UMR IRISA

Activity Report 2025

Team KERDATA

Scalable Storage for Clouds and Beyond

Joint team with Centre Inria de l'Université de Rennes

D1 – Large Scale an Secure Systems

Université de Rennes | rennes | INSA INSTITUT NATIONAL DES SCIENCES APPLIQUÉES RENNES | cnrs | Inria | Inserm | CentraleSupélec | IMT Atlantique Bretagne-Pays de la Loire École Mines-Télécom | Université Bretagne Sud ubs

Project-Team KERDATA

Creation of the Project-Team: 2025 January 01

Each year, Inria research teams publish an Activity Report presenting their work and results over the reporting period. These reports follow a common structure, with some optional sections depending on the specific team. They typically begin by outlining the overall objectives and research programme, including the main research themes, goals, and methodological approaches. They also describe the application domains targeted by the team, highlighting the scientific or societal contexts in which their work is situated. The reports then present the highlights of the year, covering major scientific achievements, software developments, or teaching contributions. When relevant, they include sections on software, platforms, and open data, detailing the tools developed and how they are shared. A substantial part is dedicated to new results, where scientific contributions are described in detail, often with subsections specifying participants and associated keywords. Finally, the Activity Report addresses funding, contracts, partnerships, and collaborations at various levels, from industrial agreements to international cooperations. It also covers dissemination and teaching activities, such as participation in scientific events, outreach, and supervision. The document concludes with a presentation of scientific production, including major publications and those produced during the year.

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. – HPC for machine learning

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

Contents

Project-Team KERDATA **1**

1 Team members, visitors, external collaborators **6**

2 Overall objectives **7**

2.1 Context: the emergence of the Edge-Cloud-HPC Continuum. 7

2.2 Application/workflow-level challenges 7

2.3 Middleware-level challenges 8

2.4 Resource management challenges 8

2.5 Approach, methodology, platforms 8

2.6 Collaboration strategy 8

2.7 Alignment with institutional, national and European strategies 9

3 Research program **9**

3.1 Axis 1: Supporting Data-Centric Applications and Workflows Running Across the Computing Continuum 10

3.2 Axis 2: Data-Aware Middleware Approaches for the Computing Continuum 12

3.3 Axis 3: Sustainable Resource Management for the Computing Continuum 15

4 Application domains **17**

4.1 Radio astronomy 17

4.2 Nuclear Fusion 18

4.3 Material Science 18

5 Social and environmental responsibility **18**

5.1 Footprint of research activities 18

5.2 Impact of research results 18

6 Highlights of the year **19**

7 Latest software developments, platforms, open data **19**

7.1 Latest software developments 19

7.1.1 Damaris 19

7.1.2 E2Clab 20

7.1.3 Fives 21

7.1.4 MOSAIC 21

7.1.5 FLAdversary 22

7.1.6 FLDrift 23

7.2 Open data 23

7.2.1 I/O Traces 23

8 New results **23**
8.1 Supporting Data-Centric Applications and Workflows Running Across the Computing Continuum . . . 23
8.1.1 On the Reproducibility Challenges of Federated Learning: Investigating the Gap between Simulation, Emulation and Real-World Deployments . . . 23
8.1.2 Evaluating Federated Learning Workflows Beyond Simulation: A Deployment-Aware Methodology . . . 24
8.1.3 Supporting SKA data processing workflows with the E2CLab approach to workflow lifecycle management across the continuum . . . 25
8.1.4 Methodology for Automated IoT Experimentation in Controlled Testbeds Prior to Real-World Deployments . . . 25
8.2 Data-Aware Middleware Approaches for the Computing Continuum . . . 25
8.2.1 Multi-level analysis of the I/O pattern of HPC applications . . . 25
8.2.2 Study of I/O interference between jobs . . . 26
8.2.3 Enabling Efficient Runtime Data Analysis to a Crystal Deformation Simulation . . . 27
8.3 Sustainable Resource Management for the Computing Continuum . . . 27
8.3.1 Result-Scalability: Following the Evolution of Selected Social Impact of HPC. . . . 27
8.3.2 Increasing the Lifetime of HPC Machines: Issues, Implications, and Open Challenges 28
8.3.3 Improving Supercomputer Usage with Aging Awareness. . . . 28
8.3.4 Priority-BF: a Task Manager for Priority-Based Scheduling . . . 28
8.3.5 Scheduling multiple task-based applications on distributed heterogeneous computing nodes . . . 29
8.4 Methodological study over the practice of HPC Research . . . 29
8.4.1 Implementing a Reproducibility Initiative in HPC: Experiences from SC24. . . . 29
8.4.2 An In-depth Study of LLM Contributions to the Bin Packing Problem . . . 29

9 Partnerships and cooperations **30**
9.1 International initiatives . . . 30
9.1.1 Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program . . . 30
9.2 International research visitors . . . 31
9.2.1 Visits of international scientists . . . 31
9.2.2 Visits to international teams . . . 31
9.3 European initiatives . . . 31
9.3.1 H2020 projects . . . 31
9.3.2 Collaborations with Major European Organizations . . . 33
9.4 National initiatives . . . 33

10 Dissemination **35**
10.1 Promoting scientific activities . . . 35
10.1.1 Scientific events: organisation . . . 35
10.1.2 Scientific events: selection . . . 37
10.1.3 Journal . . . 37
10.1.4 Invited talks . . . 37
10.1.5 Leadership within the scientific community . . . 38

10.1.6 Scientific expertise 38
10.1.7 Research administration 38
10.2 Teaching - Supervision - Juries - Educational and pedagogical outreach 39
10.2.1 Teaching 39
10.2.2 Supervision 39
10.2.3 Juries 40

11 Scientific production **40**
11.1 Major publications 40
11.2 Publications of the year 41
11.3 Cited publications 43

1 Team members, visitors, external collaborators

Research Scientists

- Gabriel Antoniu [Team leader, INRIA, Senior Researcher, HDR]
- Silvina Caino Lores [INRIA, ISFP]
- Jakob Luettgau [INRIA, Researcher, from Oct 2025]
- Jakob Luettgau [INRIA, Starting Research Position, until Sep 2025]
- Guillaume Pallez [INRIA, Researcher, HDR]
- François Tessier [INRIA, ISFP]

Faculty Member

- Alexandru Costan [INSA RENNES, Associate Professor, until Sep 2025, HDR]

PhD Students

- Robin Boezennec [INRIA]
- Arthur Jaquard [INRIA]
- Theo Jolivel [INRIA]
- Cedric Prigent [INRIA, until Feb 2025]
- Simon Renard [INRIA, from Oct 2025]
- Alix Tremodeux [UNIV RENNES, from Sep 2025]
- Mathis Valli [INRIA]

Technical Staff

- Thomas Badts [INRIA, Engineer]
- Julien Monniot [INRIA, Engineer, until May 2025]
- Jean Etienne Ndamlabin Mboula [INRIA, Engineer]

Interns and Apprentices

- Remy Chiv [INRIA, Intern, from May 2025 until Oct 2025]
- Alix Tremodeux [ENS DE LYON, Intern, until Feb 2025]

Administrative Assistants

- Laurence Dinh [INRIA]
- Armelle Mozziconacci [CNRS]
- Gunther Tessier [INRIA]

Visiting Scientists

- Elias Del Pozo Punal [UNIV CARLOS III, from Mar 2025 until Jul 2025]
- Tomasz Kanas [UNIV VARSOVIE, from Sep 2025 until Nov 2025]

2 Overall objectives

2.1 Context: the emergence of the Edge-Cloud-HPC Continuum.

As witnessed in industry and science and highlighted in strategic documents such as the European ETP4HPC Strategic Research Agenda [91], there is a clear trend to combine numerical computations, large-scale data analytics and AI techniques to improve the results and efficiency of traditional HPC applications, and to advance new applications in fields such as autonomous vehicles, digital twins, smart buildings/towns, etc. A typical scenario consists in Edge devices creating streams of input data, which are processed by data analytics and machine learning applications in the Cloud; alternatively (or in parallel!) they can feed simulations on large, specialised HPC systems, to provide insights and help for prediction of some future system state. Such emerging applications typically need to be implemented as complex workflows and require the coordinated use of supercomputers, Cloud data centres and Edge-processing devices. This assembly is called the Computing Continuum (CC). It raises challenges at multiple levels: at the application/workflow level, to bridge simulations, machine learning and data-driven analytics; at the middleware level, adequate tools must enable efficient deployment and orchestration of the workflow components across the whole distributed infrastructure; and, finally, a capable resource management system must allocate a suitable set of components of the infrastructure to run the application workflow, preferably in a dynamic and adaptive way, taking into account the specific capabilities of each component of the underlying heterogeneous infrastructure.

While each level exhibits specific associated challenges, there are also common, cross-layer concerns, among which we specifically highlight two. The first cross-layer concern regards *sustainability*, understood as an optimization goal encompassing energy efficiency and the reduction of the environmental impact. The second cross-layer concern is related to the rapid development of AI-related workflows, which creates specific needs at multiple levels.

Our objective: *Enable the Data Continuum*. Our research project aims to address some open challenges at each of the aforementioned three levels, while considering the two aforementioned transverse concerns. We specifically focus on *data-related* challenges posed by the requirements (storage, processing, analytics) of complex workflows executed on the Edge-Cloud-HPC continuum and propose innovative algorithms and software architecture solutions towards a Data Continuum.

2.2 Application/workflow-level challenges

In the current state, multitudes of software development stacks are tailored to specific use cases, with no guarantee of interoperability between them. This greatly impedes application software development for integrated CC use cases. Moreover, specific software stacks have been developed for HPC (e.g., based on optimized MPI libraries able to leverage high-end network interconnects), data analytics (e.g., based on Spark, designed for commodity clusters available in cloud datacenters) and AI (e.g., TensorFlow or PyTorch), with different requirements for their initial target execution infrastructures. Components based on such software stacks cannot be integrated efficiently together to support CC workflows, as their assumptions about the underlying infrastructure are different. Programming a complex, hybrid workflow at the highest level requires the ability to consistently combine such workflow components in a unified framework. This requires flexible programming models and supporting environments, which also safeguard performance and energy efficiency. Composability (the ability to combine multiple programming models or software stacks for a single application with defined rules) and reproducibility of workflow execution will be very valuable in this context.

2.3 Middleware-level challenges

Similarly, compatibility and interoperability across all parts of a CC infrastructure must be assured; in particular, this includes data formats, storage abstractions, communication, data processing and data analysis paradigms. It first requires a deep understanding of the I/O behaviour of the distributed workflows. As an illustrative example, upcoming Exascale HPC workflows deployed on supercomputers as part of the continuum will continue to highlight the lack of infrastructures and methodologies to store and analyze the huge results of running simulations - should this storage or analysis be performed on HPC systems or on cloud-based infrastructures. This can limit the scalability potential and lead to sub-optimal usage of the computing infrastructures. As in some cases storing all data (originated from sensors or generated by simulations) may be infeasible, thus new scalable approaches are needed. The goal is to enable processing and analysis of such massive outputs of data on various parts of the continuum infrastructure during and after the HPC simulations through asynchronous I/O and in-situ or in-transit processing inside or outside the HPC system, thus avoiding storage.

2.4 Resource management challenges

Large-scale heterogeneity must be managed in an effective and efficient way. This again cuts across compute, storage and communication systems, and the scheduling/orchestration has to optimize the mapping of workflows onto the CC resources with regard to performance and energy use. A challenge here is to enable the design of adequate data storage architectures coping in particular with capacity-related or energy-related constraints that may diversely concern certain parts of the continuum (Edge, but also energy-bound supercomputers at the post-Exascale age, where sustainability is a primary consideration).

2.5 Approach, methodology, platforms

KerData's global approach consists in modelling, designing, implementing and evaluating distributed algorithms and software architectures to address some of the data-related challenges described above. A specific description of the research questions we address is provided in the next section. We will generally focus on hybrid infrastructures (Edge/Cloud/supercomputers), although some of our research may not span across the complete spectrum of the continuum.

Our research balances theoretical modelling (thanks to the recent arrival of Guillaume Pallez) with a predominantly experimental validation methodology (traditionally carried out by most team members as part of the former KerData team). Overall, to validate our proposed algorithms and architectures, we build software prototypes, then validate them at large scale on real testbeds and experimental platforms.

We will strongly rely on the Grid'5000/SLICES FR platform. Moreover, thanks to our projects and partnerships (in particular in EuroHPC projets building pre-Exascale platforms, such as ACROSS and EUPEX), we have access to reference supercomputer testbeds, such as the Karolina[1] and Irene (CEA). More importantly, the team is leading Exa-DoST (2023-2029), the project of the NumPEx program focused on data-related challenges for the Exascale, as part of a national effort to design and build the software infrastructure for Jules Verne, the first Exascale machine to be installed in France. All these are excellent opportunities to validate our results on advanced realistic platforms.

2.6 Collaboration strategy

We chose to work in close collaboration with some of the leading international teams in the areas of data management for Edge, Clouds and HPC systems in Academia. As an example, we have been building and maintaining a long-term, privileged partnership with Argonne National Laboratory (USA), a top player in USA HPC research field, through a series of Associate Team projects (Data@Exascale, UNIFY, UNIFY 2) in the framework of the JLESC international laboratory. More recently we initiated collaborations including Oak Ridge National Laboratory (USA) - where the most powerful supercomputer available today (Frontier)

[1] www.it4i.cz/en/infrastructure/karolina

is running; we also collaborate with DFKI (Germany), a strategic Inria partner in the AI area. In industry, formal collaborations are currently in place with ATOS/Eviden, a strategic HPC stakeholder in France and DataDirectNetworks (DDN), a major storage company, in the context of national (PEPR) and European collaborative projects.

