspect of working with AI workflows is that today there are no comprehensive formalisms able to capture the complexity and relationships in workflow and model provenance data. Our ongoing work towards the definition of ontologies and taxonomies for AI workflow provenance data aims to fill this gap and serve as a theoretical foundation for developing tailored provenance data management systems tailored for the different stakeholders involved in AI applications. In addition, we are exploring new analysis and visualisation techniques to facilitate the integration and inspection of heterogeneous and complex provenance data.

8.2.3 Efficient Resource-Constrained Federated Learning Clustering with Local Data Compression on the Edge-to-Cloud Continuum

Participants: Cédric Prigent, Alexandru Costan, Gabriel Antoniu.

Collaboration. *This work has been carried out in close co-operation with Loïc Cudennec (DGA), who is co-advising the PhD thesis of Cédric Prigent, and with DFKI in the context of the ENGAGE Inria-DFKI project.*

Federated Learning (FL) has been proposed as a privacy-preserving approach for distributed learning over decentralized resources. While it can be a highly efficient tool for large-scale collaborative training of Machine Learning (ML) models, its efficiency may be strongly impacted by a high variability in data distributions among clients [20]. Clustered FL tackles this problem by grouping clients with similar data distributions and training personalized models. Despite increasing model accuracy for federated peers, existing clustering approaches overlook system and infrastructure constraints leading to sustainability problems for resource-constrained devices.

In [28], we introduce a new method for resource-constrained FL clustering. We leverage pre-trained autoencoders to compress client data into low dimensional space and build lightweight embedding vectors used to cluster federated clients. A randomized quantization approach specifically secures the client embedding vectors against data reconstruction. Extensive experiments using a multi-GPU testbed with multiple scenarios introducing concept drift between clients demonstrate the generality of our approach to personalized FL. While each of the baselines

encounters performance degradation in at least one of the scenarios, our strategy demonstrates top efficiency in all of them.

8.2.4 Efficient Data-Parallel Continual Learning with Asynchronous Distributed Rehearsal Buffers

Participants: Thomas Bouvier, Alexandru Costan, Gabriel Antoniu.

Collaboration. *This work has been carried out in the context of a JLESC project (Continual Learning at Scale) in close co-operation with Bogdan Nicolae (Argonne National Laboratory-ANL, USA), who serves as a technical advisor for the PhD work of Thomas Bouvier.*

Deep learning has emerged as a powerful method for extracting valuable information from large volumes of data. However, when new training data arrives continuously (i.e., is not fully available from the beginning), incremental training suffers from catastrophic forgetting (i.e., new patterns are reinforced at the expense of previously acquired knowledge). Training from scratch each time new training data becomes available would result in extremely long training times and massive data accumulation. Rehearsal-based continual learning has shown promise for addressing the catastrophic forgetting challenge, but research to date has not addressed performance and scalability. To fill this gap, we propose an approach based on a distributed rehearsal buffer that efficiently complements data-parallel training on multiple GPUs to achieve high accuracy, short runtime, and scalability.

This year we introduced a set of abstractions that do not impose any particular constraints on the data shape of training samples retained in the rehearsal buffer, allowing to accommodate a large range of deep learning workloads and rehearsal strategies. Besides, we introduce the concept of annotated tuples of tensors to serve representative training samples and their associated states conveniently to the AI runtime. Such an approach (1) addresses the need to support more rehearsal strategies (notably strategies leveraging knowledge distillation) while (2) transparently augmenting minibatches produced by data pipelines of data-parallel CL training instances [24, 17]

8.2.5 Comparative Analysis of Federated Learning: Simulations Versus Real-World Testbeds

Participants: Mathis Valli, Alexandru Costan, Gabriel Antoniu.

Collaboration. *This work has been carried out in close co-operation with Loïc Cudennec (DGA). Mathis Valli PhD is co-supervised by Cédric Tedeschi (Myriads team).*

Federated Learning (FL), while a significant advancement in decentralized machine learning, is predominantly assessed through simulations. These simulations often fail to capture the complexities and unpredictability of real-world environments. They usually overlook factors like network variability, device heterogeneity, and real-time operational challenges, leading to a gap between theoretical efficiency and practical applicability.

Addressing these limitations, our current work emphasizes the deployment of FL models on real-world testbeds. By moving beyond purely simulation-based analyses, we aim to better understand how factors such as bandwidth variability and device heterogeneity influence both model convergence and system-level behavior. Specifically, our experiments focus on monitoring performance metrics like model accuracy, training time, communication overhead, and energy consumption to evaluate how FL reacts to realistic conditions. By comparing these real-testbed

outcomes with simulation results, our study intends to highlight the disparities and refine the practical applicability of FL models, bridging the gap between theory and practice.

We conduct these experiments on a heterogeneous set of nodes in Grid'5000, using E2Clab2's orchestration tools to manage and record resource usage data. Through continuous observation and data collection, we are gaining insights into how real operational challenges—like transient node failures or resource contention—affect FL training processes. Our current emphasis remains on quantifying these impacts rather than deploying self-adaptive mechanisms.

While this phase of our research focuses on understanding and quantifying the influence of system constraints, we have also highlighted the potential benefits of adapting FL systems in real time. Our publication [31] presented the case for incorporating dynamic adaptation, suggesting that an adaptive approach could help FL frameworks cope more effectively with evolving operational conditions. Ultimately, we envision robust FL solutions that can autonomously fine-tune their parameters to meet the demands of dynamic environments, thus bridging the remaining gap between controlled simulations and practical, large-scale FL deployments.

8.3 Scalable I/O, in-situ Visualization and Resource Management at Large Scale

8.3.1 Multi-level analysis of the I/O pattern of HPC applications

Participants: François Tessier, Théo Jolivel, Jakob Luettgau, Julien Monniot, Gabriel Antoniu, Guillaume Pallez.

Collaboration. *This work has been carried out in close co-operation with the Inria TADaM team in Bordeaux and Ahmad Tarraf from the Technical University of Darmstadt, Germany*

While the ratio of I/O performance to computing power has declined by a factor of 10 in the last decade², the volume of data generated by scientific workflows and applications has significantly grown. In some supercomputing centers for instance, this volume has increased almost 40-fold in ten years. This has made access to storage resources a major bottleneck to scaling up applications.

Several levers exist along the data path to mitigate this burden. For example, optimizations can be applied at the I/O library level or within the application source code to improve I/O performance. At the job scheduler level, decisions can be taken when allocating resources to avoid I/O interference between jobs. However, all these optimizations require a good upstream understanding of application I/O behavior.

In this research axis, we are working on analyzing the I/O behavior of large-scale applications at various levels. The thesis that Théo Jolivel started in October 2024 proposes to tackle this question. One approach is to exploit public datasets containing several years of I/O execution traces of applications running on supercomputers. We developed multiple methodologies and tools to pre-process those datasets, extract the relevant data, and analyse the data access behavior. In particular, we introduced MOSAIC, a categorizer that detects I/O patterns from execution traces. MOSAIC extracts I/O operations contained in Darshan traces and assigns classes to describe how I/O operations are performed throughout the execution. The description is based on three distinct axes: I/O temporality (when was data read or written?), access periodicity (are there recurring operations?), and metadata overhead (what is the impact of metadata operations?). This work has been published in one of the major workshops at the Supercomputing conference [26] while a complementary work in collaboration with Inria Bordeaux and TU Darmstadt has been accepted at IPDPS, an A-rank conference in the field [37].

²<https://top500.org/>

8.3.2 Topology and affinity-aware data aggregation

Participants: François Tessier.

Collaboration. *This work is the conclusion of several years of collaboration with Emmanuel Jeannot (TADaaM team at Inria Bordeaux) and Venkatram Vishwanath (Argonne National Laboratory) as part of a JLESC project.*

Over the years, moving data from application to the storage system has become more and more challenging. While the amount of data to manage has drastically increased, the capability of the HPC architectures to absorb this burden has reduced. In order to limit concurrency or non-contiguous accesses on file systems, a preliminary phase of data aggregation is often necessary before moving data. In the context of I/O, the two-phase I/O algorithm is a common method which involves selecting a subset of the processing entities, called aggregators, to accumulate contiguous pieces of data (aggregation phase) before writing/reading them to/from the storage system (I/O phase).

As part of a long-term JLESC collaboration, we have been focusing on optimizing data movement on large-scale systems, especially through data aggregation techniques for I/O intensive applications. For that purpose, we designed and developed the TAPIOCA library. TAPIOCA is a C++ I/O library based on the two-phase I/O scheme mentioned above and performing scalable architecture-aware data aggregation. This library features several key innovations: a RDMA-based implementation reducing the data movement cost between data producers and the aggregators, a way to capture the data model and the data layout of the application to optimize the I/O scheduling, and a model and an objective function computing an architecture-aware placement of the aggregators. The work on TAPIOCA has been published in the FGCS journal [21].

