RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

2025 *Activity Report*

RESEARCH CENTRE: **Inria Centre at Université Grenoble Alpes**
IN PARTNERSHIP WITH: **Université de Grenoble Alpes, CNRS**

Project-Team

DATAMOVE

Data Aware Large Scale Computing

In collaboration with Laboratoire d'Informatique de Grenoble (LIG)

Project-Team DATAMOVE

Creation of the Project-Team: 2017 November 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

A1.1.4. – High performance computing

A1.1.5. – Exascale

A1.3.6. – Fog, Edge

A1.6. – Green Computing

A2.6.2. – Middleware

A2.6.4. – Ressource management

A7.1.1. – Distributed algorithms

A7.1.2. – Parallel algorithms

A9.7. – AI algorithmics

A9.9. – Distributed AI, Multi-agent

Other research topics and application domains

B3.3. – Geosciences

B6.4. – Internet of things

Contents

Project-Team DATAMOVE 1

1 Team members, visitors, external collaborators 5

2 Overall objectives 7

3 Research program 7
3.1 Motivation . . . 7
3.2 Strategy . . . 8
3.3 Research Directions . . . 9

4 Application domains 9
4.1 Data Aware Batch Scheduling . . . 9
4.1.1 Algorithms . . . 10
4.1.2 Locality Aware Allocations . . . 10
4.1.3 Data-Centric Processing . . . 11
4.1.4 Learning . . . 11
4.1.5 Multi-objective Optimization . . . 12
4.2 Empirical Studies of Large Scale Platforms . . . 12
4.2.1 Workload Traces with Resource Consumption . . . 13
4.2.2 Simulation . . . 13
4.2.3 Job and Platform Models . . . 14
4.2.4 Emulation and Reproducibility . . . 14
4.3 Integration of High Performance Computing and Data Analytics . . . 14
4.3.1 Programming Model and Software Architecture . . . 15
4.3.2 Resource Sharing . . . 15
4.3.3 Co-Design with Data Scientists . . . 16

5 Social and environmental responsibility 16

6 Highlights of the year 17

7 Latest software developments, platforms, open data 17
7.1 Latest software developments . . . 17
7.1.1 OAR . . . 17
7.1.2 MELISSA . . . 17
7.1.3 NixOS-Compose . . . 18
7.1.4 Batsim . . . 18
7.1.5 Kameleon . . . 18
7.1.6 alumet . . . 18
7.2 New platforms . . . 19
7.2.1 Slices-fr/Grid'5000 and Meso Center Gricad . . . 19

8 New results 19
8.1 Multimodal Vision and Attention-Based Detection . . . 19
8.2 Energy, Carbon Footprint, and Sustainability in HPC and AI . . . 19
8.2.1 Carbon Footprint of High-Performance Computing . . . 19
8.2.2 Power-Aware Scheduling and Dynamic Resource Management . . . 20
8.2.3 Environmental Impact of Generative AI . . . 20
8.3 Computing Continuum: Architectures, Testbeds, and Complexity Management . . . 20
8.3.1 Edge–Cloud and Serverless Scheduling . . . 20
8.3.2 Operational Technology Platforms and Orchestration . . . 21
8.3.3 Testbeds and Complexity Analysis for the Continuum . . . 21

9 Bilateral contracts and grants with industry **21**

10 Partnerships and cooperations **22**
10.1 European initiatives . . . 22
10.1.1 Horizon Europe . . . 22
10.2 National initiatives . . . 24
10.2.1 PEPR NUMPEX . . . 24
10.2.2 ANR . . . 24
10.3 Public policy support . . . 24

11 Dissemination **24**
11.1 Scientific events: organisation . . . 24
11.2 Scientific expertise . . . 25
11.3 Research administration . . . 25
11.4 Teaching . . . 25
11.5 Popularization . . . 26

12 Scientific production **26**
12.1 Major publications . . . 26
12.2 Publications of the year . . . 27

1 Team members, visitors, external collaborators

Research Scientists

- Bruno Raffin [Team leader, INRIA, Senior Researcher, HDR]
- Carlos Jaime Barrios Hernandez [INRIA, Advanced Research Position]
- Christophe Cerin [UNIV PARIS, until Aug 2025]
- Fanny Dufosse [INRIA, Researcher]
- Bertrand Simon [CNRS, Researcher, from Sep 2025]