2.7 Alignment with institutional, national and European strategies

Data-intensive applications exhibit several common requirements with respect to the need for data storage and I/O processing at very large scales, to support complex workflows combining scientific simulation and data analytics. While our past activity was already aligned with Inria's strategic objectives [62], which acknowledged HPC-Big Data convergence as one of the priorities of our institute, our project for the future goes beyond. It explicitly leverages the challenges identified in the latest edition of the ETP4HPC agenda [91], which highlights the evolution of HPC from a traditional supercomputer-centric vision to an enlarged vision where complex workflows are distributed across interconnected supercomputers, Clouds and Edge infrastructures. Our research program is addressing some of these challenges. In addition, at the national level, our team is leading two strategic PEPR projects whose respective scientific programs have been defined based on this continuum-aware vision. The first one is the Exa-DoST project (Exascale Data-Oriented Software and Tools), a 6.2 M€ project within the NumPEx PEPR program (2023-2029), which aims to provide the software infrastructure for the future Exascale supercomputer expected to be installed in France in 2025 (Jules Vernes). The second one is STEEL, a 2.8 M€ project (Secure and efficient daTa storagE and procEssing on cLoud-based infrastructures) within the CLOUD PEPR program (2023-2030). These projects (defined for 7+ years) are structuring many of our long-term activities.

In addition, some of our concrete collaborative projects involve 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 ATOS/Eviden, through the ACROSS and EUPEX H2020 EuroHPC projects.

3 Research program

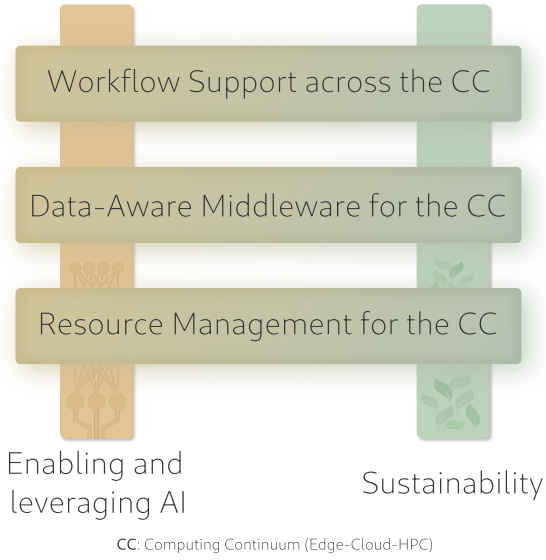


Figure 1: Overview of the research program.

The emergence of the Computing Continuum raises challenges at multiple levels: at the application/workflow level, at the middleware level and at the resource management level. We structured our research program accordingly, in three axes.

The first axis covers workflow-level/application-level research directions. It addresses questions like: how to enable workflow composition across the continuum? How to ensure the reproducibility of workflow execution? How to leverage different sources of metadata to establish a provenance chain (i.e., a record trail of the overall state of the application and its intermediate results) that builds trust on the workflow's results? How could data models support data volume and transfer reduction as a step towards resource sustainability of applications in the continuum? It also includes some more specific research directions related to the execution of distributed AI workflows across the Computing Continuum (involving parallel learning and federated learning).

The second axis addresses research challenges related to middleware-level data management across the continuum, where workflows combining simulation, data analytics and AI are being deployed. In particular, this axis plans to cover topics such as I/O behaviour characterization, storage-centric hybrid infrastructure convergence, and data interoperability across hybrid HPC/Cloud/Edge infrastructures. It also addresses the question: how to perform in-situ data analysis for post-Exascale workflows processing continuous data flows, while considering both performance and energy efficiency?

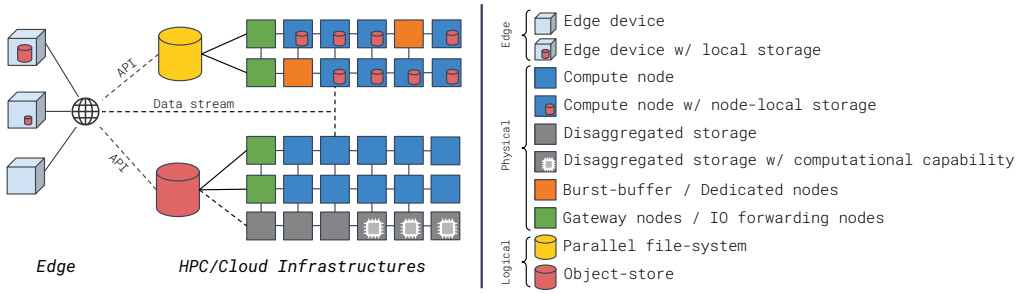


Figure 2: Storage heterogeneity across the Computing Continuum.

The third (lower-level) axis focuses on resource management, with a strong focus on storage resources, but not only. It addresses questions such as: how to provision heterogeneous storage resources across hybrid HPC/Cloud/Edge infrastructures (Figure 2)? What would be a frugal data storage architecture enabling the transition to post-Exascale workflows? How to leverage emerging storage approaches such as disaggregated storage (i.e. a set of storage units physically separated from the compute units) and computational storage (i.e. storage units augmented with some limited integrated computational capabilities)? Finally: how can resource managers and HPC transform/evolve to better adapt to climate change?

We identified two transverse (vertical) themes that are present in some of the research topics of the three (horizontal) axes: *artificial intelligence* (as a target type of workflow to be supported, but also as an enabling technique) and *sustainability* (including aspects related to energy efficiency, frugality and adaptation to emerging applications and hardware technologies in response to climate change).

3.1 Axis 1: Supporting Data-Centric Applications and Workflows Running Across the Computing Continuum

Today, there is a need to efficiently integrate simulations, data analytics and learning, which requires interoperable solutions for data processing [91]. As an example, upcoming large-scale scientific experiments like the Square Kilometer Array (SKA) [2] are expected to process raw data in the order of an exabyte per day [3]. Processing these data volumes requires complex scientific workflows able to extract knowledge and produce insight at every stage: from the instruments and devices producing data that needs to be reduced and pre-processed in situ, to the service-oriented visualization and exploration dashboards that need to be customized for the use case of the domain scientist. Existing works on workflow composition and deployment in the continuum focus on task-flow control and are disconnected from data patterns and structures beyond domain-specific applications [40, 32]. Moreover, general approaches for representing knowledge and provenance in the form of metadata are also lacking for such workflows.

[2] www.skao.int/en
[3] eas.unige.ch/EAS2021/session.jsp?id=SS5

As different communities leverage the Computing Continuum, they express the need to make their research verifiable by others. This is exacerbated by the pervasive usage of AI, as there is increasing awareness about potential ethical and practical implications [105]. The explainability (i.e., making AI's decision-making process understandable) and transparency of AI (i.e., ensuring clarity in AI's design, data and operation) are particularly concerning [31]. Advancing explainability and transparency in AI is currently an essential priority for responsible and trustworthy AI-powered applications. This requires advances in repeatability, replicability, and reproducibility (3R's) accross the Computing Continuum [41, 109].

As AI-oriented workflows tend to gain an increasing share, it becomes important to address the performance and scalability of machine learning (ML) distributed algorithms executed across the Computing Continuum. Methods like deep learning (DL) and federated learning (FL) leverage different technologies to produce insight from large volumes of data. Despite increasing convergence between DL and HPC [77, 44], the training of DL models remains time-consuming and resource-intensive. In FL, powerful facilities (Cloud or HPC) are used to train a global model, while the local, personalized training is typically done close to the data production sites on less powerful computational resources (Edge). This yields the challenge of managing heterogeneity (e.g., differences in computation capacity, network latency and node volatility) as well as variability in data distributions among clients, while respecting the privacy requirements of potentially malicious devices.

In summary, this axis covers research directions to support the composition of scalable and reproducible workflows comprising diverse applications –from simulations to AI–, while addressing the challenges of an heterogeneous environment.

Data-Centric Workflow Composition in the Computing Continuum. The scientific community has reached consensus that common interfaces for data management in the continuum are necessary [71]. Unified data abstractions can enable the interoperability of data storage and processing across the continuum and facilitate data analytics at all levels [53], alleviating the disconnect between application- and storage-oriented approaches to interoperability. However, no unified data modeling approaches exist for structuring and representing data on a logical level across the computing continuum.

The first steps in this research direction involve establishing what are the essential attributes needed to represent data in the different programming models coexisting in the continuum (e.g., ML models, simulation data, annotations resulting from analysis). We will systematically categorize these attributes to deliver data abstractions and models that can be specialized for different tasks. In addition, we will investigate how to embed metadata in these abstractions so that future work can explore new ways of describing, processing, and tracking data at the workflow level, which aligns with the topics of workflow instrumentation and reproducibility. On a longer term, we will also study how these data models could support data volume and transfer reduction as a step towards resource sustainability of applications in the continuum.

Enabling Reproducibility and Trustworthiness in Complex Workflows Across the Computing Continuum. Current approaches to support workflow reproducibility are based on workflow modelling [104] or simulation [27, 110]. These approaches raise some important challenges in terms of specification, modelling, and validation to support reproducibility in the Computing Continuum. For example, it is increasingly difficult to model the heterogeneity and volatility of Edge devices or to assess the impact of the inherent complexity of hybrid Edge-Cloud deployments on performance. With the rise of AI workflows, the issue of reproducibility is aggravated by the limitations to our ability to reason about the decision-making process of many machine learning models that act as a black box [33], and the lack of comprehensive specifications to the data that needs to be collected to establish a provenance chain (i.e., a record trail of the overall state of the application and its intermediate results) that builds trust on the workflow's results [46].

We aim to tackle these challenges through a rigorous methodology supporting the automatic deployment, the complete analysis cycle and the optimization of applications on the Computing Continuum. We started to implement this methodology in the E2Clab [103] software tool for workflow lifecycle management across the Continuum. For the short term, in the framework of the STEEL project of the PEPR Cloud, we plan to investigate how to further enrich both the methodology and the tool, in order to support the next generation of scientific testbeds (*e.g.,* SLICES-FR) and the non-trivial reproducibility of ML and DL workflows. This is a particularly challenging direction due to the increased degree of randomness (i.e., in terms of initial

parameters and hyperparameters settings) incurred by such applications. We will expand E2Clab to capture provenance metadata during the execution of AI workflows, which includes a detailed record of data sources, processing steps, model configurations, and computational resources utilized. This provenance metadata can be leveraged not only to ensure transparency and traceability throughout the AI lifecycle, but also to conduct resource, energy and performance optimizations. For the longer term we will work towards the definition of ontologies and taxonomies for AI workflow provenance data to build a theoretical foundation for developing provenance data management systems tailored for the different stakeholders involved in AI applications.

Efficient Parallel Continual Learning. Some scenarios of DL training involve the need to assimilate new training data arriving continuously. This kind of 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 mixes samples from previous training tasks with samples from new training tasks to alleviate catastrophic forgetting, but research to date has not addressed performance and scalability of these methods [39, 102, 54, 59].

We propose asynchronous data management techniques that enable the design and implementation of a scalable distributed rehearsal buffer abstraction, which is instrumental in enabling continual learning to take advantage of data-parallel techniques. So far, this solution was validated for class-incremental classification problems. The approach could however be easily applied to generative models (in which case we can simply use one class to store all representatives). This is a short-term research direction we intend to explore in the context of our new research project. For the longer term, we plan to further explore other parallelization approaches (model-based and hybrid) to address the challenges posed by evolving datasets in DL models.

Scalable, Secure and Resource-Efficient Federated Learning. FL aims to achieve an accuracy close to the one achieved by centralized models but in a scalable and resource-efficient manner. Simultaneously, FL is subject to security threats coming from the edge of the network since malicious peers may attempt to manipulate the learning process, compromise the privacy of other peers, or disrupt the training altogether. Clustered FL –grouping clients with similar data distributions and by training personalized models in each identified cluster– is a mechanism to support low resource utilization, but existing approaches [73, 51] mainly focus on the achieved accuracy of the clustering mechanisms, overlooking system and infrastructure resource constraints like energy consumption. At the same time, current threat mitigation approaches [47, 114, 36, 85] rely on robust aggregation, anomaly detection and generative models for defending against poisoning attacks. Yet, they either have limited defensive capabilities due to their underlying design or are impractical to use as they rely on constraining building blocks.

For the short term, we plan to explore approaches to scalable, secure and resource-efficient FL considering for the first time the device heterogeneity, the training accuracy and the robustness against malicious activity simultaneously. A first direction consists in devising resource-constrained clustering algorithms, specifically tailored for FL executed at the edge. The goal is to enable transparent adaptation to the execution environment (e.g., node volatility, malicious attacks, network congestion) by automatically tuning the FL parameters in order to improve user-defined performance metrics (e.g., energy efficiency, execution time, accuracy). Generative model based approaches have been gaining increasing interest, and are shown to be more resilient against a wider range of attacks. In this context, we will continue to extend FedGuard [60], our novel FL framework that utilizes the generative capabilities of Conditional Variational AutoEncoders (CVAE) to effectively defend against poisoning attacks with tuneable overhead in communication and computation. We plan to enhance the robustness of this approach through new aggregation operators and under different levels of dataset imbalance, including highly imbalanced datasets with very few samples per client. For the longer term, a more challenging direction that we plan to explore is how FedGuard and other strategies perform in a setup where clients get access to a stream of incoming data (i.e., dynamic datasets).

3.2 Axis 2: Data-Aware Middleware Approaches for the Computing Continuum

Supporting emerging scenarios over the domains of “modelling and simulation”, “AI”, “Analytics” and “Internet of things” (IoT) across the Computing Continuum leads to new data movement challenges. This is

due in part to the variety of storage systems and to the increasing gaps between the processing paradigms that developed separately in these environments. In this context, it is necessary to explore the different ways in which storage models can converge, notably through a thorough understanding of workloads on the one hand, and efficient data-aware middleware on the other. Within KerData, we propose to address this problem from multiple complementary perspectives.

Firstly, we will study the I/O behavior of scientific workloads running across the Computing Continuum. Understanding what, when, where, and how I/O-intensive applications read or write data is decisive for making the right decisions, especially when it comes to scheduling. Optimization goals need to consider both performance and energy efficiency and potential necessary trade-offs between both. We will then study I/O optimization techniques for extreme-scale workloads. Although a lot of research has been produced on this subject, the expansion to very large scale as in the case of HPC systems raises new challenges that must be addressed to accelerate the time to solution while maintaining sustainability. Next, we will look at the abstraction of distributed storage resources on hybrid infrastructures as a first step towards general-purpose middleware solutions that are necessary to interoperate the storage resources in the continuum. Finally, building on the above, we will propose a data exchange layer that can interoperate with the various platforms on the Computing Continuum. This layer will be central to the composition of hybrid workflows.