8.3.3 Scalable asynchronous I/O and in-situ processing with Damaris

Participants: Joshua Bowden, Etienne Ndamlabin, Gabriel Antoniu.

Collaboration. *This work has been done in collaboration with Atgeirr Rasmussen (SINTEF) and his team within the framework of the EuroHPC H2020 ACROSS project.*

As large-scale simulations can take a long time to run, and require significant high-performance computing resources, we investigate how asynchronous I/O and in situ processing can help improve the performance, scaling, and efficiency of workflow executions. Within the ACROSS project we consider the use case provided by the OPM Flow software, used for a Carbon Sequestration simulation. In 2023 we started to investigate how the Damaris approach developed by the KerData team could be leveraged by OPM Flow to provide asynchronous analytics, in particular to support Dask (www.dask.org), a Python-based library for scalable analytics. Dask offers a suite of useful distributed analytic methods using familiar Python-like interfaces, similar to NumPy and Pandas. Our proposed Python interface has enabled access to the suite of Python based visualization libraries and Damaris has been successfully tested with new options for in situ visualization. This work was continued in 2024 with the goal to improve Damaris's support for Dask data sharing.

The EuroHPC ACROSS project has supported this work and the results are benefiting the OPM Flow simulation software which integrates Damaris in a public release. The developed capabilities within Damaris are now studied in collaboration with CEA within the NumPEX exploratory PEPR project.

8.3.4 Qualitatively Analyzing Optimization Objectives in the Design of HPC Resource Manager

Participants: Robin Boëzennec, Guillaume Pallez.

Collaboration. *This work has been done in collaboration with Fanny Dufosse from the INRIA Datamove team.*

A correct evaluation of scheduling algorithms and a good understanding of their optimization criteria are key components of resource management in HPC. In this axis, we discuss bias and limitations of the most frequent optimization metrics from the literature. We provide elements on how to evaluate performance when studying HPC batch scheduling. We experimentally demonstrate these limitations by focusing on two use-cases: a study on the impact of runtime estimates on scheduling performance, and the reproduction of a recent high-impact work that designed an HPC batch scheduler based on a network trained with reinforcement learning. We demonstrate that focusing on quantitative optimization criterion ("our work improves the literature by X%") may hide extremely important caveat, to the point that the results obtained are opposed to the actual goals of the authors. Key findings show that mean bounded slowdown and mean response time are hazardous for a purely quantitative analysis in the context of HPC. Despite some limitations, utilization appears to be a good objective. We propose to complement it with the standard deviation of the throughput in some pathological cases. Finally, we argue for a larger use of area-weighted response time, that we find to be a very relevant objective.

8.3.5 Allocation Strategies for Disaggregated Memory in HPC Systems

Participants: Robin Boëzennec, Guillaume Pallez.

Collaboration. *This work has been done in collaboration with Fanny Dufosse and Danilo Carastan-Santos from the INRIA Datamove team.*

In this axis we consider scheduling strategies to deal with disaggregated memory for HPC systems. Disaggregated memory is an implementation of storage management that provides flexibility by giving the option to allocate storage based on system-defined parameters. In this case, we consider a memory hierarchy that allows to partition the memory resources arbitrarily amongst several nodes depending on the need. This memory can be dynamically reconfigured at a cost. We provide algorithms that pre-allocate or reconfigure dynamically the disaggregated memory based on estimated needs. We provide theoretical performance results for these algorithms. An important contribution of our work is that it shows that the system can design allocation algorithms even if user memory estimates are not accurate, and for dynamic memory patterns. These algorithms rely on statistical behavior of applications. We observe the impact on the performance of parameters of interest such as the reconfiguration cost.

8.3.6 Scheduling distributed I/O resources in HPC systems

Participants: Guillaume Pallez.

Collaboration. *This work has been done in collaboration with the TADAAM Inria team.*

Parallel file systems cut files into fixed-size stripes and distribute them across a number of storage targets (OSTs) for parallel access. Moreover, a layer of I/O nodes is often placed between compute nodes and the PFS. In this context, it is important to notice both OST and I/O nodes are potentially shared by running applications, which may lead to contention and low I/O performance.

Contention-mitigation approaches usually see the shared I/O infrastructure as a single resource capable of a certain bandwidth, whereas in practice it is a distributed set of resources from which each application can use a subset. In addition, performance measured in practice does not scale proportionally with the proportion of OST used. Indeed, depending on their characteristics, each application will be impacted differently by the number of used I/O resources.

We conducted [22] a comprehensive study of the problem of scheduling shared I/O resources — I/O nodes, OSTs, etc — to HPC applications. We tackled this problem by proposing heuristics to answer two questions: 1) how many resources should we give each application (allocation heuristics), and 2) which resources should be given to each application (placement heuristics). These questions are not independent, as using more resources often means sharing them. Nonetheless, our two-step approach allows for simpler heuristics that would be usable in practice.

In addition to overhead, an important aspect that impacts how “implementable” algorithms are is their input regarding applications’ characteristics, since this information is often not available or at least imprecise. Therefore, we proposed heuristics that use different input and studied their robustness to inaccurate information.

9 Partnerships and cooperations

9.1 International initiatives

9.1.1 Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program

UNIFY 2

Participants: Gabriel Antoniu, Thomas Bouvier, Alexandru Costan, Jakob Luetzgau, Julien Monnot, Cédric Prigent, François Tessier.

Title: Intelligent Unified Data Services for Hybrid Workflows Combining Compute-Intensive Simulations and Data-Intensive Analytics at Extreme Scales - 2

Duration: 2023 ->

Coordinator: Tom PETERKA (tpeterka@mcs.anl.gov)

Partners:

- Argonne National Laboratory Argonne (États-Unis)

Inria contact: Gabriel Antoniu

Summary: Since several years we have been witnessing the emergence of complex workflows combining simulations with data analysis, potentially leveraging machine-learning techniques. Such complex workflows seem to naturally need to jointly use supercomputers interconnected with clouds and potentially Edge-based systems. This assembly is called the Computing Continuum. In a general scheme, Edge devices create streams of input data, which are processed by data analytics and machine learning applications in the Cloud, whereas simulations on large, specialised HPC systems provide insights into and prediction of future system state. The emergence of such workflows is reshaping the traditional vision on the areas involved, as described in the ETP4HPC Research Agenda published in

2020. Building software ecosystems addressing the needs of such workflows poses multiple challenges at several levels. In this context, this Associate Team will focus on three related challenges: 1) How to adequately handle the heterogeneity of storage resources within the Computing Continuum to support complex science workflows? 2) How to efficiently support deep-learning workloads across the Computing Continuum? 3) How to provide reproducibility support for experimentation across the Computing Continuum?

9.2 International research visitors

9.2.1 Visits of international scientists

Sarah Neuwirth: Sarah Neuwirth is a Research Professor at Johannes Gutenberg University Mainz (JGU), Germany. She visited the team in December 2024 and gave a seminar entitled "Toward Explainable I/O for HPC Systems".

Frederic Suter: Frederic Suter is a Senior Research Scientist at Oak Ridge National Laboratory, USA. He visited KerData in December 2024 and gave a talk about his work on data management across large-scale workflows.

9.2.2 Visits to international teams

Research stays abroad

Cédric Prigent

Visited institution: University of Chicago

Country: USA

Dates: 08/07/2024 - 13/09/2024

Context of the visit: Research collaboration with Kate Keahey to work on the deployment of federated learning on real-world air-quality stations.

Mobility program/type of mobility: Internship

Research visits abroad

Silvina Caino Lores

Visited institution: Carlos III University of Madrid

Country: Spain

Dates: 23/09/2024 - 27/09/2024

Context of the visit: Exploration of research collaboration with the Computer Architecture (AR-COS) group. Participation in PhD jury of Dante Sanchez-Gallegos.

Mobility program/type of mobility: Meeting

9.3 European initiatives

9.3.1 H2020 projects

EUPEX

Participants: Joshua Bowden, Etienne Ndamlabin, Gabriel Antoniu.