Faculty Members

- Danilo Carastan Dos Santos [UGA, Associate Professor]
- Christophe Cerin [UNIV PARIS, Professor Delegation, from Sep 2025]
- Yves Denneulin [GRENOBLE INP, Professor, HDR]
- Pierre Dutot [UGA, Associate Professor]
- Grégory Mounié [GRENOBLE INP, Associate Professor]
- Kim Thang Nguyen [GRENOBLE INP, Professor]
- Olivier Richard [GRENOBLE INP, Associate Professor Delegation, from Sep 2025]
- Olivier Richard [GRENOBLE INP, Associate Professor, until Aug 2025]
- Denis Trystram [GRENOBLE INP, Professor, from Sep 2025, HDR]
- Denis Trystram [GRENOBLE INP, Professor Delegation, until Aug 2025, HDR]
- Frederic Wagner [GRENOBLE INP, Associate Professor]
- Philippe Waille [UGA, Associate Professor, from Feb 2025]

Post-Doctoral Fellow

- Aina Rasoldier [FLORALIS, Post-Doctoral Fellow, until Jun 2025]

PhD Students

- Abdessalam Benhari [ATOS, CIFRE, until Apr 2025]
- Jad Berjawi [UGA, CIFRE, from Oct 2025]
- Louis Boulanger [UGA, from Apr 2025 until Aug 2025]
- Louis Boulanger [INRIA, until Mar 2025]
- Louis Closson [ADEUNIS, CIFRE, until Mar 2025]
- Wenke Du [INRIA]
- Yoann Dupas [ORANGE, CIFRE, until Nov 2025]
- Sofya Dymchenko [INRIA]
- Dorian Goepp [UGA]

- Marina Gradvohl [SCHNEIDER ELECTRIC, CIFRE]
- Eniko Kevi [UGA, until Jan 2025]
- Yannick Malot [CEA]
- Guillaume Raffin [BULL, CIFRE]
- Hamza Safri [BERGER-LEVRAULT, CIFRE, until Feb 2025]
- Theo Seigneuret-Poussard [ORANGE, CIFRE]
- Yifei Sun [INRIA, from Jul 2025]
- Valentin Trophime-Gilotte [INRIA]

Technical Staff

- Fernando Ayats Llamas [INRIA, Engineer, until Aug 2025]
- Louis Beal [INRIA, Engineer]
- Andres Bermeo Marinelli [INRIA, Engineer]
- Pierre Cesar [INRIA, Engineer, from Nov 2025]
- Dominik Huber [UGA, Engineer, until Mar 2025]
- Pierre Neyron [CNRS, Engineer]
- Abhishek Purandare [INRIA, Engineer]
- Colin Regal-Mezin [INRIA, Engineer, from Feb 2025]
- Djoser Simeu [INPG SA, from Oct 2025]
- Hugo Strappazzon [INRIA, Engineer, from Apr 2025]

Interns and Apprentices

- Jad Berjawi [INRIA, Intern, from Feb 2025 until Aug 2025]
- Einar Bratthall [INRIA, Intern, from May 2025 until Jun 2025]
- Einar Bratthall [INRIA, Intern, until Apr 2025]
- Pierre Cesar [INRIA, Intern, from Apr 2025 until Sep 2025]
- Scott Douanla Meli [INRIA, Intern, from May 2025 until Sep 2025]
- Emile Dugelay [INRIA, Intern, from Feb 2025 until Jul 2025]
- Jules Dupuis [INRIA, Intern, from May 2025 until Jul 2025]
- Jules Dupuis [INRIA, Intern, from Feb 2025 until Apr 2025]
- Clement Grennerat [INRIA, Intern, from Aug 2025 until Sep 2025]
- Clement Grennerat [INRIA, Intern, from Jun 2025 until Aug 2025]
- Paul Kailer [INRIA, Intern, from Feb 2025 until Jul 2025]
- Luiz Felipe Mascarenhas Dalle Nery [INRIA, Intern, from May 2025 until Jul 2025]
- Luiz Felipe Mascarenhas Dalle Nery [INRIA, Intern, until Apr 2025]

- Louka Moroni [INRIA, Intern, from May 2025 until Jul 2025]
- Louka Moroni [INRIA, Intern, until Apr 2025]
- Matteo Rossillol–Laruelle [INRIA, Intern, from May 2025 until Jul 2025]
- Gabriella Silva Saraiva [GOUV BRESIL, Intern, from Apr 2025 until May 2025]
- Djoser Simeu [UGA, Intern, from Feb 2025 until Jul 2025]
- Adrien Vannson [ENS DE LYON, Intern, from Feb 2025 until Aug 2025]

Administrative Assistants

- Luce Coelho [INRIA]
- Annie Simon [INRIA]