Workflow I/O Behavior Analysis Methods for Sustainability. Understanding how workflows and applications use staging areas on the Computing Continuum is decisive for improving scheduling algorithms and deploying an optimized I/O software stack [63]. This requires first characterizing these applications and workflows from an I/O point of view, i.e. determine through performance evaluation and empirical study a relatively high-level set of characteristics that describes the data access pattern [101, 98]. Data collection and analysis will also leverage semi-supervised clustering methods and federated learning techniques from Axis 1. The result of this characterization can then be used to feed job and I/O scheduling algorithms and improve data movement efficiency. The preliminary step in this characterization is the collection of execution traces, from which detailed studies can be carried out [55]. In the field of high-performance computing, Darshan[107] is the reference tool for I/O monitoring.

We will extend our existing work on PyDarshan [89, 88], a Python library for querying Darshan log records, and develop new tools and abstractions applicable throughout the computing continuum that generalize "decision support services" that allow the augmentation of workflow execution plans with I/O and energy behavior information which support our resource management and system architecture research. For example, by identifying the most energy-intense operations, candidates for hardware acceleration can be determined. This short-term direction is the first step to provide support for the design of future strategies for I/O scheduling considering performance/energy trade-offs, that we further plan to investigate for the longer term. This line of research also supports Axis 3 on resource management.

Abstraction for HPC/Cloud Storage Convergence. On HPC and Cloud infrastructures, while the number of processing units has grown to meet the computing power requirements of large-scale applications, the I/O capacity as well as the I/O bandwidth per core have drastically decreased. Thus, data management and analytics becoming the critical bottleneck on large-scale systems, vendors have overcome this problem by deploying new tiers of intermediate storage between the applications and the global shared storage system, usually along with a dedicated (and sometimes proprietary) software layer. These new levels of storage hierarchy feature various capacities, characteristics and performance one has to be aware of to fully utilize them [87]. This is especially true in the context of complex hybrid workflows such as in-situ analysis, visualization or code-coupling: the unawareness of those underlying tiers is a serious loss of performance [72]. An approach focusing on storage convergence across HPC and Cloud infrastructure is decisive to glue this deep hierarchy and to make the most of these new technologies on one side and ensure effective data sharing between components running across the computing continuum [34, 43].

Identifying a good storage abstraction that is accurate enough to properly describe the wide variety of devices and sufficiently general to be portable on various systems is crucial. In that context, we will work on the development of a two-stage abstraction layer above local (system) and remote (distant platform) storage resources. To do so, we will follow a co-design approach, whereby the HPC, Cloud, and Edge-computing architectures would all benefit from an infrastructure-wide level of abstraction. This work will be a

continuation of the research undertaken on data aggregation [112, 113], for which an abstraction of network topology and memory and storage levels was necessary for the algorithm's portability, and will also build on existing work in the community on resource abstraction [52, 111]. In the longer term, we plan to extend this library, which focuses on physical resources, with a logical layer. More concretely, we will build on top of that library a tier-to-tier data transfer layer enabling compatibility between several storage paradigms (block, file, object).

Exascale In-Situ Analytics. Without a major change in practices, the increased computing capacity of the next generation of computers will lead to an explosion in the volume of data produced by numerical simulations. Managing this data, from production to analysis, is a major challenge. While it is not conceivable to do without a storage system, many experiments are aimed at reducing its use. Thus emerged approaches leveraging *in-situ processing*, *in-transit processing*, *staging nodes*, *helper cores* [45, 66]. All of these approaches aim to replace the usual write-read process by a means to perform analysis *at the same time as simulation*, a capability of particular interest to physicists. This need has led to the first implementations of in-situ or in-transit analysis systems in simulation codes, and to the creation of specific middleware for asynchronous, scalable post-Exascale systems, such as Damaris [26, 67, 49].

Developed by the KerData team since 2011, Damaris is the team's flagship software for Exascale HPC. Damaris [4] proposes a middleware-level approach to scalable asynchronous I/O management and real-time in situ processing of data from large-scale MPI-based HPC simulations. It leverages the idea of dedicated cores for such tasks performed asynchronously within multicore nodes. Initial feedback from application users clearly shows the need to design a system that can dynamically trigger the activation of new analyses during the simulation run. The timing can be decided either by the simulation code or by an analysis. To maintain high performance results, it is also essential to appropriately leverage the possibility to place analysis tasks on GPUs. These are challenges we plan to address by extending our previous work based on the Damaris approach, in the context of the Exa-DoST project of the NumPEx PEPR (2023-2030), to support the needs of Exascale workloads. In particular, two applications are targeted: SKA [65] in collaboration with the CNRS, Observatoire de Paris and Observatoire de la Côte d'Azur and Gysela [75], in collaboration with the CEA. For the longer term, we expect additional application requirements to emerge during the execution of the NumPEx PEPR program (in particular, in collaboration with the NumPEx Exa-DI project [5], which has set up a dedicated process to support new applications, not identified yet). We plan to contribute to the support of such applications that could exhibit new patterns with respect to in situ analysis.

Sustainable Interoperability Across the Computing Continuum. New endeavors towards interoperability in the continuum are addressing the need for common data spaces through federated data infrastructures in the cloud (e.g., Gaia-X [50], European Open Science Cloud (EOSC) [37], German National Research Data Infrastructure (NFDI) [76]) and converged research infrastructures for leadership-class supercomputing and cloud resources (e.g., FENIX [6], EuroHPC [106], PRACE-RI [86], European Grid Initiative (EGI) [81]). For the long term, key players in the public and private sectors are making strong investments in long-term strategic decisions about how the computing continuum will develop. Specifically, quantum computing is receiving massive support, and one of the mandates of the EuroHPC JU is to acquire and deploy quantum technologies in HPC environments once they reach sufficient maturity [7]. Furthermore, other new technologies like neuromorphic accelerators [56] will increase the heterogeneity in future HPC systems. These non-conventional architectures can also be found in highly energy-efficient IoT devices [96, 68], fast scientific instruments for large-scale science [108], and new approaches for fast and efficient artificial intelligence in cloud and HPC environments [100, 90, 97]. Current works on the integration of emerging technologies into existing computing ecosystems focus on the interoperability and performance of algorithms without considering data-oriented optimizations, and workflow-specific challenges such as task-resource mapping, automation, and provenance are rarely explored [64, 99].

Overall, the successful interoperability between the existing and projected platforms in the Computing

[4] project.inria.fr/damaris/
[5] numpex.org/exadi-development-and-integration/
[6] fenix-ri.eu/infrastructure
[7] eurohpc-ju.europa.eu/about/discover-eurohpc-ju_en#mission

Continuum will depend on middleware able to interoperate and execute in hybrid scenarios. In the short term, we plan to investigate the design a data exchange layer that will be the core of a workflow composition approach that connects with established data staging and transport layers, alleviating the disconnect between raw data management and knowledge-based workflow management in the continuum for better resource balancing, economy, provenance, and data reduction. In addition, the applications themselves could use this information hub to monitor and record the progress of their individual components in smart ways, enriching existing approaches for in situ analysis and workflow reproducibility. To ensure the sustainability of workflow software solutions in an evolving and hyper-heterogeneous landscape, on a longer term we will study the data and access patterns in hybrid workflows involving the interaction with emerging hardware technologies. The final, long-term goal is to contribute to the development of an interoperable software stack by designing data models that address the challenges in encoding, arrangement, locality, and mapping to high-level data abstractions.

3.3 Axis 3: Sustainable Resource Management for the Computing Continuum

With a growing number disciplines relying on compute services in the CC, data centers are faced with significant changes in their workload and services. In addition to "traditional" numerical simulation applications, there is a massive influx of data (sometimes coupled with remote sensors), analytical and learning applications. These applications present significant uncertainty and dynamicity in their resource requirements, due to their intrinsic behavior and data-intensive profiles. At the same time, planetary limits and the ecological crisis will also have a definite impact on the way these computing centers are managed.

Within KerData we propose to investigate next-generation resource management techniques enabled by reconfigurable software-defined hardware and by the recent convergence trend of industry standards for the flexible integration of accelerators and disaggregated memory. Memory disaggregation refers to the decoupling of physical and logical memory resulting in flexibility to leverage underutilized resources without physically needing reconfigure a distributed system. To utilize as efficiently as possible the Computing Continuum we will take advantage of compute optimization down to the lowest available levels, which is becoming feasible because open toolchains down to the chip level are maturing.

Provisioning Storage Resources on Large-Scale Infrastructures. While for years 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 extend across the Computing Continuum. This migration has been mainly motivated by the need of specialized data processing such as data filtering at the Edge or data analysis on Cloud infrastructures. From an I/O and storage perspective, this means having to deal with very different paradigms: infrastructures where direct access to resources is extremely limited due to a very high level of abstraction, on-premise supercomputers offering a low-level approach requiring tight user control, or highly-constrained devices limited in terms of access and reconfiguration. One way to address that is to converge the infrastructures composing the Computing Continuum by exploring ways to provision storage resources distributed across hybrid HPC/Cloud/Edge systems to complex scientific workflows combining data production, simulation and data analysis. However, this implies low-level access to systems that are sometimes difficult to reach, or to resources in production. Simulation is one way of exploring storage provisioning [58, 57, 83].

In this context, we will continue our work on the simulation of storage systems implemented within the StorAlloc [92, 93] simulator. This work has enabled us to demonstrate the correct sizing and partitioning of intermediate storage resources and to work on the modeling of storage systems, including the way in which they distribute data over the available storage spaces. In KerData, we will be exploring new methods for I/O-aware scheduling of jobs on hybrid infrastructures [79]. Based on post-mortem studies of storage systems, we will also work on characterizing workflows running on the Computing Continuum [70] in order to refine job scheduling decisions. In the longer term, our ambition is to propose a calibrated and validated simulator of storage systems distributed across the Computing Continuum. This simulator will enable us to predict the I/O performance and energy cost of complex workflows leveraging Edge, Cloud and HPC resources. To achieve this, we will rely on state-of-the-art simulation frameworks such as WRENCH [57] and SimGrid [58].

Storage Disaggregation and Computational Storage. A key challenge for many large-scale applications is the mismatch between compute power, the dimension of caches and buffers, and the available I/O bandwidth from the network down to the chip level. Important contributing factors to this situation include economies of scale catering to markets with different needs, a prohibitively expensive development process and lack of manufacturing capacity to consider more customized solutions. However, as computing, memory, storage, and network hardware are becoming increasingly modular and re-configurable, it is possible to consider system and storage architectures that would have been prohibitively expensive before [30, 38]. The enabling technologies for some of these developments are the emerging industry standards such as the Compute Express Link (CXL) [74, 84] for more flexible integration of accelerators and disaggregated storage, and the P4 programming language used in re-configurable networking [48].

We will identify the most energy-intensive routines used on the I/O path, both from the service perspective and from the domain perspective, and curate a portable reference library of key algorithms catering to both software and hardware acceleration. We will leverage, for example, open container standards and instruction set architectures such as RISC-V that can be applied both in data centers and resource-constrained edge contexts [116]. High-level technologies such as containers facilitate the development of algorithmic improvements but often also introduce runtime overhead, while low-level hardware acceleration allows to reduce the energy consumption of computations to a minimum. Unfortunately the hardware acceleration of all desired functionality is not possible because of the need to retain flexibility and limitations due to cost, manufacturing, and physical constraints. Instead a careful selection of routines that should be hardware accelerated needs to be performed for which the priorities shift from application to application. We will start by exploring domain and service-specific (e.g., compression, erasure coding, encryption) approaches first and then identify generalized abstractions for common functionality useful across domains. Ultimately, this research enables building reusable, abstract, and fine-grained building blocks that allow the construction of frugal computational storage architectures including the subset of functionality optimized for a particular application or workflow.

Frugal Data Storage Architectures to Support Post-Exascale Workflows. Post-exascale workflows such as digital twins and machine learning require fast access to increasing amounts of data in long-term archives which poses challenging using existing storage technologies. Especially, long-term storage is latency and bandwidth constrained while high-performance storage systems tend to be cost, energy, and capacity constrained. A major obstacle for better utilization of existing technologies lies in the requirements of legacy applications, but due to applications and workflows transitioning to new programming models it becomes possible to consider new storage system architectures. This creates an opportunity to research frugal data storage architectures that integrate computational storage allowing to avoid wait times and stress on contended resources such as the network and storage subsystems while also increasing energy efficiency through hardware acceleration.

High I/O performance and energy-efficient storage designs require taking a domain- or application-specific approaches and an extension of computational storage to long-term data archives [74, 84, 69]. By applying the methods for holistic workflow I/O behavior analysis already discussed in Axis 2 to discover bottlenecks, it becomes possible to identify service- and application-specific I/O bottlenecks and apply I/O acceleration building blocks. Using these identified building blocks we will develop software libraries that allow their remote execution and/or hardware acceleration close to the data storage location. A holistic effort is necessary to combine advances and consolidation in data and workflow management – as promoted by the FAIR principles (Findable, Accessible, Interoperable, Reusable) [117] – with emerging open technologies for computation and storage [30, 38]. To this end, we will research low-level software and hardware support for metadata queries well as aggregations on top of self-describing data formats needed by both the application workflow and data management communities. In particular, we will investigate the integration of emerging storage technologies that allow for highly parallel access as found in NAND- and NVRAM-based systems as manufacturing costs go down, as well as DNA-based storage systems when array-based synthesis methods become commercially available [82, 95].

HPC Resource Management Faced with the Environmental Crisis. It is essential to consider the evolution of HPC in the face of the climate crisis, and its impact on our research topics. As in other field, we

have to consider what a "lower-tech" version of HPC would be, how to make it usable. This is the target of this axis. The current trend in HPC has been to outbid each other for new supercomputers, renewing them every 6-7 years to make them ever more powerful. However, this policy seems hardly sustainable. Regular shortages of components [42, 78], the origin (and uniqueness) of the sources of certain materials, coupled with the geopolitical context [61] alone make this growth policy challenging. We will start by trying to evaluate the need for scale from a social perspective: what is the relation between scale and social advance. While the scientific community traditionally relies on various metrics to assess the performance of HPC systems —such as the Top500 list (based on HPL performance), HPCG, Graph500, IO500— these metrics do not capture how HPC contributes to social progress.