[EUPEX project on cordis.europa.eu](http://cordis.europa.eu)

Title: EUROPEAN PILOT FOR EXASCALE

Duration: From January 1, 2022 to December 31, 2026

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- GRAND EQUIPEMENT NATIONAL DE CALCUL INTENSIF (GENCI), France
- VSB - TECHNICAL UNIVERSITY OF OSTRAVA (VSB - TU Ostrava), Czechia
- JOHANNES GUTENBERG-UNIVERSITÄT MAINZ, Germany
- FORSCHUNGSZENTRUM JÜLICH GMBH (FZJ), Germany
- COMMISSARIAT À L'ÉNERGIE ATOMIQUE ET AUX ÉNERGIES ALTERNATIVES (CEA), France
- IDRYMA TECHNOLOGIAS KAI EREVNAS (FOUNDATION FOR RESEARCH AND TECHNOLOGYHELLAS), Greece
- SVEUCILISTE U ZAGREBU FAKULTET ELEKTROTEHNIKE I RACUNARSTVA (UNIVERSITY OF ZAGREB FACULTY OF ELECTRICAL ENGINEERING AND COMPUTING), Croatia
- UNIVERSITÀ DEGLI STUDI DI TORINO (UNITO), Italy
- Consortium Ubiquitous Technologies S.c.a.r.l. (CUBIT), Italy
- CYBELETECH, France
- UNIVERSITÀ DI PISA (UNIFI), Italy
- ISTITUTO NAZIONALE DI ASTROFISICA (INAF), Italy
- UNIVERSITÀ DEGLI STUDI DEL MOLISE, Italy
- E4 COMPUTER ENGINEERING SPA (E4), Italy
- UNIVERSITÀ DEGLI STUDI DELL'AQUILA (UNIVAQ), Italy
- JOHANN WOLFGANG GOETHE-UNIVERSITÄT FRANKFURT AM MAIN (GUF), Germany
- EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS (ECMWF), United Kingdom
- BULL SAS (BULL), France
- POLITECNICO DI MILANO (POLIMI), Italy
- EXASCALE PERFORMANCE SYSTEMS - EXAPSYS IKE, Greece
- ALMA MATER STUDIORUM - UNIVERSITÀ DI BOLOGNA (UNIBO), Italy
- PARTEC AG (PARTEC), Germany
- ISTITUTO NAZIONALE DI GEOFISICA E VULCANOLOGIA, Italy
- CINECA CONSORZIO INTERUNIVERSITARIO (CINECA), Italy
- SECO SPA (SECO SRL), Italy
- CONSORZIO INTERUNIVERSITARIO NAZIONALE PER L'INFORMATICA (CINI), Italy

Inria contact: Olivier Beaumont

Coordinator: Etienne Walter (EVIDEN)

Summary: The EUPEX consortium aims to design, build, and validate the first EU platform for HPC, covering end-to-end the spectrum of required technologies with European assets: from the architecture, processor, system software, development tools to the applications. The EUPEX prototype will be designed to be open, scalable and flexible, including the modular OpenSequana-compliant platform and the corresponding HPC software ecosystem for the Modular Supercomputing Architecture. Scientifically, EUPEX is a vehicle to prepare HPC, AI, and Big Data processing communities for upcoming European Exascale systems and technologies. The hardware platform is sized to be large enough for relevant application preparation and scalability forecast, and a proof of concept for a modular architecture relying on European technologies in general and on European Processor Technology (EPI) in particular. In this context, a strong emphasis is put on the system software stack and the applications.

Being the first of its kind, EUPEX sets the ambitious challenge of gathering, distilling and integrating European technologies that the scientific and industrial partners use to build a production-grade prototype. EUPEX will lay the foundations for Europe's future digital sovereignty. It has the potential for the creation of a sustainable European scientific and industrial HPC ecosystem and should stimulate science and technology more than any national strategy (for numerical simulation, machine learning and AI, Big Data processing).

The EUPEX consortium – constituted of key actors on the European HPC scene – has the capacity and the will to provide a fundamental contribution to the consolidation of European supercomputing ecosystem. EUPEX aims to directly support an emerging and vibrant European entrepreneurial ecosystem in AI and Big Data processing that will leverage HPC as a main enabling technology.

ACROSS

Participants: Joshua Bowden, Gabriel Antoniu, Alexandru Costan, François Tessier, Thomas Bouvier.

[ACROSS project on cordis.europa.eu](https://cordis.europa.eu/project/ACROSS)

Title: HPC BIG DATA ARTIFICIAL INTELLIGENCE CROSS STACK PLATFORM TOWARDS EXASCALE

Duration: From March 1, 2021 to February 29, 2024

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- VSB - TECHNICAL UNIVERSITY OF OSTRAVA (VSB - TU Ostrava), Czechia
- MORFO DESIGN SRL, Italy
- NEUROPUBLIC AE PLIROFORIKIS & EPIKOINONION (NEUROPUBLIC SA), Greece
- UNIVERSITA DEGLI STUDI DI FIRENZE (UNIFI), Italy
- UNIVERSITA DEGLI STUDI DI TORINO (UNITO), Italy
- SINTEF AS (SINTEF), Norway
- INSTITUT NATIONAL DES SCIENCES APPLIQUEES DE RENNES (INSA RENNES), France
- STICHTING DELTARES (Deltares), Netherlands
- EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS (ECMWF), United Kingdom

- BULL SAS (BULL), France
- GE AVIO SRL (GE AVIO SRL), Italy
- FONDAZIONE LINKS - LEADING INNOVATION & KNOWLEDGE FOR SOCIETY (FONDAZIONE LINKS), Italy
- UNIVERSITA DEGLI STUDI DI GENOVA (UNIGE), Italy
- MAX-PLANCK-GESELLSCHAFT ZUR FORDERUNG DER WISSENSCHAFTEN EV (MPG), Germany
- CINECA CONSORZIO INTERUNIVERSITARIO (CINECA), Italy
- CONSORZIO INTERUNIVERSITARIO NAZIONALE PER L'INFORMATICA (CINI), Italy

Inria contact: Gabriel Antoniu

Coordinator: Olivier Terzo (FONDAZIONE LINKS)

Summary: Supercomputers have been extensively used to solve complex scientific and engineering problems, boosting the capability to design more efficient systems. The pace at which data are generated by scientific experiments and large simulations (e.g., multiphysics, climate, weather forecast, etc.) poses new challenges in terms of capability of efficiently and effectively analysing massive data sets. Artificial Intelligence, and more specifically Machine Learning (ML) and Deep Learning (DL) recently gained momentum for boosting simulations' speed. ML/DL techniques are part of simulation processes, used to early detect patterns of interests from less accurate simulation results. To address these challenges, the ACROSS project will co-design and develop an HPC, BD, and Artificial Intelligence (AI) convergent platform, supporting applications in the Aeronautics, Climate and Weather, and Energy domains. To this end, ACROSS will leverage on next generation of pre-exascale infrastructures, still being ready for exascale systems, and on effective mechanisms to easily describe and manage complex workflows in these three domains. Energy efficiency will be achieved by massive use of specialized hardware accelerators, monitoring running systems and applying smart mechanisms of scheduling jobs. ACROSS will combine traditional HPC techniques with AI (specifically ML/DL) and BD analytic techniques to enhance the application test case outcomes (e.g., improve the existing operational system for global numerical weather prediction, climate simulations, develop an environment for user-defined in-situ data processing, improve and innovate the existing turbine aero design system, speed up the design process, etc.). The performance of ML/DL will be accelerated by using dedicated hardware devices. ACROSS will promote cooperation with other EU initiatives (e.g., BDVA, EPI) and future EuroHPC projects to foster the adoption of exascale-level computing among test case domain stakeholders.

9.3.2 Collaborations with Major European Organizations

Participants: Gabriel Antoniu, Alexandru Costan, Jakob Luettgau.

ETP4HPC: Since 2019, Gabriel Antoniu has served as a co-leader of the working group on Programming Environments, contributing to two successive versions of the Strategic Research Agenda of ETP4HPC, the latest one being published in 2024. Alexandru Costan served as a member of this working group. Jakob Luettgau served as a member of the working group on Data Storage and I/O.

9.4 National initiatives

Exa-DoST

Participants: Gabriel Antoniu, François Tessier, Julien Monniot, Joshua Bowden, Etienne Ndamlabin, Silvina Caino Lores, Guillaume Pallez.

Exa-DoST project of the NumPEX PEPR program

Title: Data-oriented Software and Tools for the Exascale

Duration: From January 1, 2023 to April 1, 2030

Partners:

- Inria
- CEA
- CNRS
- University of Bordeaux
- Observatoire de Paris
- Observatoire de la Côte d'Azur
- Data Direct Networks France (DDN)

Coordinator: Gabriel Antoniu (KerData Team, Inria)

Summary: The advent of future Exascale supercomputers raises multiple data-related challenges. To enable applications to fully leverage the upcoming infrastructures, a major challenge concerns the scalability of techniques used for data storage, transfer, processing and analytics. Additional key challenges emerge from the need to adequately exploit emerging technologies for storage and processing, leading to new, more complex storage hierarchies. Finally, it now becomes necessary to support more and more complex hybrid workflows involving at the same time simulation, analytics and learning, running at extreme scales across supercomputers interconnected to clouds and edgebased systems. The Exa-DoST project will address most of these challenges, organized in 3 areas:

- Scalable storage and I/O;
- Scalable in situ processing;
- Scalable smart analytics.