2 Overall objectives

Moving data on large supercomputers is becoming a major performance bottleneck, and the situation is expected to worsen even more at exascale and beyond. Data transfer capabilities are growing at a slower rate than processing power ones. The profusion of flops available will be difficult to use efficiently due to constrained communication capabilities. Moving data is also an important source of power consumption. The DataMove team focuses on **data aware large scale computing**, investigating approaches to reduce data movements on large scale HPC machines. We will investigate data aware scheduling algorithms for job management systems. The growing cost of data movements requires adapted scheduling policies able to take into account the influence of intra-application communications, IOs as well as contention caused by data traffic generated by other concurrent applications. At the same time experimenting new scheduling policies on real platforms is unfeasible. Simulation tools are required to probe novel scheduling policies. Our goal is to investigate how to extract information from actual compute centers traces in order to replay job allocations and executions with new scheduling policies. Schedulers need information about the jobs behavior on the target machine to actually make efficient allocation decisions. We will research approaches relying on learning techniques applied to execution traces to extract data and forecast job behaviors. In addition to traditional computation intensive numerical simulations, HPC platforms also need to execute more and more often data intensive processing tasks like data analysis. In particular, the ever growing amount of data generated by numerical simulation calls for a tighter integration between the simulation and the data analysis. The goal is to reduce the data traffic and to speed-up result analysis by processing results in-situ, i.e. as closely as possible to the locus and time of data generation. Our goal is here to investigate how to program and schedule such analysis workflows in the HPC context, requiring the development of adapted resource sharing strategies, data structures and parallel analytics schemes. To tackle these issues, we will intertwine theoretical research and practical developments to elaborate solutions generic and effective enough to be of practical interest. Algorithms with performance guarantees will be designed and experimented on large scale platforms with realistic usage scenarios developed with partner scientists or based on logs of the biggest available computing platforms. Conversely, our strong experimental expertise will enable to feed theoretical models with sound hypotheses, to twist proven algorithms with practical heuristics that could be further retro-feeded into adequate theoretical models.

3 Research program

3.1 Motivation

Today's largest supercomputers are composed of few millions of cores, with performances reaching 1 ExaFlops [1] for the largest machines. Moving data in such large supercomputers is becoming a major

[1] 10^{18} floating point operations per second

performance bottleneck, and the situation is expected to worsen even more at exascale and beyond. The data transfer capabilities are growing at a slower rate than processing power ones. The profusion of available flops will very likely be underused due to constrained communication capabilities. It is commonly admitted that data movements account for 50% to 70% of the global power consumption. Thus, data movements are potentially one of the most important source of savings for enabling supercomputers to stay in the commonly adopted energy barrier of 20 MegaWatts. In the mid to long term, non volatile memory (NVRAM) is expected to deeply change the machine I/Os. Data distribution will shift from disk arrays with an access time often considered as uniform, towards permanent storage capabilities at each node of the machine, making data locality an even more prevalent paradigm.

The proposed DataMove team will work on **optimizing data movements for large scale computing** mainly at two related levels:

- Resource allocation
- Integration of numerical simulation and data analysis

The resource and job management system (also called batch scheduler or RJMS) is in charge of allocating resources upon user requests for executing their parallel applications. The growing cost of data movements requires adapted scheduling policies able to take into account the influence of intra-application communications, I/Os as well as contention caused by data traffic generated by other concurrent applications. Modelling the application behavior to anticipate its actual resource usage on such architecture is known to be challenging, but it becomes critical for improving performances (execution time, energy, or any other relevant objective). The job management system also needs to handle new types of workloads: high performance platforms now need to execute more and more often data intensive processing tasks like data analysis in addition to traditional computation intensive numerical simulations. In particular, the ever growing amount of data generated by numerical simulation calls for a tighter integration between the simulation and the data analysis. The challenge here is to reduce data traffic and to speed-up result analysis by performing result processing (compression, indexation, analysis, visualization, etc.) as closely as possible to the locus and time of data generation. This emerging trend called *in-situ analytics* requires to revisit the traditional workflow (loop of batch processing followed by postmortem analysis). The application becomes a whole including the simulation, in-situ processing and I/Os. This motivates the development of new well-adapted resource sharing strategies, data structures and parallel analytics schemes to efficiently interleave the different components of the application and globally improve the performance.

3.2 Strategy

DataMove targets HPC (High Performance Computing) at Exascale. But such machines and the associated applications are expected to be available only in 5 to 10 years. Meanwhile, we expect to see a growing number of petaflop machines to answer the needs for advanced numerical simulations. A sustainable exploitation of these petaflop machines is a real and hard challenge that we will address. We may also see in the coming years a convergence between HPC and Big Data, HPC platforms becoming more elastic and supporting Big Data jobs, or HPC applications being more commonly executed on cloud like architectures. We will contribute to that convergence at our level, considering more dynamic and versatile target platforms and types of workloads.