Then, in front of the lack of resources, we expect manufacturers to need to take these risks into account in their future machines. American HPC laboratories already have constraints for the 2050 horizon, such as zero-emission procurement. So, the first trend we can expect is an extension of HPC machine lifetimes. This could be followed by a move towards refurbished machines, i.e. machines that use components from other machines. These changes, and the introduction of second-hand hardware, should open up several challenges for system managers. Until now, the number of faults has grown linearly with the number of resources [35]. HPC fault tolerance mechanism assume that the Mean Time Between Failure is large in front of system characteristic time (such as the time to checkpoint data). With second-hand material, the number of fault may increase at a much higher rate, while machine performance would not improve since we are not updating the machines. This would render obsolete existing fault-tolerance mechanisms. To this end we will explore new fault-tolerance mechanisms that could be applicable. Heterogeneity linked to resource unavailability and the increased computational complexity motivates the need for a precise description of available resources. In this context, we will explore alternatives for an efficient design of resource management systems to optimize the use of these resources. In addition, non-fatal faults may be invisible, typically slowdowns/varying performance due to wear and tear. We will investigate how one can detect and manage resources that slow down the calculations performed on it. Of course, one of the challenge of this axis will be to work on defining metrics to evaluate the benefits of various solutions. Indeed, not only is it a multi-dimensional problem (TCO analysis), but it should also consider what has been long known on optimization and Jevons paradox [29].

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 [119] 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 [118].

4.2 Nuclear Fusion

GYSELA-X[8] is our second use case explored in the Exa-DoST project. It is a plasma simulation code developed at CEA as part of several national and international collaborations. This code is also exhibiting data-related challenges with respect to the scalability of I/O, storage and in-situ processing. It is part of a demonstrator for the Alice Recoque Exascale supercomputer.

4.3 Material Science

Coddex (Code de Dynamique des Discontinuités pour l'Étude des cristaux) is a simulation code developed at CEA that solves the equations of continuum mechanics in dynamic hyperelasticity (for instance shocks or rapid loading). It also incorporates the description of behavioral discontinuities of change. Within the Exa-DoST project, this application serves to evaluate the approaches that we propose for efficient on-demand in-situ data analysis. The PhD thesis of Arthur Jaquard explores this research direction.

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 [4] 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. In addition, in the new research program of the team, the whole third research axis is dedicated to frugal and sustainable HPC.

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

Our scientific project includes 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. There is a question of what *sufficient* HPC would mean.

Social impact. When considering sufficiency in HPC, we need to question the use of the resources and if we can reduce them. This is the main challenge of the project on result-scalability [11]: we aim at proposing ways to correctly resize HPC computations by focusing on an evaluation of the output rather than by considering input-based scaling models.

Environmental impact. Part of our research focuses on extending the lifespan of HPC machines, in the hope that it could reduce the environmental impact of the field. We have set-up a working group with different teams at Inria Rennes (PACAP, TARAN) to study the challenges extending the life of supercomputers would raise.

[8] gyselax.github.io

6 Highlights of the year

- Silvina Caino Lores was the recipient of an ANR JCJC project.
- François Tessier served as a Program Chair of the HiPC'25 international conference.
- Guillaume Pallez was appointed Associate Editor at IEEE TPDS and IEEE TOPC. He was nominated a member of the steering committee of SC.
- François Tessier, Alexandru Costan, with the help of Jakob Luettgau, organized the 24th IEEE International Symposium on Parallel and Distributed Computing (ISPDC) in Rennes, France.
- Jakob Luettgau was hired as a permanent Inria Researcher on 1 October 2025.
- Alexandru Costan left the team on 1 October 2025.

7 Latest software developments, platforms, open data

7.1 Latest software developments

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.

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 (Alice Recoque, 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: v1.12.1 of Damaris provides basis for an overhaul of the the Plugin layer: adding event triggers on specific hooks, reorganizing event functioning, and enabling/adding data dependency for events. It includes also the (missing) implementions of the management of Parameter for the string and label types, and the handling of some typos and bugs.

News of the Year: In 2025, an extendable Scheduling layer has been added (yet to be released): to reduce communication costs. Also, to enable dynamic analysis handling capability in Damaris, two main activities have been carry out (yet to be released). The development of a dynamic expression module, and an overhaul of the the Plugin layer (Harmonization of plugin definition, with possibility to pass specific data to each plugin, dynamic event creation, triggers (with condition), event/data dependency, data availability across iteration with 'sliding window'). Furthermore, in the context of NumPEx PEPR project, we enhanced the Damaris / PDI interoperability. We continued the development of the Damaris plugin for PDI (to be release soon), and started working on the PDI plugin in Damaris. With this, the simulation instrumented with PDI (https://pdi.dev/main/), can use Damaris to perform asynchronous data analysis using dedicated resources, and the ones instrumented with Damaris could have full access to PDI ecosystem. In addition, Damaris is now part of the NumPEx Software Catalog (https://numpex-pc5.gitlabpages.inria.fr/tutorials/projects/catalog/index.html).

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

Contact: Gabriel Antoniu

Participant: 8 anonymous participants

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 Pl@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 date 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 coordinated within the PEPR Cloud, we have worked towards adapting E2Clab to run experiment leveraging commercial computing resources provided by Scaleway. Enabling users to occasionally leverage resources provided by Scaleway would give them access to state-of-the art GPU nodes and diversity of computing resources.

Additional contributions include: - Ongoing experiments with the ECLAT laboratory to provide experimental support to their simulation pipeline - Improved software reliability through testing and usability through easy ssh access to deployed resources - Documented use of the software for new use-cases such as Federated Learning.

Latest release archive: https://gitlab.inria.fr/E2Clab/e2clab/-/releases/v3.6.0

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

Participant: 5 anonymous participants

7.1.3 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

7.1.4 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: Support of file temperature, better detection of periodic behavior and improved performance for very large datasets were implemented in 2025. An intermediate data format, based on the so-called Trace Event Format, was developed for MOSAIC to convert traces from I/O monitoring tools (such as Darshan, Recorder, and so on) to a common abstraction.

Publication: hal-04808300

Contact: François Tessier

7.1.5 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)

7.1.6 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)

7.2 Open data

7.2.1 I/O Traces

For our IPDPS'25 paper [13], we used traces of I/O activity from four different systems to answer a set of questions about temporal I/O behavior. To focus on realistic workloads, we gathered traces from jobs running over a period of time instead of profiling a limited set of selected applications.

Two of these data sets were Darshan traces available online (from the Intrepid[9] and Blue Waters systems[10]), while two others were obtained by us, using file system monitoring tools:

- PlaFRIM (BeeGFS): a 192-nodes experimental platform in Bordeaux, monitored during 26 months (2022–2024).
- SDumont (Lustre): the largest supercomputer in Latin America, monitored during 12 months (2020).

The collected file system data was correlated with the batch scheduler logs to obtain two time series of I/O bandwidth per job (for "reads" and "writes"), with a value per second for PlaFRIM and a value every 15 seconds for SDumont. The two datasets as well as all code and instructions on how to reproduce our experiments are provided in Zenodo: https://doi.org/10.5281/zenodo.14965920. As explained in the instructions, additional information can be obtained from https://github.com/tuda-parallel/FTIO/tree/main/artifacts/ipdps25 for FTIO, and https://zenodo.org/records/13785395 for MOSAIC.

8 New results

8.1 Supporting Data-Centric Applications and Workflows Running Across the Computing Continuum

8.1.1 On the Reproducibility Challenges of Federated Learning: Investigating the Gap between Simulation, Emulation and Real-World Deployments

[9] https://reports.alcf.anl.gov/data/
[10] https://bluewaters.ncsa.illinois.edu/data-sets.html

Participants: Cédric Prigent, Alexandru Costan, Gabriel Antoniu.

Collaboration. *This work has been carried out in co-operation with Cédric Tedeschi (University of Rennes, MAGELLAN team), Loïc Cudennec (DGA MI) and Kate Keahey (Argonne National Laboratory), in the framework of the STEEL project of the PEPR CLOUD program and of the UNIFY 2 Associate Team with ANL, associated to teh JLESC international laboratory.*

Federated Learning (FL) is an emerging paradigm for decentralized training of Machine Learning models. It has been the subject of a large corpus of research due to its innovative approach to handling sensitive data. A common practice in the 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 has been published as [16].

8.1.2 Evaluating Federated Learning Workflows Beyond Simulation: A Deployment-Aware Methodology

Participants: Mathis Valli, Alexandru Costan, Gabriel Antoniu.

Collaboration. *This work has been carried out in co-operation with Cédric Tedeschi (University of Rennes, MAGELLAN team) and Loïc Cudennec (DGA MI).*

Federated Learning (FL) is often evaluated in simulation, which overlooks network variability, system heterogeneity, and energy costs in geo-distributed settings. We propose a deployment-aware methodology that triangulates *analytical modeling*, *simulation*, and *real-world deployments* within a unified FL evaluation framework. For a given series of experimental scenarios, the methodology allows to assess the consistency of performance trends across the three evaluation approaches, quantifying deviations in key metrics such as run time, communication overhead, and energy consumption. This further enables cross-validation of the reliability of multiple measurement tools, highlighting discrepancies in commonly reported metrics such as the energy usage.

The methodology is validated on FL workloads by comparing analytical predictions and simulations against large-scale deployments on the Grid'5000 testbed, spanning 51 nodes across four geographically distant sites. By varying key FL components such as aggregation algorithms, client sampling rates, and datasets, we characterize how different FL design choices affect the reliability of the three evaluation approaches. Our findings reveal significant divergences: analytical models accurately capture communication patterns and preserve the relative performance of the scenarios, simulations reflect broad trends but often lead to performance rankings of different configurations inconsistent with those found through actual deployment, while only the latter uncovers hidden costs, such as increased energy consumption due to data imbalances.

This work has been submitted for publication to a conference (currently under evaluation).

8.1.3 Supporting SKA data processing workflows with the E2CLab approach to workflow lifecycle management across the continuum

Participants: Thomas Badts, Gabriel Antoniu.

Collaboration. *This work has been carried out in co-operation with Baptiste Besnard and Damien Gratadour* (*LIRA, Observatoire de Paris*)

Tu support automatic deployment, the complete analysis cycle and the optimization of applications on the Computing, we have proposed the E2Clab methodology and its supporting software tool for workflow lifecycle management across the Continuum. We aim to assist the execution of the Karabo pipeline [115] for radioastronomy simulation by enabling reproducible distributed deployments and experiments on academic testbeds. The Karabo pipeline is being developped to support simulation of the future SKA radiotelescope within the ECLAT laboratory. E2Clab also provides the workflow capabilities to run optimization loops over end-to-end experiments and improve parameter discovery and fine-tuning in a complex, cross-disciplinary, stack of software components ranging from distributed computing frameworks to astrophysics simulations.

This collaboration started in 2005 is still in the exploratory stages, further work is expected in the following year.

8.1.4 Methodology for Automated IoT Experimentation in Controlled Testbeds Prior to Real-World Deployments

Participants: Elias Del Pozo Punal, Silvina Caino Lores, Thomas Badts, Gabriel Antoniu.

Collaboration. *This work has been carried out in co-operation with Felix Garcia-Carballeira and Alejandro Calderon from University Carlos III of Madrid.*

Several tools and frameworks have been proposed to automate deployments in distributed systems. Infrastructure-as-Code (IaC) approaches such as Ansible, Puppet, Salstack, or Chef are widely used to abstract low-level configuration details. In parallel, some frameworks support experiment description and execution in specific research testbeds, such as the cOntrol and Management Framework (OMF) and the OMF Measurement Library (OML) [28]. Despite these advances, existing solutions often remain limited to specific domains or infrastructures, and integrating heterogeneous environments remains a challenge when considering the broader computing continuum, and in particular for IoT deployments. As a result, researchers frequently rely on fragmented tools or manual procedures, which hinder the repeatability and scalability of experiments and ultimately limit the ability to perform consistent pre-deployment validation of IoT systems in controlled environments.

This work proposes a general methodology for automated IoT experimentation and validation in controlled environments. The approach provides a structured workflow for designing, deploying, and executing experiments across different testbeds in a reproducible, scalable manner. It allows researchers to evaluate IoT deployments through controlled simulations and pre-deployment testing, bridging the gap between conceptual design and real-world implementation.

8.2 Data-Aware Middleware Approaches for the Computing Continuum

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

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

Collaboration. *This work has been carried out in close co-operation with Philippe Deniel from CEA, the Inria TADaaM team in Bordeaux within the Exa-DoST project of the PEPR NumPEx program. It also involves a collaboration with 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 decade [11], 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 extended MOSAIC [80], a categorizer that detects I/O patterns from execution traces. MOSAIC extracts I/O operations contained in I/O 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 extended version is under submission in a conference [22] and has been presented as a poster during the annual meeting of the ExaDoST project [21] (an updated version of this poster is also under submission for the PASC 2026 conference). A complementary work on the temporal I/O behavior of HPC applications, in collaboration with Inria Bordeaux and TU Darmsdadt, has been presented at IPDPS'2025, an A-rank conference in the field [19].

8.2.2 Study of I/O interference between jobs

Participants: François Tessier, Méline Trochon.

Collaboration. *This work has been carried out in close co-operation with the Inria TADaaM team in Bordeaux and Jean-Thomas Acquaviva, from DDN within the Exa-DoST project of the PEPR NumPEx program.*

High-performance computing is a key component for accelerating scientific discovery and innovation by enabling rapid processing of complex simulations and large-scale data analyses. As HPC applications grow in scale, the performance of the underlying storage infrastructure, particularly parallel file systems (PFS), becomes critical. These shared systems distribute data across multiple storage targets (OST), but concurrent access by multiple jobs can lead to interference, reducing performance compared to isolated operations. Interference varies depending on application characteristics, often degrading overall bandwidth and causing significant performance variability, sometimes by orders of magnitude.

In the context of Méline Trochon's PhD thesis (CIFRE DDN-Inria) we studied how interference impacts checkpointing, a key fault-tolerance technique in HPC. Checkpointing involves periodically saving application data to persistent storage to recover from failures. As applications handle more data, checkpoint files grow

[11] https://top500.org/

larger, making I/O performance even more crucial. Interference during these operations can severely affect their efficiency, highlighting the need to understand and mitigate its effects.