As part of the NumPEX program, Exa-DoST targets a much higher technology readiness level than previous national projects concerning the HPC software stack. It will address the major data challenges by proposing operational solutions co-designed and validated in French and European applications. This will allow filling the gap left by previous international projects to ensure that French and European needs are taken into account in the roadmaps for building the data-oriented Exascale software stack.

STEEL

Participants: Gabriel Antoniu, Alexandru Costan, Jakob Luetzgau, François Tessier, Mathis Valli, Thomas Badts.

Title: Secure and efficient daTa storageE and procEssing on cLoud-based infrastructures

Duration: From June 1, 2023 to 31 August 2030

Partners:

- Inria
- CNRS
- Institut Mines Télécom (IMT)
- University of Bordeaux
- University of Rennes
- INSA Rennes
- INSA Lyon

Coordinator: Gabriel Antoniu (KerData Team, Inria)

Summary: The strong development of cloud computing since its emergence in 2007 and its massive adoption for the storage of unprecedented volumes of data in a growing number of domains has brought to light major technological challenges. In this project we will address several of these challenges, organized in three research directions. The first direction concerns the exploitation of emerging technologies for efficient storage on cloud infrastructures. We will address this challenge through NVRAM-based distributed performance storage solutions, as close as possible to data production and consumption locations (disaggregation principle) and develop strategies to optimize the trade-off between data consistency and access performance. The second direction concerns the efficient storage and processing of data on hybrid, heterogeneous infrastructures within the digital edge-cloud-supercomputer continuum. In many domains (autonomous cars, predictive maintenance, intelligent buildings, etc.) we are witnessing the emergence of hybrid workflows combining simulations, analysis of sensor data flows and machine learning. Their execution requires storage resources ranging from the edge to cloud infrastructures, and even to supercomputers, which poses challenges for unified data storage and processing. The third research direction is dedicated to confidential storage, in connection with the need to store and analyze large volumes of data of strategic interest or of a personal nature. For all of these directions, the project will take into account the need to propose and validate interoperable approaches with a potential for transfer to major French or European industrial players in cloud computing.

ECLAT

Participants: François Tessier, Gabriel Antoniu, Théo Jolivel, Jakob Luettgau.

Title: Extreme Computing Laboratory for Astronomical Telescopes

Duration: Since May, 2024

Partners:

- Inria
- CNRS
- Université de Rennes
- Eviden
- Observatoire de la Côte d'Azur
- Observatoire de Paris
- Université Paris-Saclay

- Centrale Supelec

Coordinator: Gabriel Antoniu (KerData Team, Inria)

Summary: ECLAT is positioned as a center of excellence dedicated to High-Performance Computing (HPC) and Artificial Intelligence (AI) technologies and techniques applied to astronomical instrumentation. This project brings together sixteen partner laboratories and teams around a common roadmap, aimed at strengthening research and development (R&D) collaborations. The aim is to design and build future cyber-physical systems for astronomy, capable of managing, processing and optimizing gigantic volumes of data.

Grid'5000 We are members of Grid'5000 community and run experiments on the Grid'5000 platform on a daily basis.

Inria Exploratory program: Repas

Participants: Guillaume Pallez.

Project Acronym: REPAS

Title: New Portrayal of HPC Applications

Coordinator: Guillaume PALLEZ

Collaboration: This is done in collaboration with the team DATAMOVE (Inria Grenoble)

Duration: 2022-2025

What is the right way to represent an application in order to run it on a highly parallel (typically exascale) machine? The idea of project is to completely review the models used in the development scheduling algorithms and software solutions to take into account the real needs of new users of HPC platforms.

10 Dissemination

10.1 Promoting scientific activities

Participants: Gabriel Antoniu, Silvina Caino Lores, Alexandru Costan, Jakob Luetzgau, Julien Monniot, Guillaume Pallez, Cédric Prigent, François Tessier, Mathis Valli.

10.1.1 Scientific events: organisation

General chair, scientific chair

François Tessier:

- General Chair of [ESSA 2024](#), the 5th Workshop on Extreme-Scale Storage and Analysis, held in conjunction with IPDPS 2024 (San Francisco, USA).
- Co-Chair of [SuperCompCloud](#), the 8th Workshop on Interoperability of Supercomputing and Cloud Technologies, held in conjunction with SC'24 (Atlanta, USA).
- General Co-chair of [ISPDC 2025](#), the 24th International Symposium on Parallel and Distributed Computing (Rennes, France), to be held in 2025.

Alexandru Costan:

- General Co-chair of **ISPDC 2025**, the 24th International Symposium on Parallel and Distributed Computing (Rennes, France), to be held in 2025.
- General Co-Chair of **FlexScience 2024**, the 14th Workshop on AI and Scientific Computing at Scale using Flexible Computing Infrastructures, held in conjunction with HPDC 2024 (Pisa, Italy).

Silvina Caino-Lores:

- General Co-Chair of **WORKS 2024**, the 19th Workshop on Workflows in Support of Large-Scale Science, held in conjunction with SC 2024 (Atlanta, USA).

Gabriel Antoniu:

- Steering Committee Chair of the ESSA Workshop series on High-Performance Storage, held in conjunction with the IEEE IPDPS conference since 2020.

Guillaume Pallez:

- Technical Program Chair of SC'24
- Steering Committee of ICPP

Member of the organizing committees**François Tessier:**

- Co-organizer of the 7th **SuperCompCloud** workshop held in conjunction with ISC 2024 (Hamburg, Germany).

Silvina Caino-Lores:

- Workshops and Minisymposia Co-Chair of the **30th International European Conference on Parallel and Distributed Computing (Euro-Par 2024)** (Madrid, Spain).
- Reproducibility Challenge Co-Chair of the **The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC 2024)** (Atlanta, USA).
- Co-organizer of the 1st **Workshop on Workflow Monitoring, Observability, and In Situ Analytics (WOWMON)** held in conjunction with ICPP 2024 (Gotland, Sweden).

Jakob Luettgau:

- Co-organizer of the First Symposium on Ethical, Social and Policy issues in HPC **ESP-HPC** held in conjunction with SC 2024 (Atlanta, USA)

10.1.2 Scientific events: selection**Chair of conference program committees****François Tessier:**

- HPC System Software Track Chair of **HiPC 2024**, the 31st edition of the IEEE International Conference on High Performance Computing, Data, and Analytics (Bengalore, India).

Silvina Caino-Lores:

- Software Systems and Platforms Track Co-Chair of **CCGRID 2025**, the 25th IEEE International Symposium on Cluster, Cloud, and Internet Computing (Tromso, Norway).

Member of the conference program committees

François Tessier: SC'24 (Data Analytics, Visualization and Storage track), ISC'24 (workshop proposals), WORKS 2024, COMPAS 2024

Alexandru Costan: SC'24 (Data Analytics, Visualization and Storage track), IPDPS 2024, Euro-Par 2024, UCC 2024, IEEE BigData 2024, CloudCom 2024

Jakob Luetzgau: SC'24 (Reproducibility Challenge), ICPP24, ISC'24 (Birds of a Feather), PASC'24 (ACM Posters - Student Research Competition), SBAC-PAD 24, CCGrid 24, HiPC'24 IPDPS'25, WORKS 2024, XLOOP 2024, WOCC'24

Silvina Caino-Lores: HPDC 2024, CARLA 2024, WIDE 2024, PDP 2025, IPDPS 2025

Gabriel Antoniu: HPDC 2024, IPDPS 2024

Reviewer

Cédric Prigent: Euro-Par 2024, IEEE BigData 2024

Julien Monniot: IEEE/ACM SC24

Mathis Valli: IEEE BigData 2024

10.1.3 Journal

Member of the editorial boards

Jakob Luetzgau: Guest editor of Volume 26, Issue 3 for IEEE Computing in Science & Engineering (CiSE) on "Converged Computing: A Best of Both Worlds of High-Performance Computing and Cloud"

Reviewer - reviewing activities

Alexandru Costan: IEEE Transactions on Parallel and Distributed Systems, Future Generation Computer Systems, Concurrency and Computation Practice and Experience, IEEE Transactions on Cloud Computing, Journal of Parallel and Distributed Computing.

Silvina Caino-Lores: IEEE Transactions on Services Computing

10.1.4 Invited talks

François Tessier:

- Talk at the monthly I/O Seminar of Inria Bordeaux about the provisioning of heterogeneous storage resources on supercomputers.

Silvina Caino-Lores:

- *Unified Data Abstractions for Scientific Workflow Composition in the Computing Continuum, Driving Scientific Workflows from the Data Plane* Minisymposium at SIAM PP 2024 (Baltimore, USA).