Our approaches should entail minimal modifications on the code of numerical simulations. Often large scale numerical simulations are complex domain specific codes with a long life span. We assume these codes as being sufficiently optimized. We will influence the behavior of numerical simulations through resource allocation at the job management system level or when interleaving them with analytics code.

To tackle these issues, we propose to intertwine theoretical research and practical developments in an agile mode. Algorithms with performance guarantees will be designed and experimented on large scale platforms with realistic usage scenarios developed with partner scientists or based on logs of the biggest available computing platforms (national supercomputers like Curie, or the BlueWaters machine accessible through our collaboration with Argonne National Lab). Conversely, a strong experimental expertise will enable to feed theoretical models with sound hypotheses, to twist proven algorithms with practical heuristics that could be further retro-feeded into adequate theoretical models.

A central scientific question is to make the relevant choices for optimizing performance (in a broad sense) in a reasonable time. HPC architectures and applications are increasingly complex systems (heterogeneity, dynamicity, uncertainties), which leads to consider the **optimization of resource allocation based on multiple objectives**, often contradictory (like energy and run-time for instance). Focusing on the optimization of one particular objective usually leads to worsen the others. The historical positioning of some members of the team who are specialists in multi-objective optimization is to generate a (limited) set of trade-off configurations, called *Pareto points*, and choose when required the most suitable trade-off between all the objectives. This methodology differs from the classical approaches, which simplify the problem into a single objective one (focus on a particular objective, combining the various objectives or agglomerate them). The real challenge is thus to combine algorithmic techniques to account for this diversity while guaranteeing a target efficiency for all the various objectives.

The DataMove team aims to elaborate generic and effective solutions of practical interest. We will make our new algorithms accessible through the team flagship software tools, **the OAR batch scheduler and the Ensemble run online data processing framework Melissa**. We will maintain and enforce strong links with teams closely connected with large architecture design and operation (CEA DAM, BULL, Argonne National Lab), as well as scientists of other disciplines, in particular computational biologists, with whom we will elaborate and validate new usage scenarios (IBPC, CEA DAM, EDF).

3.3 Research Directions

DataMove research activity is organized around three directions:

1. When a parallel job executes on a machine, it triggers data movements through the input data it needs to read, the results it produces (simulation results as well as traces) that need to be stored in the file system, as well as internal communications and temporary storage (for fault tolerance related data for instance). Modeling in details the simulation and the target machines to analyze scheduling policies is not feasible at large scales. We propose to investigate alternative approaches, including learning approaches, to capture and model the influence of data movements on the performance metrics of each job execution to develop **Data Aware Batch Scheduling** models and algorithms (Sec. 4.1).

2. Experimenting new scheduling policies on real platforms at scale is unfeasible. Theoretical performance guarantees are not sufficient to ensure a new algorithm will actually perform as expected on a real platform. An intermediate evaluation level is required to probe novel scheduling policies. The second research axe focuses on the **Empirical Studies of Large Scale Platforms** (Sec. 4.2). The goal is to investigate how we could extract from actual computing centers traces information to replay the job allocations and executions on a simulated or emulated platform with new scheduling policies. Schedulers need information about jobs behavior on target machines to actually be able to make efficient allocation decisions. Asking users to characterize jobs often does not lead to reliable information.

3. The third research direction **Integration of High Performance Computing and Data Analytics** (Sec. 4.3) addresses the data movement issue from a different perspective. New data analysis techniques on the HPC platform introduce new type of workloads, potentially more data than compute intensive, but could also enable to reduce data movements by directly enabling to pipe-line simulation execution with a live (in situ) analysis of the produced results. Our goal is here to investigate how to program and schedule such analysis workflows in the HPC context.

4 Application domains

4.1 Data Aware Batch Scheduling

Large scale high performance computing platforms are becoming increasingly complex. Determining efficient allocation and scheduling strategies that can adapt to technological evolutions is a strategic and difficult challenge. We are interested in scheduling jobs in hierarchical and heterogeneous large scale platforms. On such platforms, application developers typically submit their jobs in centralized waiting queues. The job management system aims at determining a suitable allocation for the jobs, which all compete against each other for the available computing resources. Performances are measured using different classical metrics like maximum completion time or slowdown. Current systems make use of very simple (but fast) algorithms that however rely on simplistic platform and execution models, and thus, have limited performances.