To do this, we launched a large number of experiments with an application that simulates checkpoint phases and one or more applications that simulate interference. Since the checkpoint application has fixed parameters, we looked at how different configurations of interference workloads may or may not affect I/O performance and to what extent. This work is currently being finalized and a paper is expected to be published in 2026. A pre-print is already available online [23]. This work will continue in 2026, notably through the development of a simulator that will allow us to test more configurations.

8.2.3 Enabling Efficient Runtime Data Analysis to a Crystal Deformation Simulation

Participants: Arthur Jaquard, Silvina Caino Lores, Gabriel Antoniu.

Collaboration. *This work has been carried out in close co-operation with Laurent Colombet (from CEA DAM) and Julien Bigot (CEA/Maison de la Simulation) within the Exa-DoST project of the PEPR NumPEx program.*

Exascale simulations generate massive data volumes that strain I/O and post-hoc analysis. In the framework of Arthur Jaquard's PhD thesis we explore how in-situ analysis as supported by the Damaris in situ middleware can benefit to Coddex, a crystal deformation code, to offload data movement and analysis to dedicated processes. This is achieved by enabling runtime extraction of key diagnostics without writing intermediate files. We evaluated tin hysteresis cases on CEA's INTI cluster (with 14 nodes, 1,728 cores) and compare against a ParaView-based post-hoc pipeline. In situ analysis eliminates per-iteration I/O stalls and reduces output time by up to 5x while preserving overall iteration time, with benefits increasing with the number of tracked variables. This work is conducted within the Exa-DoST project of the PEPR NumPEx program, which aims to build the software infrastructure for the first Exascale machine expected to be set up in France (Alice Recoque, Jules Verne project). It has been published as a poster at the SC25 conference [24].

8.3 Sustainable Resource Management for the Computing Continuum

8.3.1 Result-Scalability: Following the Evolution of Selected Social Impact of HPC.

Participants: Guillaume Pallez.

Collaboration. *This work has been carried out in collaboration with Sally Rose Ellingson from the medical college of the University of Kentucky.*

While the scientific community traditionally relies on various computational metrics to assess the performance of HPC systems –such as the TOP500 list (based on HPL performance), HPCG, Graph500, IO500– these metrics do not capture how HPC contributes to social progress. We propose [11] a novel approach to follow how the growth of HPC systems and the advances of HPC research address concrete social challenges. The uniqueness of these new metrics lies in their ability to not only measure the capabilities of HPC architectures but also to gauge the concrete social advancements achieved through their use: it focuses on the output of the computation instead of its input. Contrarily to current measure, it also promotes the diversity of machines by evaluating the Pareto front created between size and result. We emphasize the need for dynamic, community-driven metrics that can evolve with emerging social needs.

8.3.2 Increasing the Lifetime of HPC Machines: Issues, Implications, and Open Challenges

Participants: Guillaume Pallez, Robin Boezennec.

Collaboration. *This work has been carried out as a large collaboration in Rennes including two different teams*: *PACAP and TARAN, as well as with Brice Goglin* (*TADAAM in Bordeaux*)

Extending the lifetime of High-Performance Computing (HPC) machines is becoming an important concern for a variety of reasons. These include the environmental and human costs associated with chip manufacturing, the rising demands by AI workloads, the soaring prices of accelerator chips, political blocks, and delays in the delivery of next-generation supercomputers. As a community, we must reconsider the traditional HPC paradigm and explore new strategies for making existing HPC infrastructure viable for longer periods. In [18], we highlight the current barriers in prolonging HPC machines lifespan and discuss key technical and operational challenges towards this goal.

8.3.3 Improving Supercomputer Usage with Aging Awareness.

Participants: Guillaume Pallez, Robin Boezennec, Alix Tremodeux.

Lifetime of electronic devices has a critical impact on their environmental footprint. In addition, the high-demand by AI companies of GPU has reduced tremendously their availability for supercomputing centers. Consequently, improving the duration of CPUs and GPUs is becoming a major issue in High Performance Computing (HPC) domain. This contribution [12] investigates the optimization of a machine usage before a fatal failure and the trade-offs with performance. The lifetime of computing devices is strongly connected with the temperature and thus with the running frequency. We investigate the node frequency reconfiguration to optimize HPC usage. We estimate the benefit of a dedicated scheduling algorithm compared with a constant frequency.

We show that a correct decision can increase considerably the number of FLOP of a machine with a trade-off in terms of performance. Because aging models are currently inaccurate, we consider different models and discuss the robustness of our algorithms to inaccuracy

8.3.4 Priority-BF: a Task Manager for Priority-Based Scheduling

Participants: Guillaume Pallez.

Collaboration. *This work has been carried out with Ana Gainaru and Scott Patkin* (*Oak Ridge National Laboratory*).

The increasing demand for computational resources, particularly in High-Performance Computing environments, necessitates to rethink how we handle job scheduling strategies. In [14], we address the challenge of managing concurrent jobs with differing priorities on overloaded parallel systems, where strict QoS constraints are often difficult for users to define. Our solution relies on a qualitative description of priorities and pulls from two key approaches: the Easy-BF algorithm and the Conservative Backfilling algorithms. This solution improves the response time for high-priority jobs by 50% without affecting the overall system utilization. We show its applicability in several critical scenarios such as High-Performance Computing (HPC) resource management and in-situ computing.

8.3.5 Scheduling multiple task-based applications on distributed heterogeneous computing nodes

Participants: Etienne Ndamlabin.

Collaboration. *This work has been carried out with Bérenger Bramas* (*CAMUS team, Inria Nancy*).

Modern high-performance computing platforms combine extreme parallelism with growing size, complexity, and cost, making inefficient resource usage increasingly critical in terms of performance and energy. Our research addresses this challenge by focusing on the concurrent execution of multiple task-based applications on shared heterogeneous (CPU/GPU) environments. We created load-balancing heuristics to distribute the task graphs over the processing units and designed and implemented RSCHED, an adaptive scheduling framework integrated into the StarPU runtime system [94]. RSCHED dynamically reorganizes resource allocation in response to application progress and completion, while jointly optimizing application makespan and resource utilization during concurrent execution. Experimental results real applications show up ,to a 10× reduction in overall makespan compared to consecutive execution, while increasing resource utilization. RSCHED also highlights the benefits of system-level coordination on top of independent application schedulers, compared to unsupervised concurrent execution.

8.4 Methodological study over the practice of HPC Research

The following contributions are not necessarily building on the team project but are more adjacent. They both discuss how our community performs research, the first one by studying the reproducibility evaluation process of a large HPC conference (SC'24), and the second one by stuying some claims behind the use of LLM to generate scheduling algorithms.

8.4.1 Implementing a Reproducibility Initiative in HPC: Experiences from SC24.

Participants: Guillaume Pallez.

Collaboration. *This work has been carried out with Sascha Hunold* (*University of Vienna*) *and Judith Hill* (*Lawrence Livermore National Laboratory*).

Reproducibility is fundamental to scientific research, but can be particularly challenging in research that involves High Performance Computing (HPC) due to the unique characteristics of supercomputers. Performance-based metrics such as execution time, energy consumption, and throughput further complicate reproducibility, especially on shared systems. In [15], we present our experience implementing a reproducibility initiative at SC24, with particular emphasis on changes made compared to prior SC conferences. We outline HPC-specific challenges, describe the measures adopted to address them, and reflect on the limitations of reproducibility badges. Faced with the constraints of the existing badging nomenclature, we discuss our implementation of a reproducibility report, which aims to provide more context about the reproducibility of each paper. We conclude by recommending that the "Artifact Replicable" badge be dropped by HPC conferences at this time, and discuss alternate ways of ensuring replicability evaluation.

8.4.2 An In-depth Study of LLM Contributions to the Bin Packing Problem

Participants: Guillaume Pallez.

Collaboration. *This work has been carried out with Julien Herrmann (CNRS).*

Recent studies have suggested that Large Language Models (LLMs) could provide interesting ideas contributing to mathematical discovery. This claim was motivated by reports that LLM-based genetic algorithms produced heuristics offering new insights into the online bin packing problem under uniform and Weibull distributions. In [20], we reassess this claim through a detailed analysis of the heuristics produced by LLMs, examining both their behavior and interpretability. Despite being human-readable, these heuristics remain largely opaque even to domain experts. Building on this analysis, we propose a new class of algorithms tailored to these specific bin packing instances. The derived algorithms are significantly simpler, more efficient, more interpretable, and more generalizable, suggesting that the considered instances are themselves relatively simple. We then discuss the limitations of the claim regarding LLMs' contribution to this problem, which appears to rest on the mistaken assumption that the instances had previously been studied. Our findings instead emphasize the need for rigorous validation and contextualization when assessing the scientific value of LLM-generated outputs.

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

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

Swann Perarnau

Status Senior Scientist

Institution of origin: Argonne National Laboratory

Country: USA

Dates: Dec 8-10, 2025

Context of the visit: Jury for PhD of Robin Boezennec

Mobility program/type of mobility: lecture

9.2.2 Visits to international teams

Research visits abroad

Gabriel Antoniu, Jakob Luettgau, Arthur Jaquard, Robin Boezennec

Visited institution: Argonne National Laboratory
Country: USA
Dates: 13-15 May 2025
Context of the visit: Exploration of research collaboration on in situ processing with Tom Peterka and Orçun Yildiz.
Mobility program/type of mobility: Visit during the JLESC workshop.

9.3 European initiatives

9.3.1 H2020 projects

EUPEX EUPEX project on 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-UNIVERSITAT MAINZ, Germany
- FORSCHUNGSZENTRUM JULICH GMBH (FZJ), Germany
- COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES (CEA), France
- IDRYMA TECHNOLOGIAS KAI EREVNAS (FOUNDATION FOR RESEARCH AND TECHNOLOGYHELLAS), Greece
- SVEUCILISTE U ZAGREBU FAKULTET ELEKTROTEHNIKE I RACUNARSTVA (UNIZG-FER), Croatia

- UNIVERSITA DEGLI STUDI DI TORINO (UNITO), Italy
- Consortium Ubiquitous Technologies S.c.a.r.l. (CUBIT), Italy
- CYBELETECH, France
- UNIVERSITA DI PISA (UNIPI), Italy
- GRAN SASSO SCIENCE INSTITUTE (GSSI), Italy
- ISTITUTO NAZIONALE DI ASTROFISICA (INAF), Italy
- UNIVERSITA DEGLI STUDI DEL MOLISE, Italy
- E 4 COMPUTER ENGINEERING SPA (E4), Italy
- CONSIGLIO NAZIONALE DELLE RICERCHE (CNR), Italy
- JOHANN WOLFGANG GOETHE-UNIVERSITAET 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 - UNIVERSITA 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.

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. 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. A white paper of this group [25] was published in 2025.

9.4 National initiatives

Exa-DoST

Participants: Gabriel Antoniu, François Tessier, Julien Monniot, Jakob Luetgau, Etienne Ndamlabin, Silvina Caino Lores, Guilaume 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'Azure
- 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 Luettgau, François Tessier, Mathis Valli, Thomas Badts.

Title: Secure and efficient daTa storagE 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, Thomas Badts.

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
Summary: 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

10.1.1 Scientific events: organisation

General chair, scientific chair

François Tessier

- General co-Chair of ISPDC 2025, the 24th IEEE International Symposium on Parallel and Distributed Computing (Rennes, France).
- Workshop co-Chair of ESSA 2025, the 6th Workshop on Extreme-Scale Storage and Analysis held in conjunction with IPDPS 2025 (Milan, Italy).
- Workshop co-Chair of Supercompcloud, the 9th Workshop on Interoperability of Supercomputing and Cloud Technologies combined with OpenCHAMI held in conjunction with ISC 2025 (Hamburg, Germany).

Alexandru Costan:

- General co-Chair of ISPDC 2025, the 24th IEEE International Symposium on Parallel and Distributed Computing (Rennes, France).
- Workshop co-Chair of FlexScience 2025, the 15th Workshop on AI and Scientific Computing at Scale using Flexible Computing Infrastructures, held in conjuncciton with ACM HPDC 2025 (Notre Dame, USA).

Guillaume Pallez

- Co-General chair of IPDPS'26, 40th IEEE International Parallel & Distributed Processing Symposium (New Orleans, USA).
- Member of the Steering Committee of ICPP, International Conference on Parallel Processing.

Silvina Caino Lores

- General Co-Chair of WORKS 2025, the 20th Workshop on Workflows in Support of Large-Scale Science, held in conjunction with SC 2025 (St. Louis, 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.
- General Co-Chair of the 1st Workshop on Research Infrastructures for Experimenting across the HPC-Cloud-Edge Continuum(ContinuumRI), held in conjunction with the ACM/IEEE CCGRI 2025.

Member of the organizing committees

Jakob Luettgau:

- Proceedings Chair of ISPDC 2025, the 24th IEEE International Symposium on Parallel and Distributed Computing (Rennes, France).
- Co-organizer of the Birds of a Feather Session Ethics in HPC held in conjunction with ISC 2025 (Hamburg, Germany)
- Co-organizer of the Minsymposium Ëthical and Societal Considerations for Scientific Computing¨ held in conjunction with ISC 2025 (Brugg, Switzerland)
- Co-organizer of the Birds of a Feather Session: CSx4HPC: Computational Storage for High-Performance Computing held in conjunction with SC 2025 (St. Louis, USA)
- Co-organizer of the Birds of a Feather Session: Ethics in HPC held in conjunction with SC 2025 (St. Louis, USA)
- Co-organizer of the Ethics in HPC Birds of a Feather Session BoF: Ethics in HPC held in conjunction with SC 2025 (St. Louis, USA)

Gabriel Antoniu:

- Co-Leader of the Working group on Data management and Computing Continuum at the InPEx workshop on Post-Exascale Computing organized in Kanagawa, Japan.

Théo Jolivel:

- Web Chair of ISPDC 2025, the 24th IEEE International Symposium on Parallel and Distributed Computing (Rennes, France).

Arthur Jaquard:

- Web Chair of WORKS 2025, the 20th Workshop on Workflows in Support of Large-Scale Science (St. Louis, MO, USA)

10.1.2 Scientific events: selection

Chair of conference program committees

François Tessier

- Program Co-Chair of HiPC 2025, the 32nd edition of the IEEE International Conference on High Performance Computing, Data, and Analytics (Hyderabad, India).