Jakob Luetzgau:

- Talk at the **Per3S** workshop on Performance and Scalability of Storage Systems in Paris, France

Guillaume Pallez:

- "How to evaluate HPC"; Invited Talk at **CCDSC'24**

10.1.5 Leadership within the scientific community

Gabriel Antoniu:

- Large-wingspan National project management: Coordinator of ExaDoST, one of the 5 targeted projects of the NumPEX PEPR project (started in 2023, budget: 6.2 M€). Coordinator of STEEL, one of the 7 high-priority projects of the CLOUD PEPR project (started in 2023, budget: 2.8 M€).
- ETP4HPC: Since 2019, co-leader of the working group on Programming Environments, lead co-author of the corresponding chapter of the Strategic Research Agenda of ETP4HPC (latest edition published in 2024).
- International lab management: Executive Director of JLESC for Inria since April 2024 (previously Vice Executive Director). JLESC is the Joint Inria-Illinois-ANL-BSC-JSCRRIKEN/AICS Laboratory for Extreme-Scale Computing. Within JLESC, he also serves as a Topic Leader for Data storage, I/O and in situ processing for Inria.
- Bilateral Inria-DFKI project management: French coordinator of the ENGAGE project (2022-2024).
- Team management: Head of the KerData Project-Team (INRIA-INSA Rennes).
- International Associate Team management: Leader of the UNIFY Associate Team with Argonne National Lab (2019–2022).

François Tessier:

- Work package co-leader with Francieli Zanon-Boito (Associate Professor, University of Bordeaux) within the NumPEX ExaDoST project.

Alexandru Costan:

- Work package co-leader with René Schubotz(DFKI) within the **ENGAGE** Inria-DFKI project.
- Work package leader within the **PEPR CLOUD STEEL** project.

10.1.6 Scientific expertise

Guillaume Pallez:

- Member of the Inria Scientific Board

Gabriel Antoniu:

- Evaluator for a Horizon Europe project (HORIZON-CL4-2021-HUMAN-01 call)
- Evaluator for several projects submitted to FFPlus, a European initiative highlighting and promoting the adoption of High-Performance Computing (HPC) by SMEs and start-ups across Europe)

10.1.7 Research administration

François Tessier:

- Member of the Commission on Health, Safety and Working Conditions (now called FSS) within the Inria center of Rennes

Guillaume Pallez:

- Member of the National Commission on Health, Safety and Working Conditions (now called FS)

Gabriel Antoniu:

- Member of the Inria HRS4R Steering Committee (HRS4R: European Human Resources Strategy for Research)

10.2 Teaching - Supervision - Juries

Participants: Gabriel Antoniu, Thomas Bouvier, Silvina Caino Lores, Alexandru Costan, Arthur Jaquard, Théo Jolivel, Julien Monniot, Cédric Prigent, François Tessier, Mathis Valli.

10.2.1 Teaching

Alexandru Costan:

- Bachelor: Software Engineering and Java Programming, 28 hours (lab sessions), L3, INSA Rennes.
- Bachelor: Databases, 68 hours (lectures and lab sessions), L2, INSA Rennes.
- Bachelor: Practical case studies, 24 hours (project), L3, INSA Rennes.
- Master: Big Data Storage and Processing, 28h hours (lectures, lab sessions), M1, INSA Rennes.
- Master: Algorithms for Big Data, 28 hours (lectures, lab sessions), M2, INSA Rennes.
- Master: Big Data Project, 28 hours (project), M2, INSA Rennes.

Gabriel Antoniu:

- Master (Engineering Degree, 5th year): NoSQL and Cloud technologies, 20 hours (lectures), M2 level, ENSAI (*École nationale supérieure de la statistique et de l'analyse de l'information*), Bruz.
- Master: Infrastructures for Big Data, 14 hours (lectures), M1 level, IBD Module, University of Rennes.
- Master: Cloud Computing and Big Data, 14 hours (lectures), M2 level, Cloud Module, MIAGE Master Program, University of Rennes.

François Tessier:

- Bachelor: Computer science discovery, 15 hours (lab sessions), L1 level, DIE Module, ISTIC, University of Rennes.
- Master: Cloud Computing and Big Data, 15 hours (lectures), M2 level, Cloud Module, MIAGE Master Program, University of Rennes.
- Master (Engineering Degree, 4th year): Storage on Clouds, 5 hours (lecture and lab session), M2 level, IMT Atlantique, Rennes.

Silvina Caino-Lores:

- Master: Processing Artificial Intelligence and Machine Learning Workloads at Scale, 9 hours (lectures), M2 level, Big Data Storage and Processing Infrastructures Module, Cloud and Network Infrastructures Master Program, EIT Digital School, ISTIC, University of Rennes.
- Master: Processing Artificial Intelligence and Machine Learning Workloads at Scale, 9 hours (lectures) and 15 hours (tutorials), M1 level, Artificial Intelligence Master Program, ISTIC, University of Rennes.

Thomas Bouvier:

- Master: Stream Processing, 12 hours (lectures and lab sessions), M2 level, INSA Rennes.
- Master: Database optimizations, 30 hours (lab sessions), M1 level, ISTIC, University of Rennes.

Cédric Prigent:

- Master: Cloud Computing and Big Data, 36 hours (lab sessions), M2 level, Cloud Module, MIAGE Master Program, University of Rennes.

Théo Jolivel:

- Master: Cloud Computing and Big Data, 36 hours (lab sessions), M2 level, Cloud Module, MIAGE Master Program, University of Rennes.

Mathis Valli:

- Bachelor: Databases, 12 hours (lab sessions), L3, INSA Rennes.

Arthur Jaquard:

- Master: Processing Artificial Intelligence and Machine Learning Workloads at Scale, 13.5 hours (tutorials), M1 level, Artificial Intelligence Master Program, ISTIC, University of Rennes.

10.2.2 Supervision**PhD:**

- Thomas Bouvier, "Supporting Continual Learning across the Computing Continuum", thesis defended in November 2024, co-advised by Alexandru Costan and Gabriel Antoniu.
- Julien Monnot, "Enabling accurate simulation of HPC storage systems: methodology and practical techniques", thesis defended in December 2024, co-advised by François Tessier and Gabriel Antoniu. "
- Alexis Bandet, thesis defended in December 2024, co-advised by Guillaume Pallez and Francieli Zanon Boito.

PhD in progress:

- Cédric Prigent, "Supporting Online Learning and Inference in Parallel across the Digital Continuum", thesis started in November 2021, co-advised by Alexandru Costan, Gabriel Antoniu and Loïc Cudennec (DGA)
- Mathis Valli, "Comparative Analysis of Federated Learning: Simulations Versus Real-World Testbeds in dynamic settings", thesis started in April 2023, co-advised by Alexandru Costan, Cédric Tedeschi (Myriads) and Loïc Cudennec (DGA).
- Théo Jolivel, "Modeling and Simulation of Exascale Storage Systems", thesis started in October 2024, co-advised by François Tessier, Gabriel Antoniu and Philippe Deniel (CEA).
- Arthur Jaquard, "Dynamic in situ and in transit data analysis for Exascale Computing using Damaris", thesis started in October 2024, co-advised by Gabriel Antoniu, Laurent Colombet (CEA), Silvina Caino-Lores, and Julien Bigot (CEA).
- Robin Boezennec, thesis started in November 2022, co-advised by Guillaume Pallez and Fanny Dufossé (Datamove, Grenoble).

Internships:

- Théo Jolivel, "An Abstraction Layer for I/O Characterization of Large-Scale Applications", 5-month Master 2 internship started in March 2024, co-advised by François Tessier and Guillaume Pallez.
- Hugo Thay, "Analysis of the I/O Access Pattern of a Radio Astronomy Application", 10-week Master 1 internship started in May 2024, supervised by François Tessier.
- Nicolas Vincent, "Computational Storage Programming Paradigms and Execution Engines", 7-month Master 1 project internship started in September 2024, supervised by Jakob Luetzgau.
- Alix Trémondeux, "Refurbished HPC", 5-month post-M2 internship started in September 2024, supervised by Guillaume Pallez.

10.2.3 Juries

Alexandru Costan:

- GDR RSD - Member of the juries for the PhD award and Young/Senior Researcher award

Silvina Caino-Lores:

- PhD thesis: *New techniques to build and manage agnostic workflows for the processing of digital products*, Dante Sanchez-Gallegos, University Carlos III of Madrid, Spain

10.3 Popularization

Participants: Silvina Caino Lores.

10.3.1 Productions (articles, videos, podcasts, serious games, ...)

Silvina Caino-Lores:

- Featured profile in the 2024 edition of the *Women Leaders of the 21st Century* publication by the IT Innovation Foundation (Spain).