For all target scheduling problems we aim to provide both theoretical analysis and complementary analysis through simulations. Achieving meaningful results will require strong improvements on existing models (on power for example) and the design of new approximation algorithms with various objectives such as stretch, reliability, throughput or energy consumption, while keeping in focus the need for a low-degree polynomial complexity.

4.1.1 Algorithms

The most common batch scheduling policy is to consider the jobs according to the First Come First Served order (FCFS) with backfilling (BF). BF is the most widely used policy due to its easy and robust implementation and known benefits such as high system utilization. It is well-known that this strategy does not optimize any sophisticated function, but it is simple to implement and it guarantees that there is no starvation (i.e. every job will be scheduled at some moment).

More advanced algorithms are seldom used on production platforms due to both the gap between theoretical models and practical systems and speed constraints. When looking at theoretical scheduling problems, the generally accepted goal is to provide polynomial algorithms (in the number of submitted jobs and the number of involved computing units). However, with millions of processing cores where every process and data transfer have to be individually scheduled, polynomial algorithms are prohibitive as soon as the polynomial degree is too large. The model of *parallel tasks* simplifies this problem by bundling many threads and communications into single boxes, either rigid, rectangular or malleable. Especially malleable tasks capture the dynamicity of the execution. Yet these models are ill-adapted to heterogeneous platforms, as the running time depends on more than simply the number of allotted resources, and some of the common underlying assumptions on the speed-up functions (such as monotony or concavity) are most often only partially verified.

In practice, the job execution times depend on their allocation (due to communication interferences and heterogeneity in both computation and communication), while theoretical models of parallel jobs usually consider jobs as black boxes with a fixed (maximum) execution time. Though interesting and powerful, the classical models (namely, synchronous PRAM model, delay, LogP) and their variants (such as hierarchical delay), are not well-suited to large scale parallelism on platforms where the cost of moving data is significant, non uniform and may change over time. Recent studies are still refining such models in order to take into account communication contentions more accurately while remaining tractable enough to provide a useful tool for algorithm design.

Today, all algorithms in use in production systems are oblivious to communications. One of our main goals is to **design a new generation of scheduling algorithms fitting more closely job schedules according to platform topologies**.

4.1.2 Locality Aware Allocations

Recently, we developed modifications of the standard back-filling algorithm taking into account platform topologies. The proposed algorithms take into account locality and contiguity in order to hide communication patterns within parallel tasks. The main result here is to establish good lower bounds and small approximation ratios for policies respecting the locality constraints. The algorithms work in an online fashion, improving the global behavior of the system while still keeping a low running time. These improvements rely mainly on our past experience in designing approximation algorithms. Instead of relying on complex networking models and communication patterns for estimating execution times, the communications are disconnected from the execution time. Then, the scheduling problem leads to a trade-off: optimizing locality of communications on one side and a performance objective (like the makespan or stretch) on the other side.

In the perspective of taking care of locality, other ongoing works include the study of schedulers for platforms whose interconnection network is a static structured topology (like the 3D-torus of the BlueWaters platform we work on in collaboration with the Argonne National Laboratory). One main characteristic of this 3D-torus platform is to provide I/O nodes at specific locations in the topology. Applications generate and access specific data and are thus bounded to specific I/O nodes. Resource allocations are constrained in a strong and unusual way. This problem is close for actual hierarchical platforms. The scheduler needs to compute a schedule such that I/O nodes requirements are filled for each application while at the same time avoiding communication interferences. Moreover, extra constraints can arise for applications requiring accelerators that are gathered on the nodes at the edge of the network topology.

While current results are encouraging, they are however limited in performance by the low amount of information available to the scheduler. We look forward to extend ongoing work by progressively increasing application and network knowledge (by technical mechanisms like profiling or monitoring or by more sophisticated methods like learning). It is also important to anticipate on application resource usage in terms of compute units, memory as well as network and I/Os to efficiently schedule a mix of applications with different profiles. For instance, a simple solution is to partition the jobs as "communication intensive" or "low communications". Such a tag could be achieved by the users them selves or obtained by learning techniques. We could then schedule low communications jobs using leftover spaces while taking care of high communication jobs. More sophisticated options are possible, for instance those that use more detailed communication patterns and networking models. Such options would leverage the work proposed in Section 4.2 for gathering application traces.

4.1.3 Data-Centric Processing

Exascale computing is shifting away from the traditional compute-centric models to a more data-centric one. This is driven by the evolving nature of large scale distributed computing, no longer dominated by pure computations but also by the need to handle and analyze large volumes of data. These data can be large databases of results, data streamed from a running application or another scientific instrument (collider for instance). These new workloads call for specific resource allocation strategies.