Member of the conference program committees

François Tessier: CCGrid2025, ISC25 (Workshop proposals)
Gabriel Antoniu: HPDC 2025, Cluster 2025
Alexandru Costan: SC'25 (Posters and ACM SRC track), IPDPS 25 (PhD Forum), EuroPar 2025, BigData 2025, HiPC 2025, CCGrid 2025

Reviewer

Théo Jolivel:

- IEEE CCGrid25

Mathis Valli:

- IEEE BigData 2025

Arthur Jaquard:

- CCGRID2025

10.1.3 Journal

Member of the editorial boards

Guillaume Pallez:

- IEEE TPDS
- IEEE TOPC

10.1.4 Invited talks

Guillaume Pallez:

- « Vers un calcul intensif plus sobre », organisé par *Laboratoire 1.5*
- "Model (co)-Design and Accuracy for Resource Management in HPC" at Co-Design workshop (Osaka, Jn) co-organized by Jack Dongarra and the Chinese Academy of Science

François Tessier:

- "The Difficult Task of Understanding I/O Behavior on Large-scale Systems", Keynote talk at the 3rd NHR Conference, Germany

10.1.5 Leadership within the scientific community

Gabriel Antoniu:

- Large 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.
- International lab management: Executive Director of JLESC for Inria since April 2024 (previously Vice Executive Director). JLESC is the Joint Inria-Illinois-ANL-BSC-JSC-RIKEN/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.
- International Working Group management: Co-Leader of the Working group on Data management and Computing Continuum within the InPEX International Post-Exascale Project.
- Team management: Head of the KerData Project-Team (INRIA-INSA Rennes).
- International Associate Team management: Leader of the UNIFY2 Associate Team with Argonne National Lab (2013–2025).

François Tessier:

- Work package co-leader with Francieli Boito (Associate Professor, University of Bordeaux) within the NumPEX ExaDoST project.
- Leader for KerData in the ECLAT joint laboratory.

Alexandru Costan:

- Work package leader of WP2 within the PEPR CLOUD STEEL project.

10.1.6 Scientific expertise

Gabriel Antoniu:

- Evaluator for a Horizon Europe project (HORIZON-CL4-2021-HUMAN-01 call)

Alexandru Costan:

- 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)
- Member of the jury for GDR RSD Prix de thèse, Prix chercheur

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)
- Member of the Scientific Board of Inria

Gabriel Antoniu:

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

10.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

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, 21 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.

Jakob Luettgau:

- Master: Cloud and Network Infrastructures (CNI), 4 hours (lectures), M2 level, 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.

10.2.2 Supervision

Defended PhD theses:

- 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). Defended on 25 May 2025.
- Robin Boezennec, “Reducing HPC Resource Consumption”, defended on December 10th, 2025, co-advised by Guillaume Pallez and Fanny Dufossé (Datamove, Grenoble). Defended on 10 December 2025.

PhD in progress:

- 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).
- Méline Trochon, "Adaptive Checkpoint-Restart System with Knowledge of the Network Load", CIFRE thesis started in February 2025, located at Inria Bordeaux, co-supervised by Francieli Boito, Brice Goglin (TADaaM - Inria Bordeaux), Jean-Thomas Acquaviva (DDN) and François Tessier.
- Serge Meurrens, "Ordonnancement des E/S adapté aux applications dans les systèmes HPC", thesis started in December 2025, located at Inria Bordeaux, co-supervised by Francieli Boito, Luan Teylo (TADaaM - Inria Bordeaux) and François Tessier.
- Simon Renard, “Data Interfaces for Hybrid Quantum-Classical Computational Workflows”, thesis started on October 2025, co-supervised by Silvina Caino Lores,Gabriel Antoniu and Marc Baboulin (Inria Paris-Saclay).
- Alix Tremodeux, “Etude des conséquences du vieillissement sur les machines HPC”, thesis started on September 2025, co-supervised by Guillaume Pallez and Erven Rohou (PACAP - Inria Rennes).

Internships:

- Remy Chiv, "Analyse et optimisation des entrées/sorties d’un pipeline de traitement de données pour la radio-astronomie à grande échelle", 5-month Master 2 internship started in May 2025, supervised by François Tessier.

10.2.3 Juries

Gabriel Antoniu:

- HDR: *Towards Better I/O Resource Usage in HPC*, Francieli Zanon Boito, Université de Bordeaux, defended on 5 December 2025.

Alexandru Costan:

- PhD: *Complexity and Algorithmic results for Translocation Distances*, Maria Constantinescu, University of Bucharest, defended on 29 May 2025.

11 Scientific production

11.1 Major publications

[1] G. Antoniu, P. Valduriez, H.-C. Hoppe and J. Krüger. *Towards Integrated Hardware/Software Ecosystems for the Edge-Cloud-HPC Continuum*. 2021. DOI: 10.5281/zenodo.5534464. URL: https://hal.archives-ouvertes.fr/hal-03358930.

[2] 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* 9.4 (2024), pp. 1–28. DOI: 10.1145/3701986. URL: https://hal.science/hal-04187517.

[3] F. Boito, L. Teylo, M. Popov, T. Jolivel, F. Tessier, J. Luettgau, J. Monniot, A. Tarraf, A. Carneiro and C. Osthoff. 'A Deep Look Into the Temporal I/O Behavior of HPC Applications'. In: *39th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*. 39th IEEE International Parallel & Distributed Processing Symposium (IPDPS). Milan, Italy, 3rd June 2025. DOI: 10.1109/IPDPS64566.2025.00072. URL: https://inria.hal.science/hal-04887809.

[4] 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. URL: https://hal.inria.fr/hal-01353890 (cit. on p. 18).

[5] A. Gainaru, B. Goglin, V. Honoré and G. Pallez. 'Profiles of upcoming HPC Applications and their Impact on Reservation Strategies'. In: *IEEE Transactions on Parallel and Distributed Systems* 32.5 (May 2021), pp. 1178–1190. DOI: 10.1109/TPDS.2020.3039728. URL: https://inria.hal.science/hal-03010676.

[6] 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. URL: https://hal.science/hal-03937825.

[7] S. Neuwirth, P. Deniel, J.-T. Acquaviva, M. Golasowski, M. Hennecke, A. Jackson, T. Leibovici, J. Luettgau and R. Nou. *ETP4HPC SRA 6 White Paper - I/O and Storage*. Jan. 2025. DOI: 10.5281/zenodo.14605692. URL: https://hal.science/hal-04905038.

[8] C. Prigent, K. Keahey, A. Costan, L. Cudennec and G. Antoniu. 'On the Reproducibility Challenges of Federated Learning: Investigating the Gap between Simulation, Emulation and Real-World Deployments'. In: CCGrid 2025 - IEEE 25th International Symposium on Cluster, Cloud and Internet Computing. Tromso, Norway, 2025, pp. 185–194. DOI: 10.1109/ccgrid64434.2025.00054. URL: https://hal.science/hal-04997547.

[9] 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. URL: https://hal.archives-ouvertes.fr/hal-02916032.

[10] 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. URL: https://hal.science/hal-04902079.

11.2 Publications of the year

International journals

[11] S. Ellingson and G. Pallez. 'Result-Scalability: Following the Evolution of Selected Social Impact of HPC'. In: *International Journal of High Performance Computing Applications* 39.5 (24th Apr. 2025), pp. 713–721. DOI: 10.1177/10943420251338168. URL: https://inria.hal.science/hal-05037241 (cit. on pp. 18, 27).

International peer-reviewed conferences

[12] R. Boëzennec, F. Dufossé, G. Pallez and A. Tremodeux. 'Improving Supercomputer Usage with Aging Awareness'. In: *SC Workshops '25: Proceedings of the SC '25 Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis*. Sustainable Supercomputing (Workshop of SC25). St. Louis, Missouri, United States: ACM, 15th Nov. 2025, pp. 1980–1989. DOI: 10.1145/3731599.3767561. URL: https://hal.science/hal-05109521 (cit. on p. 28).

[13] F. Boito, L. Teylo, M. Popov, T. Jolivel, F. Tessier, J. Luettgau, J. Monniot, A. Tarraf, A. Carneiro and C. Osthoff. 'A Deep Look Into the Temporal I/O Behavior of HPC Applications'. In: *39th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*. 39th IEEE International Parallel & Distributed Processing Symposium (IPDPS). Milan, Italy, 3rd June 2025. DOI: 10.1109/IPDPS64566.2025.00072. URL: https://inria.hal.science/hal-04887809 (cit. on p. 23).

[14] A. Gainaru, S. Klasky and G. Pallez. 'Priority-BF: a Task Manager for Priority-Based Scheduling'. In: *31st International European Conference on Parallel and Distributed Computing*. EURO-PAR 2025 - 31st International European Conference on Parallel and Distributed Computing. Dresden, Germany, 2026, pp. 219–232. DOI: 10.1007/978-3-031-99854-6_15. URL: https://inria.hal.science/hal-05109295 (cit. on p. 28).

[15] G. Pallez, J. Hill and S. Hunold. 'Implementing a Reproducibility Initiative in HPC: Experiences from SC24'. In: REP 2025 - 3rd ACM Conference on Reproducibility and Replicability. Vancouver (British Columbia), Canada, July 2025, pp. 1–8. DOI: 10.1145/3736731.3746148. URL: https://hal.science/hal-05148685 (cit. on p. 29).

[16] C. Prigent, K. Keahey, A. Costan, L. Cudennec and G. Antoniu. 'On the Reproducibility Challenges of Federated Learning: Investigating the Gap between Simulation, Emulation and Real-World Deployments'. In: CCGrid 2025 - IEEE 25th International Symposium on Cluster, Cloud and Internet Computing. Tromso, Norway, 2025, pp. 185–194. DOI: 10.1109/ccgrid64434.2025.00054. URL: https://hal.science/hal-04997547 (cit. on p. 24).

Doctoral dissertations and habilitation theses

[17] C. Prigent. 'Towards Efficient and Trustworthy Federated Learning on the Computing Continuum'. INSA de Rennes, 26th May 2025. URL: https://theses.hal.science/tel-05279213.

Reports & preprints

[18] R. Boëzennec, F. Fernandes dos Santos, B. Goglin, A. Kritikakou, G. Pallez, E. Rohou, O. Sentieys and M. Traiola. *Increasing the Lifetime of HPC Machines: Issues, Implications, and Open Challenges*. 2025. URL: https://hal.science/hal-05312072 (cit. on p. 28).

[19] F. Boito, L. Teylo, M. Popov, T. Jolivel, F. Tessier, J. Luettgau, J. Monniot, A. Tarraf, A. Carneiro and C. Osthoff. *A deep look into the temporal I/O behavior of HPC applications -extended version*. RR-9577. Inria & Labri, Univ. Bordeaux, 5th Mar. 2025, pp. 1–42. URL: https://inria.hal.science/hal-04978752 (cit. on p. 26).

[20] J. Herrmann and G. Pallez. *An In-depth Study of LLM Contributions to the Bin Packing Problem*. 30th Oct. 2025. URL: https://hal.science/hal-05340506 (cit. on p. 30).

[21] T. Jolivel and F. Tessier. *MOSAIC: Automatic categorization of I/O patterns*. Ed. by G. Antoniu and P. Deniel. 1st Nov. 2025. URL: https://hal.science/hal-05465856 (cit. on p. 26).

[22] T. Jolivel, F. Tessier, J. Luettgau, G. Antoniu and P. Deniel. *A Methodology for System-Scale I/O Pattern Taxonomy for HPC Workloads*. 2025. URL: https://hal.science/hal-05448514 (cit. on p. 26).

[23] M. Trochon, J.-T. Acquaviva, F. Boito, B. Goglin, F. Tessier and L. Teylo. *On the Impact of Interference from Concurrent Jobs on Checkpointing Performance*. 2025. URL: https://hal.science/hal-05294610 (cit. on p. 27).

Other scientific publications

[24] A. Jaquard. 'Enabling Efficient Runtime Data Analysis to a Crystal Deformation Simulation'. In: SC 2025 - International Conference for High Performance Computing, Networking, Storage, and Analysis. Ed. by G. Antoniu, L. Colombet, S. Caino-Lores and J. Bigot. Saint Louis, MO, USA, United States, 2025. URL: https://inria.hal.science/hal-05465846 (cit. on p. 27).

[25] S. Neuwirth, P. Deniel, J.-T. Acquaviva, M. Golasowski, M. Hennecke, A. Jackson, T. Leibovici, J. Luettgau and R. Nou. *ETP4HPC SRA 6 White Paper - I/O and Storage*. Jan. 2025. DOI: 10.5281/zenodo.14605692. URL: https://hal.science/hal-04905038 (cit. on p. 33).

11.3 Cited publications

[26] M. Dorier, G. Antoniu, F. Cappello, M. Snir and L. Orf. 'Damaris: How to Efficiently Leverage Multicore Parallelism to Achieve Scalable, Jitter-free I/O'. In: CLUSTER 2012 - IEEE International Conference on Cluster Computi,ng. Beijing, China: IEEE, 25th Sept. 2012. URL: https://hal.inria.fr/hal-00715252 (cit. on p. 14).

[27] D. P. Abreu, K. Velasquez, M. Curado and E. Monteiro. 'A Comparative Analysis of Simulators for the Cloud to Fog Continuum'. In: *Simulation Modelling Practice and Theory* (2019), p. 102029 (cit. on p. 11).

[28] F. Abut and M. Kızıldağ. 'Design and Implementation of a Reconfigurable Test Environment for Network Measurement Tools Based on a Control and Management Framework'. In: *Applied Sciences* 15.1 (2025). DOI: 10.3390/app15010487. URL: https://www.mdpi.com/2076-3417/15/1/487 (cit. on p. 25).

[29] B. Alcott, M. Giampietro, K. Mayumi and J. Polimeni. *The Jevons paradox and the myth of resource efficiency improvements*. Routledge, 2012 (cit. on p. 17).

[30] A. Amid, D. Biancolin, A. Gonzalez, D. Grubb, S. Karandikar, H. Liew, A. Magyar, H. Mao, A. Ou, N. Pemberton, P. Rigge, C. Schmidt, J. Wright, J. Zhao, Y. S. Shao, K. Asanović and B. Nikolić. 'Chipyard: Integrated Design, Simulation, and Implementation Framework for Custom SoCs'. In: *IEEE Micro* 40.4 (2020), pp. 10–21. DOI: 10.1109/MM.2020.2996616 (cit. on p. 16).