11 Scientific production

11.1 Major publications

- [1] F. Boito, G. Pallez, L. Teylo and N. Vidal. 'IO-SETS: Simple and efficient approaches for I/O bandwidth management'. In: *IEEE Transactions on Parallel and Distributed Systems* 34.10 (15th Aug. 2023), pp. 2783–2796. DOI: [10.1109/TPDS.2023.3305028](https://doi.org/10.1109/TPDS.2023.3305028). URL: <https://inria.hal.science/hal-03648225>.
- [2] N. Cherière, M. Dorier and G. Antoniu. 'How fast can one resize a distributed file system?' In: *Journal of Parallel and Distributed Computing* 140 (June 2020), pp. 80–98. DOI: [10.1016/j.jpdc.2020.02.001](https://doi.org/10.1016/j.jpdc.2020.02.001). URL: <https://hal.archives-ouvertes.fr/hal-02961875>.
- [3] M. Dorier, G. Antoniu, F. Cappello, M. Snir, R. Sisneros, O. Yildiz, S. Ibrahim, T. Peterka and L. Orf. 'Damaris: Addressing Performance Variability in Data Management for Post-Petascale Simulations'. In: *ACM Transactions on Parallel Computing* 3.3 (2016), p. 15. DOI: [10.1145/2987371](https://doi.org/10.1145/2987371). URL: <https://hal.inria.fr/hal-01353890> (cit. on p. 11).
- [4] M. Dorier, S. Ibrahim, G. Antoniu and R. Ross. 'Using Formal Grammars to Predict I/O Behaviors in HPC: the Omnisc'IO Approach'. In: *IEEE Transactions on Parallel and Distributed Systems* (2016). DOI: [10.1109/TPDS.2015.2485980](https://doi.org/10.1109/TPDS.2015.2485980). URL: <https://hal.inria.fr/hal-01238103>.
- [5] K. Fauvel, D. Balouek-Thomert, D. Melgar, P. Silva, A. Simonet, G. Antoniu, A. Costan, V. Masson, M. Parashar, I. Rodero and A. Termier. 'A Distributed Multi-Sensor Machine Learning Approach to Earthquake Early Warning'. In: In Proceedings of the 34th AAAI Conference on Artificial Intelligence. New York, United States, 7th Feb. 2020, pp. 403–411. DOI: [10.1609/aaai.v34i01.5376](https://doi.org/10.1609/aaai.v34i01.5376). URL: <https://hal.archives-ouvertes.fr/hal-02373429> (cit. on p. 9).

- [6] J. Luettgau, S. Snyder, T. Reddy, N. Awtrey, K. Harms, J. L. Bez, R. Wang, R. Latham and P. Carns. ‘Enabling Agile Analysis of I/O Performance Data with PyDarshan’. In: SC-W 2023: Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis. Denver CO USA, France: ACM, 12th Nov. 2023, pp. 1380–1391. DOI: [10.1145/3624062.3624207](https://doi.org/10.1145/3624062.3624207). URL: <https://inria.hal.science/hal-04390702>.
- [7] M. Malms, L. Cargemel, E. Suarez, N. Mittenzwey, M. Duranton, S. Sezer, C. Prunty, P. Rossé-Laurent, M. Pérez-Harnandez, M. Marazakis, G. Lonsdale, P. Carpenter, G. Antoniu, S. Narasimhamurthy, A. Brinkman, D. Pleiter, U.-U. Haus, J. Krueger, H.-C. Hoppe, E. Laure, A. Wierse, V. Bartsch, K. Michielsen, C. Allouche, T. Becker and R. Haas. *ETP4HPC’s SRA 5 - Strategic Research Agenda for High-Performance Computing in Europe - 2022*. Zenodo, 2022. DOI: [10.5281/zenodo.7347009](https://doi.org/10.5281/zenodo.7347009). URL: <https://hal.science/hal-03937825> (cit. on p. 3).
- [8] O.-C. Marcu, A. Costan, G. Antoniu, M. S. Pérez-Hernández, B. Nicolae, R. Tudoran and S. Bortoli. ‘KerA: Scalable Data Ingestion for Stream Processing’. In: ICDCS 2018 - 38th IEEE International Conference on Distributed Computing Systems. Vienna, Austria: IEEE, 2nd July 2018, pp. 1480–1485. DOI: [10.1109/ICDCS.2018.00152](https://doi.org/10.1109/ICDCS.2018.00152). URL: <https://hal.inria.fr/hal-01773799>.
- [9] O.-C. Marcu, A. Costan, G. Antoniu and M. S. Pérez-Hernández. ‘Spark versus Flink: Understanding Performance in Big Data Analytics Frameworks’. In: Cluster 2016 - The IEEE 2016 International Conference on Cluster Computing. Taipei, Taiwan, 12th Sept. 2016. DOI: [10.1109/cluster.2016.22](https://doi.org/10.1109/cluster.2016.22). URL: <https://hal.inria.fr/hal-01347638>.
- [10] P. Matri, Y. Alforov, A. Brandon, M. Pérez, A. Costan, G. Antoniu, M. Kuhn, P. Carns and T. Ludwig. ‘Mission Possible: Unify HPC and Big Data Stacks Towards Application-Defined Blobs at the Storage Layer’. In: *Future Generation Computer Systems* 109 (Aug. 2020), pp. 668–677. DOI: [10.1016/j.future.2018.07.035](https://doi.org/10.1016/j.future.2018.07.035). URL: <https://hal.archives-ouvertes.fr/hal-01892682>.
- [11] P. Matri, A. Costan, G. Antoniu, J. Montes and M. S. Pérez. ‘Tyr: Blob Storage Meets Built-In Transactions’. In: IEEE ACM SC16 - The International Conference for High Performance Computing, Networking, Storage and Analysis 2016. Salt Lake City, United States, 13th Nov. 2016. DOI: [10.1109/sc.2016.48](https://doi.org/10.1109/sc.2016.48). URL: <https://hal.inria.fr/hal-01347652>.
- [12] S. Narasimhamurthy, N. Mittenzwey, F. Magugliani, M. Duranton, C. Prunty, P. Rossé-Laurent, M. Marazakis, P. Carpenter, G. Antoniu, S. Neuwirth, P. Deniel, D. Pleiter, U.-U. Haus, E. Laure, A. Wierse, T. Becker, R. Haas, M. Malms, H.-C. Hoppe, V. Bartsch, S. Dolas, O. Vysocky, M. Perez, A. Forrester, K. Michielsen and E. Suarez. *ETP4HPC’s SRA 6 - Strategic Research Agenda for High Performance Computing in Europe*. ETP4HPC Strategic Research Agenda. Zenodo, 2024. DOI: [10.5281/zenodo.14268783](https://doi.org/10.5281/zenodo.14268783). URL: <https://hal.science/hal-04846393>.
- [13] D. Rosendo, A. Costan, G. Antoniu, M. Simonin, J.-C. Lombardo, A. Joly and P. Valduriez. ‘Reproducible Performance Optimization of Complex Applications on the Edge-to-Cloud Continuum’. In: Cluster 2021 - IEEE International Conference on Cluster Computing. Portland, OR, United States, 2021, pp. 23–34. DOI: [10.1109/Cluster48925.2021.00043](https://doi.org/10.1109/Cluster48925.2021.00043). URL: <https://hal.archives-ouvertes.fr/hal-03310540>.
- [14] D. Rosendo, P. Silva, M. Simonin, A. Costan and G. Antoniu. ‘E2Clab: Exploring the Computing Continuum through Repeatable, Replicable and Reproducible Edge-to-Cloud Experiments’. In: Cluster 2020 - IEEE International Conference on Cluster Computing. Kobe, Japan, 14th Sept. 2020, pp. 1–11. DOI: [10.1109/CLUSTER49012.2020.00028](https://doi.org/10.1109/CLUSTER49012.2020.00028). URL: <https://hal.archives-ouvertes.fr/hal-02916032> (cit. on p. 11).
- [15] Y. Taleb, R. Stutsman, G. Antoniu and T. Cortes. ‘Tailwind: Fast and Atomic RDMA-based Replication’. In: ATC ‘18 - USENIX Annual Technical Conference. Boston, United States, 13th July 2018, pp. 850–863. URL: <https://hal.inria.fr/hal-01676502>.