Data movements and storage are expected to be a major energy and performance bottleneck on next generation platforms. Storage architectures are also evolving, the standard centralized parallel file system being complemented with local persistent storage (Burst Buffers, NVRAM). Thus, one data producer can stage data on some nodes' local storage, requiring to schedule close by the associated analytics tasks to limit data movements. This kind of configuration, often referred as *in-situ analytics*, is expected to become common as it enables to switch from the traditional I/O intensive workflow (batch-processing followed by *post mortem* analysis and visualization) to a more storage conscious approach where data are processed as closely as possible to where and when they are produced (in-situ processing is addressed in details in section 4.3). By reducing data movements and scheduling the extra processing on resources not fully exploited yet, in-situ processing is expected to have also a significant positive energetic impact. Analytics codes can be executed in the same nodes than the application, often on dedicated cores commonly called helper cores, or on dedicated nodes called stagging nodes. The results are either forwarded to the users for visualization or saved to disk through I/O nodes. In-situ analytics can also take benefit of node local disks or burst buffers to reduce data movements. Future job scheduling strategies should take into account in-situ processes in addition to the job allocation to optimize both energy consumption and execution time. On the one hand, this problem can be reduced to an allocation problem of extra asynchronous tasks to idle computing units. But on the other hand, embedding analytics in applications brings extra difficulties by making the application more heterogeneous and imposing more constraints (data affinity) on the required resources. Thus, the main point here is to develop efficient algorithms for dealing with heterogeneity without increasing the global computational cost.

4.1.4 Learning

Another important issue is to adapt the job management system to deal with the bad effects of uncertainties, which may be catastrophic in large scale heterogeneous HPC platforms (jobs delayed arbitrarly far or jobs killed). A natural question is then: *is it possible to have a good estimation of the job and platform parameters in order to be able to obtain a better scheduling ?* Many important parameters (like the number or type of required resources or the estimated running time of the jobs) are asked to the users when they submit their jobs. However, some of these values are not accurate and in many cases, they are not even provided by the end-users. In DataMove, we propose to study new methods for a better prediction of the characteristics of the jobs and their execution in order to improve the optimization process. In particular, the methods well-studied in the field of big data (in supervised Machine Learning, like classical regression methods, Support Vector Methods, random forests, learning to rank techniques or deep learning) could and must be used to improve job scheduling in large scale HPC platforms. This topic received a great attention recently in the field of parallel and distributed processing. A preliminary study has been done recently by our team with the target of predicting the job running times (called wall times). We succeeded to improve significantly in average the reference EASY Back Filling algorithm by estimating the wall time of the jobs, however, this method leads to big delay for the stretch of few jobs. Even if we succeed in determining more precisely hidden parameters, like the wall time of the jobs, this is not enough to determine an optimized solution. The shift is not only to learn on dedicated parameters but also on the scheduling policy. The data collected from the accounting and profiling of jobs can be used to better understand the needs of the jobs and through learning to propose adaptations for future submissions. The goal is to propose extensions to further improve the job scheduling and improve the performance and energy efficiency of the application. For instance preference learning may enable to compute on-line new priorities to back-fill the ready jobs.

4.1.5 Multi-objective Optimization

Several optimization questions that arise in allocation and scheduling problems lead to the study of several objectives at the same time. The goal is then not a single optimal solution, but a more complicated mathematical object that captures the notion of trade-off. In broader terms, the goal of multi-objective optimization is not to externally arbitrate on disputes between entities with different goals, but rather to explore the possible solutions to highlight the whole range of interesting compromises. A classical tool for studying such multi-objective optimization problems is to use *Pareto curves*. However, the full description of the Pareto curve can be very hard because of both the number of solutions and the hardness of computing each point. Addressing this problem will opens new methodologies for the analysis of algorithms.

To further illustrate this point here are three possible case studies with emphasis on conflicting interests measured with different objectives. While these cases are good representatives of our HPC context, there are other pertinent trade-offs we may investigate depending on the technology evolution in the coming years. This enumeration is certainly not limitative.

Energy versus Performance. The classical scheduling algorithms designed for the purpose of performance can no longer be used because performance and energy are contradictory objectives to some extent. The scheduling problem with energy becomes a multi-objective problem in nature since the energy consumption should be considered as equally important as performance at exascale. A global constraint on energy could be a first idea for determining trade-offs but the knowledge of the Pareto set (or an approximation of it) is also very useful.