[31] M. Anagnostou, O. Karvounidou, C. Katritzidaki, C. Kechagia, K. Melidou, E. Mpeza, I. Konstantinidis, E. Kapantai, C. Berberidis, I. Magnisalis et al. 'Characteristics and challenges in the industries towards responsible AI: a systematic literature review'. In: *Ethics and Information Technology* 24.3 (2022), p. 37 (cit. on p. 11).

[32] G. Antoniu, P. Valduriez, H.-C. Hoppe and J. KrÃŒger. *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. URL: https://hal.archives-ouvertes.fr/hal-03358930 (cit. on p. 10).

[33] V. Arya, R. K. Bellamy, P.-Y. Chen, A. Dhurandhar, M. Hind, S. C. Hoffman, S. Houde, Q. V. Liao, R. Luss, A. Mojsilović et al. 'One explanation does not fit all: A toolkit and taxonomy of ai explainability techniques'. In: *arXiv preprint arXiv:1909.03012* (2019) (cit. on p. 11).

[34] M. Asch, T. Moore, R. Badia, M. Beck, P. Beckman, T. Bidot, F. Bodin, F. Cappello, A. Choudhary, B. De Supinski et al. 'Big data and extreme-scale computing: Pathways to convergence-toward a shaping strategy for a future software and data ecosystem for scientific inquiry'. In: *The International Journal of High Performance Computing Applications* 32.4 (2018), pp. 435–479 (cit. on p. 13).

[35] G. Aupy. 'Resilient and energy-efficient scheduling algorithms at scale'. PhD thesis. École Normale Supérieure de Lyon, 2014 (cit. on p. 17).

[36] S. Awan, B. Luo and F. Li. 'Contra: Defending against poisoning attacks in federated learning'. In: *Computer Security–ESORICS 2021: 26th European Symposium on Research in Computer Security, Darmstadt, Germany, October 4–8, 2021, Proceedings, Part I 26*. Springer. 2021, pp. 455–475 (cit. on p. 12).

[37] P. Ayris, J.-Y. Berthou, R. Bruce, S. Lindstaedt, A. Monreale, B. Mons, Y. Murayama, C. Soedergaard, K. Tochtermann and R. Wilkinson. 'Realising the European open science cloud'. In: (2016) (cit. on p. 14).

[38] J. Bachrach, H. Vo, B. Richards, Y. Lee, A. Waterman, R. Avižienis, J. Wawrzynek and K. Asanović. 'Chisel: constructing hardware in a Scala embedded language'. In: *Proceedings of the 49th Annual Design Automation Conference*. DAC '12. San Francisco, California: Association for Computing Machinery, 2012, pp. 1216–1225. DOI: 10.1145/2228360.2228584. URL: https://doi.org/10.1145/2228360.2228584 (cit. on p. 16).

[39] Y. Balaji, M. Farajtabar, D. Yin, A. Mott and A. Li. 'The effectiveness of memory replay in large scale continual learning'. In: *arXiv preprint arXiv:2010.02418* (2020) (cit. on p. 12).

[40] D. Balouek-Thomert, E. G. Renart, A. R. Zamani, A. Simonet and M. Parashar. 'Towards a computing continuum: Enabling edge-to-cloud integration for data-driven workflows'. In: *The International Journal of High Performance Computing Applications* 33.6 (2019), pp. 1159–1174 (cit. on p. 10).

[41] L. A. Barba and G. K. Thiruvathukal. 'Reproducible Research for Computing in Science Engineering'. In: *Computing in Science Engineering* 19.6 (2017), pp. 85–87 (cit. on p. 11).

[42] E. Barrett. *Taiwan's drought is exposing just how much water chipmakers like TSMC use (and reuse)*. https://fortune.com/2021/06/12/chip-shortage-taiwan-drought-tsmc-water-usage/. 2021 (cit. on p. 17).

[43] M. Beck, T. Moore, P. Luszczek and A. Danalis. 'Interoperable Convergence of Storage, Networking, and Computation'. In: *Advances in Information and Communication*. Ed. by K. Arai and R. Bhatia. Cham: Springer International Publishing, 2020, pp. 667–690 (cit. on p. 13).

[44] T. Ben-Nun and T. Hoefler. 'Demystifying parallel and distributed deep learning: An in-depth concurrency analysis'. In: *ACM Computing Surveys* 52.4 (2019), pp. 1–43 (cit. on p. 11).

[45] J. C. Bennett, H. Abbasi, P.-T. Bremer, R. Grout, A. Gyulassy, T. Jin, S. Klasky, H. Kolla, M. Parashar, V. Pascucci, P. Pebay, D. Thompson, H. Yu, F. Zhang and J. Chen. 'Combining in-situ and in-transit processing to enable extreme-scale scientific analysis'. In: *SC '12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. 2012, pp. 1–9. DOI: 10.1109/SC.2012.31 (cit. on p. 14).

[46] E. Bertino, S. Bhattacharya, E. Ferrari and D. Milojicic. 'Trustworthy AI and Data Lineage'. In: *IEEE Internet Computing* 27.6 (2023), pp. 5–6 (cit. on p. 11).

[47] P. Blanchard, E. M. El Mhamdi, R. Guerraoui and J. Stainer. 'Machine learning with adversaries: Byzantine tolerant gradient descent'. In: *Advances in neural information processing systems* 30 (2017) (cit. on p. 12).

[48] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese and D. Walker. 'P4: Programming Protocol-Independent Packet Processors'. In: *ACM SIGCOMM Computer Communication Review* 44.3 (July 2014), pp. 87–95. DOI: 10.1145/2656877.2656890. URL: https://dl.acm.org/doi/10.1145/2656877.2656890 (visited on 06/03/2024) (cit. on p. 16).

[49] J. C. Bowden, F. Tessier, C. Deltel, S. Bnà and G. Antoniu. *In-situ visualization using Damaris: the Code Saturne use case*. Ed. by P. P. for Advanced Computing in Europe. Research Report. PRACE White Paper. Sept. 2021. URL: https://inria.hal.science/hal-03354035 (cit. on p. 14).

[50] A. Braud, G. Fromentoux, B. Radier and O. Le Grand. 'The road to European digital sovereignty with Gaia-X and IDSA'. In: *IEEE network* 35.2 (2021), pp. 4–5 (cit. on p. 14).

[51] C. Briggs, Z. Fan and P. András. 'Federated learning with hierarchical clustering of local updates to improve training on non-IID data'. In: *2020 International Joint Conference on Neural Networks (IJCNN)* (2020), pp. 1–9. URL: https://api.semanticscholar.org/CorpusID:216144447 (cit. on p. 12).

[52] F. Broquedis, J. Clet-Ortega, S. Moreaud, N. Furmento, B. Goglin, G. Mercier, S. Thibault and R. Namyst. 'hwloc: A Generic Framework for Managing Hardware Affinities in HPC Applications'. In: *2010 18th Euromicro Conference on Parallel, Distributed and Network-based Processing*. 2010, pp. 180–186. DOI: 10.1109/PDP.2010.67 (cit. on p. 14).

[53] P. Brox, J. Garcia-Blas, D. E. Singh and J. Carretero. 'DICE: Generic Data Abstraction for Enhancing the Convergence of HPC and Big Data'. In: *High Performance Computing: 8th Latin American Conference, CARLA 2021, Guadalajara, Mexico, October 6–8, 2021, Revised Selected Papers*. Springer. 2022, pp. 106–119 (cit. on p. 11).

[54] P. Buzzega, M. Boschini and S. Calderara. 'Rethinking experience replay: A bag of tricks for continual learning'. In: *25th International Conference on Pattern Recognition (ICPR)*. 2021, pp. 2180–2187 (cit. on p. 12).

[55] P. Carns, R. Latham, R. Ross, K. Iskra, S. Lang and K. Riley. '24/7 characterization of petascale I/O workloads'. In: *2009 IEEE International Conference on Cluster Computing and Workshops*. IEEE. 2009, pp. 1–10 (cit. on p. 13).

[56] P. Carpenter, U.-U. Haus, E. Laure, S. Narasimhamurthy and E. Suarez. *Heterogeneity is here to stay: Challenges and Opportunities in HPC*. Feb. 2022. URL: https://www.etp4hpc.eu/pujades/files/ETP4HPC_WP_Heterogeneous-HPC_20220216.pdf (cit. on p. 14).

[57] 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 (cit. on p. 15).

[58] H. Casanova, A. Giersch, A. Legrand, M. Quinson and F. Suter. 'Versatile, Scalable, and Accurate Simulation of Distributed Applications and Platforms'. In: *Journal of Parallel and Distributed Computing* 74.10 (June 2014), pp. 2899–2917. URL: http://hal.inria.fr/hal-01017319 (cit. on p. 15).

[59] A. Chaudhry, A. Gordo, P. K. Dokania, P. Torr and D. Lopez-Paz. 'Using hindsight to anchor past knowledge in continual learning'. In: *arXiv preprint arXiv:2002.08165* 3 (2020) (cit. on p. 12).

[60] M. Chelli, C. Prigent, R. Schubotz, A. Costan, G. Antoniu, L. Cudennec and P. Slusallek. 'FedGuard: Selective Parameter Aggregation for Poisoning Attack Mitigation in Federated Learning'. In: *Cluster 2023 - IEEE International Conference on Cluster Computing*. Santa Fe, New Mexico, United States: IEEE, Oct. 2023, pp. 1–10. URL: https://hal.science/hal-04208787 (cit. on p. 12).

[61] E. Commission. *Critical Raw Materials Resilience: Charting a Path towards greater Security and Sustainability*. https://publications.europa.eu/resource/cellar/160da878-edc7-11ea-991b-01aa75ed71a1.0015.03/DOC_1. 2020 (cit. on p. 17).

[62] *Contrat d'objectifs et de performance 2019-2023 Entre l'État et Inria*. https://www.inria.fr/sites/default/files/2020-02/COP_INRIA_2019-2023_version-finalesssignatures.pdf/. 2019 (cit. on p. 9).

[63] C. Daley, D. Ghoshal, G. Lockwood, S. Dosanjh, L. Ramakrishnan and N. Wright. 'Performance characterization of scientific workflows for the optimal use of Burst Buffers'. In: *Future Generation Computer Systems* 110 (2020), pp. 468–480. DOI: https://doi.org/10.1016/j.future.2017.12.022. URL: https://www.sciencedirect.com/science/article/pii/S0167739X16308287 (cit. on p. 13).

[64] A. Deshpande. 'Assessing the quantum-computing landscape'. In: *Communications of the ACM* 65.10 (2022), pp. 57–65 (cit. on p. 14).

[65] P. E. Dewdney, P. J. Hall, R. T. Schilizzi and T. J. L. W. Lazio. 'The Square Kilometre Array'. In: *Proceedings of the IEEE* 97.8 (2009), pp. 1482–1496. DOI: 10.1109/JPROC.2009.2021005 (cit. on p. 14).

[66] E. Dirand, L. Colombet and B. Raffin. 'TINS: A Task-Based Dynamic Helper Core Strategy for In Situ Analytics'. In: *SCA18 - Supercomputing Frontiers Asia 2018*. Singapore, Singapore, Mar. 2018, pp. 159–178. DOI: 10.1007/978-3-319-69953-0_10. URL: https://inria.hal.science/hal-01730910 (cit. on p. 14).

[67] 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. URL: https://inria.hal.science/hal-01353890 (cit. on p. 14).

[68] European Association on Smart Systems Integration. *Strategic Research and Innovation Agenda*. 2023. URL: https://ecssria.eu/ECS-SRIA%5C%202023.pdf (cit. on p. 14).

[69] D. Fakhry, M. Abdelsalam, M. W. El-Kharashi and M. Safar. 'A Review on Computational Storage Devices and near Memory Computing for High Performance Applications'. In: *Memories - Materials, Devices, Circuits and Systems* 4 (July 2023), p. 100051. DOI: 10.1016/j.memori.2023.100051. URL: https://www.sciencedirect.com/science/article/pii/S2773064623000282 (visited on 09/01/2024) (cit. on p. 16).

[70] R. Ferreira da Silva, H. Casanova, A.-C. Orgerie, R. Tanaka, E. Deelman and F. Suter. 'Characterizing, Modeling, and Accurately Simulating Power and Energy Consumption of I/O-intensive Scientific Workflows'. In: *Journal of Computational Science* 44 (2020), p. 101157. DOI: https://doi.org/10.1016/j.jocs.2020.101157. URL: https://www.sciencedirect.com/science/article/pii/S1877750320304580 (cit. on p. 15).

[71] A. Gainaru, L. Wan, R. Wang, E. Suchyta, J. Chen, N. Podhorszki, J. Kress, D. Pugmire and S. Klasky. 'Understanding the Impact of Data Staging for Coupled Scientific Workflows'. In: *IEEE Transactions on Parallel and Distributed Systems* 33.12 (2022), pp. 4134–4147 (cit. on p. 11).

[72] A. Gainaru, L. Wan, R. Wang, E. Suchyta, J. Chen, N. Podhorszki, J. Kress, D. Pugmire and S. Klasky. 'Understanding the Impact of Data Staging for Coupled Scientific Workflows'. In: *IEEE Transactions on Parallel and Distributed Systems* 33.12 (2022), pp. 4134–4147. DOI: 10.1109/TPDS.2022.3179989 (cit. on p. 13).

[73] A. Ghosh, J. Chung, D. Yin and K. Ramchandran. 'An Efficient Framework for Clustered Federated Learning'. In: *IEEE Transactions on Information Theory* 68.12 (2022), pp. 8076–8091. DOI: 10.1109/TIT.2022.3192506 (cit. on p. 12).

[74] D. Gouk, S. Lee, M. Kwon and M. Jung. 'Direct Access, High-Performance Memory Disaggregation with DirectCXL'. In: *2022 USENIX Annual Technical Conference (USENIX ATC 22)*. Carlsbad, CA: USENIX Association, July 2022, pp. 287–294. URL: https://www.usenix.org/conference/atc22/presentation/gouk (cit. on p. 16).