11.2 Publications of the year

International journals

- [16] R. Boëzennec, F. Dufossé and G. Pallez. ‘Qualitatively Analyzing Optimization Objectives in the Design of HPC Resource Manager’. In: *ACM Transactions on Modeling and Performance Evaluation of Computing Systems* (2024), pp. 1–28. URL: <https://hal.science/hal-04187517>. In press.
- [17] T. Bouvier, B. Nicolae, A. Costan, T. Bicer, I. Foster and G. Antoniu. ‘Efficient Distributed Continual Learning for Steering Experiments in Real-Time’. In: *Future Generation Computer Systems* (July 2024), pp. 1–19. DOI: [10.1016/j.future.2024.07.016](https://doi.org/10.1016/j.future.2024.07.016). URL: <https://inria.hal.science/hal-04664176> (cit. on p. 21).
- [18] D. I. Dolci, J. R. Maddison, D. A. Ham, G. Pallez and J. Herrmann. ‘checkpoint_schedules: schedules for incremental checkpointing of adjoint simulations’. In: *Journal of Open Source Software* 9 (22nd Mar. 2024), pp. 1–4. DOI: [10.21105/joss.06148](https://doi.org/10.21105/joss.06148). URL: <https://inria.hal.science/hal-04854312>.
- [19] Y. Du, L. Marchal, G. Pallez and Y. Robert. ‘Improving Batch Schedulers with Node Stealing for Failed Jobs’. In: *Concurrency and Computation: Practice and Experience* 36.12 (2024), pp. 1–36. DOI: [10.1002/cpe.8043](https://doi.org/10.1002/cpe.8043). URL: <https://inria.hal.science/hal-03643403>.
- [20] C. Prigent, A. Costan, G. Antoniu and L. Cudennec. ‘Enabling Federated Learning across the Computing Continuum: Systems, Challenges and Future Directions’. In: *Future Generation Computer Systems* 160 (29th June 2024), pp. 767–783. DOI: [10.1016/j.future.2024.06.043](https://doi.org/10.1016/j.future.2024.06.043). URL: <https://hal.science/hal-04659211> (cit. on p. 20).
- [21] F. Tessier, V. Vishwanath and E. Jeannot. ‘Adding topology and memory awareness in data aggregation algorithms’. In: *Future Generation Computer Systems* 159 (Oct. 2024), pp. 188–203. DOI: [10.1016/j.future.2024.05.016](https://doi.org/10.1016/j.future.2024.05.016). URL: <https://hal.science/hal-04783379> (cit. on p. 23).

International peer-reviewed conferences

- [22] A. Bandet, F. Boito and G. Pallez. ‘Scheduling distributed I/O resources in HPC systems’. In: 30th International European Conference on Parallel and Distributed Computing 26 - 30 August 2024 Madrid, Spain 30th International European Conference on Parallel and Distributed Computing. Madrid, Spain, 26th Aug. 2024. URL: <https://inria.hal.science/hal-04394004> (cit. on p. 25).
- [23] R. Boëzennec, D. Carastan-Santos, F. Dufossé and G. Pallez. ‘Allocation Strategies for Disaggregated Memory in HPC Systems’. In: HiPC 2024 - 31st IEEE International Conference on High Performance Computing, Data, and Analytics. Bangalore, India: IEEE, 2024, pp. 1–11. URL: <https://hal.science/hal-04815672>.
- [24] T. Bouvier, B. Nicolae, H. Chaugier, A. Costan, I. Foster and G. Antoniu. ‘Efficient Data-Parallel Continual Learning with Asynchronous Distributed Rehearsal Buffers’. In: CCGrid 2024 - IEEE 24th International Symposium on Cluster, Cloud and Internet Computing. Philadelphia (PA), United States, 2024, pp. 1–10. DOI: [10.1109/CCGrid59990.2024.00036](https://doi.org/10.1109/CCGrid59990.2024.00036). URL: <https://inria.hal.science/hal-04600107> (cit. on p. 21).
- [25] S. Caino-Lores, D. Claudino, E. Dumitrescu, T. S. Humble, S. L. Alarcon and E. Wong. ‘Rethinking Programming Paradigms in the QC-HPC Context’. In: WAMTA 2024 - 2nd International Workshop on Asynchronous Many-Task Systems and Applications. Vol. 14626. Lecture Notes in Computer Science. Knoxville, TN, United States: Springer Nature Switzerland, 30th May 2024, pp. 84–91. DOI: [10.1007/978-3-031-61763-8_8](https://doi.org/10.1007/978-3-031-61763-8_8). URL: <https://hal.science/hal-04902105> (cit. on p. 19).

- [26] T. Jolivel, F. Tessier, J. Monniot and G. Pallez. ‘MOSAIC: Detection and Categorization of I/O Patterns in HPC Applications’. In: *SC24-W: Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis*. PDSW 2024 - 9th International Parallel Data Systems Workshop. Atlanta, United States, 2024, pp. 1–7. DOI: [10.1109/SCW63240.2024.00172](https://doi.org/10.1109/SCW63240.2024.00172). URL: <https://hal.science/hal-04808300> (cit. on p. 22).
- [27] J. Monniot, F. Tessier, H. Casanova and G. Antoniu. ‘Simulation of Large-Scale HPC Storage Systems: Challenges and Methodologies’. In: *HiPC 2024 - 31st IEEE International Conference on High Performance Computing, Data, and Analytics*. Bangalore, India, 2024, pp. 1–11. URL: <https://inria.hal.science/hal-04784808> (cit. on pp. 12, 18).
- [28] C. Prigent, M. Chelli, A. Costan, L. Cudennec, R. Schubotz and G. Antoniu. ‘Efficient Resource-Constrained Federated Learning Clustering with Local Data Compression on the Edge-to-Cloud Continuum’. In: *HiPC 2024 - 31st IEEE International Conference on High Performance Computing, Data, and Analytics*. Bengaluru (Bangalore), India, 2024, pp. 1–11. URL: <https://hal.science/hal-04779813> (cit. on pp. 12, 19, 20).
- [29] R. Souza, S. Caino-Lores, M. Coletti, T. J. Skluzacek, A. Costan, F. Suter, M. Mattoso and R. F. da Silva. ‘Workflow Provenance in the Computing Continuum for Responsible, Trustworthy, and Energy-Efficient AI’. In: *e-Science 2024 - 20th IEEE International Conference on e-Science*. Osaka, Japan: IEEE, 16th Sept. 2024, pp. 1–7. DOI: [10.1109/e-Science62913.2024.10678731](https://doi.org/10.1109/e-Science62913.2024.10678731). URL: <https://hal.science/hal-04902079> (cit. on p. 20).
- [30] A. Tarraf, A. Bandet, F. Zanon Boito, G. Pallez and F. Wolf. ‘Capturing Periodic I/O Using Frequency Techniques’. In: *IPDPS 2024 - 38th IEEE International Parallel & Distributed Processing Symposium*. San Francisco, United States, 2024, pp. 1–13. URL: <https://inria.hal.science/hal-04382142>.
- [31] M. Valli, A. Costan, C. Tedeschi and L. Cudennec. ‘Towards Efficient Learning on the Computing Continuum: Advancing Dynamic Adaptation of Federated Learning’. In: *FlexScience’24: Proceedings of the 14th Workshop on AI and Scientific Computing at Scale using Flexible Computing Infrastructures*. FlexScience 2024 - 14th Workshop on AI and Scientific Computing at Scale using Flexible Computing Infrastructures. Pisa, Italy: ACM, 2024, pp. 42–49. DOI: [10.1145/3659995.3660042](https://doi.org/10.1145/3659995.3660042). URL: <https://hal.science/hal-04698619> (cit. on p. 22).

Scientific books

- [32] S. Narasimhamurthy, N. Mittenzwey, F. Magugliani, M. Durantou, C. Prunty, P. Rossé-Laurent, M. Marazakis, P. Carpenter, G. Antoniu, S. Neuwirth, P. Deniel, D. Pleiter, U.-U. Haus, E. Laure, A. Wierse, T. Becker, R. Haas, M. Malms, H.-C. Hoppe, V. Bartsch, S. Dolas, O. Vysocky, M. Perez, A. Forrester, K. Michielsen and E. Suarez. *ETP4HPC’s SRA 6 - Strategic Research Agenda for High Performance Computing in Europe*. ETP4HPC Strategic Research Agenda. Zenodo, 2024. DOI: [10.5281/zenodo.14268783](https://doi.org/10.5281/zenodo.14268783). URL: <https://hal.science/hal-04846393>.

Reports & preprints

- [33] A. Bandet, F. Boito and G. Pallez. *Prediction and Interpretability of HPC I/O Resources Usage with Machine Learning*. 2024. URL: <https://inria.hal.science/hal-04698511>.
- [34] A. Bandet, F. Zanon Boito and G. Pallez. *Allocation and Placement Algorithms for Scheduling Distributed I/O Resources in HPC Systems*. RR-9549. Inria Bordeaux; Inria Rennes, 28th May 2024, pp. 1–27. URL: <https://hal.science/hal-04593977>.
- [35] C. Barthélemy, F. Zanon Boito, E. Jeannot, G. Pallez and L. Teylo. *Implementation of an unbalanced I/O Bandwidth Management system in a Parallel File System*. RR-9537. Inria, Jan. 2024. URL: <https://inria.hal.science/hal-04417412>.