Administrators versus application developers. Both are naturally interested in different objectives: In current algorithms, the performance is mainly computed from the point of view of administrators, but the users should be in the loop since they can give useful information and help to the construction of better schedules. Hence, we face again a multi-objective problem where, as in the above case, the approximation of the Pareto set provides the trade-off between the administrator view and user demands. Moreover, the objectives are usually of the same nature. For example, *max stretch* and *average stretch* are two objectives based on the slowdown factor that can interest administrators and users, respectively. In this case the study of the norm of stretch can be also used to describe the trade-off (recall that the L_1-norm corresponds to the average objective while the L_∞-norm to the max objective). Ideally, we would like to design an algorithm that gives good approximate solutions at the same time for all norms. The L_2 or L_3-norm are useful since they describe the performance of the whole schedule from the administrator point of view as well as they provide a fairness indication to the users. The hard point here is to derive theoretical analysis for such complicated tools.

In general, resource augmentation can explain the intuitive good behavior of some greedy algorithms while, more interestingly, it can give ideas for new algorithms. For example, in the rejection context we could dedicate a small number of nodes for the usually problematic rejected jobs. Some initial experiments show that this can lead to a schedule for the remaining jobs that is very close to the optimal one.

4.2 Empirical Studies of Large Scale Platforms

Experiments or realistic simulations are required to take into account the impact of allocations and assess the real behavior of scheduling algorithms. While theoretical models still have their interest to lay the groundwork for algorithmic designs, the models are necessarily reflecting a purified view of the reality. As transferring our algorithm in a more practical setting is an important part of our creed, we need to ensure that the theoretical results found using simplified models can really be transposed to real situations. On the way to exascale computing, large scale systems become harder to study, to develop or to calibrate because of the costs in both time and energy of such processes. It is often impossible to convince managers to use a production cluster for several hours simply to test modifications in the RJMS. Moreover, as the existing RJMS production systems need to be highly reliable, each evolution requires several real scale test iterations. The consequence is that scheduling algorithms used in production systems are mostly outdated and not customized correctly. To circumvent this pitfall, we need to develop tools and methodologies for alternative empirical studies, from analysis of workload traces, to job models, simulation and emulation with reproducibility concerns.

4.2.1 Workload Traces with Resource Consumption

Workload traces are the base element to capture the behavior of complete systems composed of submitted jobs, running applications, and operating tools. These traces must be obtained on production platforms to provide relevant and representative data. To get a better understanding of the use of such systems, we need to look at both, how the jobs interact with the job management system, and how they use the allocated resources. We propose a general workload trace format that adds jobs resource consumption to the commonly used Standard Workload Format workload trace format. This requires to instrument the platforms, in particular to trace resource consumptions like CPU, data movements at memory, network and I/O levels, with an acceptable performance impact. In a previous work we studied and proposed a dedicated job monitoring tool whose impact on the system has been measured as lightweight (0.35% speed-down) with a 1 minute sampling rate. Other tools also explore job monitoring, like TACC Stats. A unique feature from our tool is its ability to monitor distinctly jobs sharing common nodes.

Collected workload traces with jobs resource consumption will be publicly released and serve to provide data for works presented in Section 4.1. The trace analysis is expected to give valuable insights to define models encompassing complex behaviours like network topology sensitivity, network congestion and resource interferences.

4.2.2 Simulation

Simulations of large scale systems are faster by multiple orders of magnitude than real experiments. Unfortunately, replacing experiments with simulations is not as easy as it may sound, as it brings a host of new problems to address in order to ensure that the simulations are closely approximating the execution of typical workloads on real production clusters. Most of these problems are actually not directly related to scheduling algorithms assessment, in the sense that the workload and platform models should be defined independently from the algorithm evaluations, in order to ensure a fair assessment of the algorithms' strengths and weaknesses. These research topics (namely platform modeling, job models and simulator calibration) are addressed in the other subsections.

We developed an open source platform simulator within DataMove (in conjunction with the OAR development team) to provide a widely distributable test bed for reproducible scheduling algorithm evaluation. Our simulator, named Batsim, allows to simulate the behavior of a computational platform executing a workload scheduled by any given scheduling algorithm. To obtain sound simulation results and to broaden the scope of the experiments that can be done thanks to Batsim, we did not chose to create a (necessarily limited) simulator from scratch, but instead to build on top of the SimGrid simulation framework.

To be open to as many batch schedulers as possible, Batsim decouples the platform simulation and the scheduling decisions in two clearly-separated software components communicating through a complete and documented protocol. The Batsim component is in charge of simulating the computational resources behaviour whereas the scheduler component is in charge of taking scheduling decisions. The scheduler component may be both a resource and a job management system. For jobs, scheduling decisions can be to execute a job, to delay its execution or simply to reject it. For resources, other decisions can be taken, for example to change the power state of a machine i.e. to change its speed (in order to lower its energy consumption) or to switch it on or off. This separation of concerns also enables interfacing with potentially any commercial RJMS, as long as the communication protocol with Batsim is implemented. A proof of concept is already available with the OAR RJMS.