[75] V. Grandgirard, Y. Sarazin, X. Garbet, G. Dif-Pradalier, P. Ghendrih, N. Crouseilles, G. Latu, E. Sonnendrücker, N. Besse and P. Bertrand. 'GYSELA, a full-f global gyrokinetic Semi-Lagrangian code for ITG turbulence simulations'. In: *AIP Conference Proceedings* 871.1 (2006), pp. 100–111. DOI: http://dx.doi.org/10.1063/1.2404543. URL: http://scitation.aip.org/content/aip/proceeding/aipcp/10.1063/1.2404543 (cit. on p. 14).

[76] N. Hartl, E. Wössner and Y. Sure-Vetter. 'Nationale Forschungsdateninfrastruktur (NFDI)'. In: *Informatik Spektrum* 44.5 (2021), pp. 370–373 (cit. on p. 14).

[77] E. A. Huerta, A. Khan, E. Davis, C. Bushell, W. D. Gropp, D. S. Katz, V. V. Kindratenko, S. Koric, W. T. C. Kramer, B. McGinty, K. McHenry and A. Saxton. 'Convergence of artificial intelligence and high performance computing on NSF-supported cyberinfrastructure'. In: *Journal of Big Data* 7.1 (2020), p. 88 (cit. on p. 11).

[78] K. Jacobs, S. Chopra, A. Barr and B. Boucher. *Supply shortages and an inflexible market give rise to high power transformer lead times*. https://www.woodmac.com/news/opinion/supply-shortages-and-an-inflexible-market-give-rise-to-high-power-transformer-lead-times/. 2021 (cit. on p. 17).

[79] E. Jeannot, G. Pallez and N. Vidal. 'IO-aware Job-Scheduling: Exploiting the Impacts of Workload Characterizations to select the Mapping Strategy'. In: *International Journal of High Performance Computing Applications* (2023), pp. 1–13. DOI: 10.1177/10943420231175854. URL: https://inria.hal.science/hal-04098706 (cit. on p. 15).

[80] 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*. Atlanta, United States, Nov. 2024, pp. 1–7. DOI: 10.1109/SCW63240.2024.00172. URL: https://hal.science/hal-04808300 (cit. on p. 26).

[81] D. Kranzlmueller, J. M. de Lucas and P. Oester. 'The European Grid Initiative (EGI) Towards a Sustainable Grid Infrastructure'. In: *Remote Instrumentation and Virtual Laboratories: Service Architecture and Networking*. Springer. 2010, pp. 61–66 (cit. on p. 14).

[82] D. Landsman and K. Strauss. 'The DNA Data Storage Model'. In: *Computer* 56.7 (July 2023), pp. 78–85. DOI: 10.1109/MC.2023.3272188. URL: https://ieeexplore.ieee.org/document/10154188/ (visited on 11/12/2023) (cit. on p. 16).

[83] A. Lebre, A. Legrand, F. Suter and P. Veyre. 'Adding Storage Simulation Capacities to the SimGrid Toolkit: Concepts, Models, and API'. In: *2015 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*. 2015, pp. 251–260. DOI: 10.1109/CCGrid.2015.134 (cit. on p. 15).

[84] H. Li, D. S. Berger, S. Novakovic, L. Hsu, D. Ernst, P. Zardoshti, M. Shah, S. Rajadnya, S. Lee, I. Agarwal, M. D. Hill, M. Fontoura and R. Bianchini. *Pond: CXL-Based Memory Pooling Systems for Cloud Platforms*. Oct. 2022. DOI: 10.48550/arXiv.2203.00241. arXiv: 2203.00241 [cs]. URL: http://arxiv.org/abs/2203.00241 (visited on 06/03/2024) (cit. on p. 16).

[85] S. Li, Y. Cheng, W. Wang, Y. Liu and T. Chen. 'Learning to Detect Malicious Clients for Robust Federated Learning'. In: *CoRR* abs/2002.00211 (2020). arXiv: 2002.00211. URL: https://arxiv.org/abs/2002.00211 (cit. on p. 12).

[86] T. Lippert, T. Eickermann and D. Erwin. 'PRACE: Europe's supercomputing research infrastructure'. In: *Applications, Tools and Techniques on the Road to Exascale Computing*. IOS Press, 2012, pp. 7–18 (cit. on p. 14).

[87] G. Lockwood, D. Hazen, Q. Koziol, R. Canon, K. Antypas and J. Balewski. 'Storage 2020: A Vision for the Future of HPC Storage'. In: *Report: LBNL-2001072*. Lawrence Berkeley National Laboratory, 2017. URL: https://escholarship.org/uc/item/744479dp#author (cit. on p. 13).

[88] J. Luettgau, S. Snyder, P. Carns, J. M. Wozniak, J. Kunkel and T. Ludwig. 'Toward Understanding I/O Behavior in HPC Workflows'. In: *2018 IEEE/ACM 3rd International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems (PDSW-DISCS)*. Dallas, TX, USA, Nov. 2018, pp. 64–75. DOI: 10.1109/PDSW-DISCS.2018.00012 (cit. on p. 13).

[89] 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: *Proceedings of the SC '23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis*. SC-W '23. New York, NY, USA: Association for Computing Machinery, Nov. 2023, pp. 1380–1391. DOI: 10.1145/3624062.3624207. URL: https://doi.org/10.1145/3624062.3624207 (visited on 02/12/2023) (cit. on p. 13).

[90] T. Luo, W.-F. Wong, R. S. M. Goh, A. T. Do, Z. Chen, H. Li, W. Jiang and W. Yau. 'Achieving Green AI with Energy-Efficient Deep Learning Using Neuromorphic Computing'. In: *Commun. ACM* 66.7 (June 2023), pp. 52–57. DOI: 10.1145/3588591. URL: https://doi.org/10.1145/3588591 (cit. on p. 14).

[91] 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. URL: https://hal.science/hal-03937825 (cit. on pp. 7, 9, 10).

[92] 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. 15).

[93] 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* 35.28 (2023), e7890. DOI: https://doi.org/10.1002/cpe.7890. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cpe.7890. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.7890 (cit. on p. 15).

[94] E. Ndamlabin and B. Bramas. 'RSCHED: An Effective Heterogeneous Resource Management for Simultaneous Execution of Task-Based Applications'. In: *International Journal of Advanced Computer Science and Applications* 16.2 (2025). DOI: 10.14569/IJACSA.2025.0160206. URL: http://dx.doi.org/10.14569/IJACSA.2025.0160206 (cit. on p. 29).

[95] B. H. Nguyen, C. N. Takahashi, G. Gupta, J. A. Smith, R. Rouse, P. Berndt, S. Yekhanin, D. P. Ward, S. D. Ang, P. Garvan, H.-Y. Parker, R. Carlson, D. Carmean, L. Ceze and K. Strauss. 'Scaling DNA Data Storage with Nanoscale Electrode Wells'. In: *Science Advances* 7.48 (Nov. 2021), eabi6714. DOI: 10.1126/sciadv.abi6714. URL: https://www.science.org/doi/10.1126/sciadv.abi6714 (visited on 01/08/2023) (cit. on p. 16).

[96] C. I. Nwakanma, J.-W. Kim, J.-M. Lee and D.-S. Kim. 'Edge AI prospect using the NeuroEdge computing system: Introducing a novel neuromorphic technology'. In: *ICT Express* 7.2 (2021), pp. 152–157 (cit. on p. 14).

[97] H. Padmanaban. 'Quantum Computing and AI in the Cloud'. In: *Journal of Computational Intelligence and Robotics* 4.1 (Mar. 2024), pp. 14–32. URL: https://thesciencebrigade.com/jcir/article/view/116 (cit. on p. 14).

[98] F. Pan, Y. Yue, J. Xiong and D. Hao. 'I/O Characterization of Big Data Workloads in Data Centers'. In: *Big Data Benchmarks, Performance Optimization, and Emerging Hardware*. Ed. by J. Zhan, R. Han and C. Weng. Cham: Springer International Publishing, 2014, pp. 85–97 (cit. on p. 13).

[99] R. Patton, C. Schuman, S. Kulkarni, M. Parsa, J. P. Mitchell, N. Q. Haas, C. Stahl, S. Paulissen, P. Date, T. Potok et al. 'Neuromorphic computing for autonomous racing'. In: *International Conference on Neuromorphic Systems 2021*. 2021, pp. 1–5 (cit. on p. 14).

[100] M. Raparthi. 'Real-Time AI Decision Making in IoT with Quantum Computing: Investigating & Exploring the Development and Implementation of Quantum-Supported AI Inference Systems for IoT Applications'. In: *Internet of Things and Edge Computing Journal* 1.1 (Mar. 2021), pp. 18–27. URL: https://thesciencebrigade.com/iotecj/article/view/130 (cit. on p. 14).

[101] G. P. Rodrigo Álvarez, P.-O. Östberg, E. Elmroth, K. Antypas, R. Gerber and L. Ramakrishnan. 'HPC System Lifetime Story: Workload Characterization and Evolutionary Analyses on NERSC Systems'. In: *Proceedings of the 24th International Symposium on High-Performance Parallel and Distributed Computing*. HPDC '15. Portland, Oregon, USA: Association for Computing Machinery, 2015, pp. 57–60. DOI: 10.1145/2749246.2749270. URL: https://doi.org/10.1145/2749246.2749270 (cit. on p. 13).

[102] D. Rolnick and A. Ahuja. 'Experience replay for continual learning'. In: *Advances in Neural Information Processing Systems* 32 (2019) (cit. on p. 12).

[103] 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: *2020 IEEE International Conference on Cluster Computing (CLUSTER)*. 2020, pp. 176–186. DOI: 10.1109/CLUSTER49012.2020.00028 (cit. on p. 11).

[104] S. Sadiq, M. Orlowska, W. Sadiq and C. Foulger. 'Data Flow and Validation in Workflow Modelling'. In: *Proceedings of the 15th Australasian database conference-Volume 27*. 2004, pp. 207–214 (cit. on p. 11).

[105] C. Sanderson, Q. Lu, D. Douglas, X. Xu, L. Zhu and J. Whittle. 'Towards Implementing Responsible AI'. In: *2022 IEEE International Conference on Big Data (Big Data)*. IEEE. 2022, pp. 5076–5081 (cit. on p. 11).

[106] T. Skordas. 'Toward a european exascale ecosystem: the eurohpc joint undertaking'. In: *Communications of the ACM* 62.4 (2019), pp. 70–70 (cit. on p. 14).

[107] S. Snyder, P. Carns, K. Harms, R. Ross, G. K. Lockwood and N. J. Wright. 'Modular HPC I/O Characterization with Darshan'. In: *2016 5th Workshop on Extreme-Scale Programming Tools (ESPT)*. 2016, pp. 9–17. DOI: 10.1109/ESPT.2016.006 (cit. on p. 13).

[108] L. Song, F. Chen, S. R. Young, C. D. Schuman, G. Perdue and T. E. Potok. 'Deep learning for vertex reconstruction of neutrino-nucleus interaction events with combined energy and time data'. In: *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE. 2019, pp. 3882–3886 (cit. on p. 14).

[109] V. Stodden and S. Miguez. 'Best Practices for Computational Science: Software Infrastructure and Environments for Reproducible and Extensible Research'. In: *Available at SSRN 2322276* (2013) (cit. on p. 11).

[110] S. Svorobej, P. Takako Endo, M. Bendechache, C. Filelis-Papadopoulos, K. M. Giannoutakis, G. A. Gravvanis, D. Tzovaras, J. Byrne and T. Lynn. 'Simulating Fog and Edge Computing Scenarios: An Overview and Research Challenges'. In: *Future Internet* 11.3 (2019), p. 55 (cit. on p. 11).

[111] H. Tang, S. Byna, F. Tessier, T. Wang, B. Dong, J. Mu, Q. Koziol, J. Soumagne, V. Vishwanath, J. Liu and R. Warren. 'Toward Scalable and Asynchronous Object-Centric Data Management for HPC'. In: *2018 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)*. 2018, pp. 113–122. DOI: 10.1109/CCGRID.2018.00026 (cit. on p. 14).

[112] F. Tessier, V. Vishwanath and E. Jeannot. 'TAPIOCA: An I/O Library for Optimized Topology-Aware Data Aggregation on Large-Scale Supercomputers'. In: *2017 IEEE International Conference on Cluster Computing (CLUSTER)*. Sept. 2017, pp. 70–80. DOI: 10.1109/CLUSTER.2017.80 (cit. on p. 14).

[113] F. Tessier, P. Gressier and V. Vishwanath. 'Optimizing Data Aggregation by Leveraging the Deep Memory Hierarchy on Large-scale Systems'. In: *Proceedings of the 2018 International Conference on Supercomputing*. ICS '18. Beijing, China: ACM, 2018, pp. 229–239. DOI: 10.1145/3205289.3205316. URL: http://doi.acm.org/10.1145/3205289.3205316 (cit. on p. 14).

[114] V. Tolpegin, S. Truex, M. E. Gursoy and L. Liu. 'Data Poisoning Attacks Against Federated Learning Systems'. en. In: *Computer Security – ESORICS 2020*. Ed. by L. Chen, N. Li, K. Liang and S. Schneider. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2020, pp. 480–501. DOI: 10.1007/978-3-030-58951-6_24 (cit. on p. 12).

[115] [SW] University of Applied Sciences Northwestern Switzerland, *Karabo-Pipeline* version v0.34.0. LIC: MIT. URL: https://github.com/i4Ds/Karabo-Pipeline (cit. on p. 25).

[116] A. Waterman, Y. Lee, R. Avizienis, H. Cook, D. Patterson and K. Asanovic. 'The RISC-V instruction set'. In: *2013 IEEE Hot Chips 25 Symposium (HCS)*. 2013, pp. 1–1. DOI: 10.1109/HOTCHIPS.2013.7478332 (cit. on p. 16).

[117] M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L. B. da Silva Santos, P. E. Bourne et al. 'The FAIR Guiding Principles for Scientific Data Management and Stewardship'. In: *Scientific Data* 3.1 (Mar. 2016), p. 160018. DOI: 10.1038/sdata.2016.18. URL: https://www.nature.com/articles/sdata201618 (visited on 06/03/2024) (cit. on p. 16).

[118] *ECLAT - Extreme Computing Lab for Astronomical Telescopes*. 2024. URL: https://eclat-lab.fr/ (cit. on p. 17).

[119] *SKA - Square Kilometre Array*. 2024. URL: https://www.skao.int/en (cit. on p. 17).