- [36] R. Ferreira da Silva, D. Bard, K. Chard, d. W. Shaun, I. Foster, T. Gibbs, C. Goble, W. Godoy, J. Gustafsson, U.-U. Haus et al. *Workflows Community Summit 2024: Future Trends and Challenges in Scientific Workflows*. Oak Ridge National Laboratory, USA, 18th Oct. 2024. DOI: [10.5281/zenodo.13844758](https://doi.org/10.5281/zenodo.13844758). URL: <https://hal.science/hal-04902062> (cit. on p. 19).
- [37] F. Zanon Boito, L. Teylo, M. Popov, T. Jolivel, F. Tessier, J. Luetzgau, J. Monniot, A. Tarraf, A. Carneiro and C. Osthoff. *A Deep Look Into the Temporal I/O Behavior of HPC Applications*. 15th Jan. 2025. URL: <https://inria.hal.science/hal-04887809> (cit. on p. 22).

Other scientific publications

- [38] P. Carpenter, G. Antoniu, M. Arenaz, O. Aumage, J. Beránek, A. Buttari, A. Costan, S. Happ, V. Kannan, C. Perez, A. Peña, A. Scionti, X. Vigouroux and P. Viviani. *ETP4HPC SRA White Paper - Programming Environment*. Dec. 2024. DOI: [10.5281/zenodo.14446622](https://doi.org/10.5281/zenodo.14446622). URL: <https://hal.science/hal-04905035>.
- [39] S. Neuwirth, P. Deniel, J.-T. Acquaviva, M. Golasowski, M. Hennecke, A. Jackson, T. Leibovici, J. Luetzgau and R. Nou. *ETP4HPC SRA 6 White Paper - I/O and Storage*. Jan. 2025. DOI: [10.5281/zenodo.14605692](https://doi.org/10.5281/zenodo.14605692). URL: <https://hal.science/hal-04905038>.

11.3 Cited publications

- [40] G. Antoniu, P. Valduriez, H.-C. Hoppe and J. KrÄEger. *Towards Integrated Hardware/Software Ecosystems for the Edge-Cloud-HPC Continuum*. ETP4HPC White Papers. ETP4HPC: European Technology Platform for High Performance Computing, 2021. DOI: [10.5281/zenodo.5534464](https://doi.org/10.5281/zenodo.5534464). URL: <https://hal.archives-ouvertes.fr/hal-03358930> (cit. on p. 5).
- [41] T. Beck, A. Baroni, R. Bennink, G. Buchs, E. A. C. Pérez, M. Eisenbach, R. F. da Silva, M. G. Meena, K. Gottiparthi, P. Groszkowski et al. ‘Integrating quantum computing resources into scientific HPC ecosystems’. In: *Future Generation Computer Systems* 161 (2024), pp. 11–25 (cit. on p. 19).
- [42] R. Bolze, F. Cappello, E. Caron, M. Dayde, F. Desprez, E. Jeannot, Y. Jégou, S. Lanteri, J. Leduc, N. Melab, G. Mornet, R. Namyst, P. Primet, B. Quétier, O. Richard, E.-G. Talbi and I. Touche. ‘Grid’5000: A Large Scale And Highly Reconfigurable Experimental Grid Testbed’. In: *International Journal of High Performance Computing Applications* 20.4 (2006), pp. 481–494. DOI: [10.1177/1094342006070078](https://doi.org/10.1177/1094342006070078). URL: <https://hal.inria.fr/hal-00684943> (cit. on p. 11).
- [43] H. Casanova, R. Ferreira da Silva, R. Tanaka, S. Pandey, G. Jethwani, W. Koch, S. Albrecht, J. Oeth and F. Suter. ‘Developing Accurate and Scalable Simulators of Production Workflow Management Systems with WRENCH’. In: *Future Generation Computer Systems* 112 (2020), pp. 162–175. DOI: [10.1016/j.future.2020.05.030](https://doi.org/10.1016/j.future.2020.05.030) (cit. on p. 18).
- [44] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio et al. ‘Variational quantum algorithms’. In: *Nature Reviews Physics* 3.9 (2021), pp. 625–644 (cit. on p. 19).
- [45] A. Joly, P. Bonnet, H. Goëau, J. Barbe, S. Selmi, J. Champ, S. Dufour-Kowalski, A. Affouard, J. Carré, J.-F. o. Molino, N. Boujemaa and D. Barthélémy. ‘A look inside the Pl@ntNet experience’. In: *Multimedia Systems* 22.6 (2016), pp. 751–766. DOI: [10.1007/s00530-015-0462-9](https://doi.org/10.1007/s00530-015-0462-9). URL: <https://hal.inria.fr/hal-01182775> (cit. on p. 10).
- [46] M. Malms, L. Cargemel, E. Suarez, N. Mittenzwey, M. Duranton, S. Sezer, C. Prunty, P. Rossé-Laurent, M. Pérez-Harnandez, M. Marazakis, G. Lonsdale, P. Carpenter, G. Antoniu, S. Narasimharmurthy, A. Brinkman, D. Pleiter, U.-U. Haus, J. Krueger, H.-C. Hoppe, E. Laure, A. Wierse, V. Bartsch, K. Michielsen, C. Allouche, T. Becker and R. Haas. *ETP4HPC’s SRA 5 - Strategic Research Agenda for High-Performance Computing in Europe - 2022*. Zenodo, 2022. DOI: [10.5281/zenodo.7347009](https://doi.org/10.5281/zenodo.7347009). URL: <https://hal.science/hal-03937825> (cit. on p. 5).

- [47] M. Malms, M. Ostasz, M. Gilliot, P. Bernier-Bruna, L. Cargemel, E. Suarez, H. Cornelius, M. Durantou, B. Koren, P. Rosse-Laurent, M. S. Párez-Hernández, M. Marazakis, G. Lonsdale, P. Carpenter, G. Antoniu, S. Narasimhamurthy, A. Brinkman, D. Pleiter, A. Tate, J. Krueger, H.-C. Hoppe, E. Laure and A. Wierse. *ETP4HPC's Strategic Research Agenda for High-Performance Computing in Europe 4*. Ed. by w. t. s. o. t. E.-2. p. ETP4HPC: European Technology Platform for High Performance Computing. ETP4HPC White Papers. 2020. DOI: [10.5281/zenodo.4605343](https://doi.org/10.5281/zenodo.4605343). URL: <https://hal.inria.fr/hal-03354396> (cit. on p. 5).
- [48] J. Monniot, F. Tessier and G. Antoniu. *Modeling Allocation of Heterogeneous Storage Resources on HPC Systems*. SC 2022 - International Conference for High Performance Computing, Networking, Storage, and Analysis (Posters). Poster. Nov. 2022. URL: <https://hal.science/hal-03878252> (cit. on p. 18).
- [49] J. Monniot, F. Tessier, M. Robert and G. Antoniu. 'StorAlloc: A Simulator for Job Scheduling on Heterogeneous Storage Resources'. In: *HeteroPar 2022*. Glasgow, United Kingdom, Aug. 2022. URL: <https://inria.hal.science/hal-03683568> (cit. on p. 18).
- [50] J. Monniot, F. Tessier, M. Robert and G. Antoniu. 'Supporting Dynamic Allocation of Heterogeneous Storage Resources on HPC Systems'. In: *Concurrency and Computation: Practice and Experience*. Special Issue:S2 World 2020. CUG 2021 & 2022. PN_HCP. HeteroPar 2022 35.28 (Aug. 2023), pp. 1–16. DOI: [10.1002/cpe.7890](https://doi.org/10.1002/cpe.7890). URL: <https://hal.science/hal-03922866> (cit. on p. 18).
- [51] *The TransContinuum Initiative vision paper*. 2020. URL: <https://www.etp4hpc.eu/tci-vision.html> (cit. on p. 5).
- [52] B. Weder, J. Barzen, F. Leymann and D. Vietz. 'Quantum software development lifecycle'. In: *Quantum Software Engineering*. Springer, 2022, pp. 61–83 (cit. on p. 19).
- [53] *Chameleon Cloud*. 2021. URL: <https://www.chameleoncloud.org/> (cit. on p. 4).
- [54] *Cybeletech - Digital technologies for the plant world*. 2021. URL: <https://www.cybeletech.com/en/home/> (cit. on p. 10).
- [55] *ECLAT - Extreme Computing Lab for Astronomical Telescopes*. 2024. URL: <https://eclat-lab.fr/> (cit. on p. 9).
- [56] *ECMWF - European Centre for Medium-Range Weather Forecasts*. 2021. URL: <https://www.ecmwf.int/> (cit. on pp. 9, 10).
- [57] *European Exascale Software Initiative*. 2013. URL: <http://www.eesi-project.eu/> (cit. on p. 4).
- [58] *The European Technology Platform for High-Performance Computing*. 2012. URL: <http://www.etp4hpc.eu/> (cit. on p. 4).
- [59] *International Exascale Software Program*. 2011. URL: <http://www.exascale.org/iesp/> (cit. on p. 4).
- [60] *SKA - Square Kilometre Array*. 2024. URL: <https://www.skao.int/en> (cit. on p. 8).