Using this test bed opens new research perspectives. It allows to test a large range of platforms and workloads to better understand the real behavior of our algorithms in a production setting. In turn, this opens the possibility to tailor algorithms for a particular platform or application, and to precisely identify the possible shortcomings of the theoretical models used.

4.2.3 Job and Platform Models

The central purpose of the Batsim simulator is to simulate job behaviors on a given target platform under a given resource allocation policy. Depending on the workload, a significant number of jobs are parallel applications with communications and file system accesses. It is not conceivable to simulate individually all these operations for each job on large plaforms with their associated workload due to implied simulation complexity. The challenge is to define a coarse grain job model accurate enough to reproduce parallel application behavior according to the target platform characteristics. We will explore models similar to the BSP (Bulk Synchronous Program) approach that decomposes an application in local computation supersteps ended by global communications and a global synchronization. The model parameters will be established by means of trace analysis as discussed previously, but also by instrumenting some parallel applications to capture communication patterns. This instrumentation will have a significant impact on the concerned application performance, restricting its use to a few applications only. There are a lot of recurrent applications executed on HPC platform, this fact will help to reduce the required number of instrumentations and captures. To assign each job a model, we are considering to adapt the concept of application signatures as proposed in. Platform models and their calibration are also required. Large parts of these models, like those related to network, are provided by Simgrid. Other parts as the filesystem and energy models are comparatively recent and will need to be enhanced or reworked to reflect the HPC platform evolutions. These models are then generally calibrated by running suitable benchmarks.

4.2.4 Emulation and Reproducibility

The use of coarse models in simulation implies to set aside some details. This simplification may hide system behaviors that could impact significantly and negatively the metrics we try to enhance. This issue is particularly relevant when large scale platforms are considered due to the impossibility to run tests at nominal scale on these real platforms. A common approach to circumvent this issue is the use of emulation techniques to reproduce, under certain conditions, the behavior of large platforms on smaller ones. Emulation represents a natural complement to simulation by allowing to execute directly large parts of the actual evaluated software and system, but at the price of larger compute times and a need for more resources. The emulation approach was chosen in to compare two job management systems from workload traces of the CURIE supercomputer (80000 cores). The challenge is to design methods and tools to emulate with sufficient accuracy the platform and the workload (data movement, I/O transfers, communication, applications interference). We will also intend to leverage emulation tools like Distem from the MADYNES team. It is also important to note that the Batsim simulator also uses emulation techniques to support the core scheduling module from actual RJMS. But the integration level is not the same when considering emulation for larger parts of the system (RJMS, compute node, network and filesystem).

Replaying traces implies to prepare and manage complex software stacks including the OS, the resource management system, the distributed filesystem and the applications as well as the tools required to conduct experiments. Preparing these stacks generate specific issues, one of the major one being the support for reproducibility. We propose to further develop the concept of reconstructability to improve experiment reproducibility by capturing the build process of the complete software stack. This approach ensures reproducibility over time better than other ways by keeping all data (original packages, build recipe and Kameleon engine) needed to build the software stack.

In this context, the Grid'5000 (see Sec. 7.2) experimentation infrastructure that gives users the control on the complete software stack is a crucial tool for our research goals. We will pursue our strong implication in this infrastructure.

4.3 Integration of High Performance Computing and Data Analytics

Data produced by large simulations are traditionally handled by an I/O layer that moves them from the compute cores to the file system. Analysis of these data are performed after reading them back from files, using some domain specific codes or some scientific visualisation libraries like VTK. But writing and then reading back these data generates a lot of data movements and puts under pressure the file system. To reduce these data movements, **the in situ analytics paradigm proposes to process the data as closely as possible to where and when the data are produced**. Some early solutions emerged either as extensions of visualisation tools or of I/O libraries like ADIOS. But significant progresses are still required to provide efficient and flexible high performance scientific data analysis tools. Integrating data analytics in the HPC context will have an impact on resource allocation strategies, analysis algorithms, data storage and access, as well as computer architectures and software infrastructures. But this paradigm shift imposed by the machine performance also sets the basis for a deep change on the way users work with numerical simulations. The traditional workflow needs to be reinvented to make HPC more user-centric, more interactive and turn HPC into a commodity tool for scientific discovery and engineering developments. In this context DataMove aims at investigating programming environments for in situ analytics with a specific focus on task scheduling in particular, to ensure an efficient sharing of resources with the simulation.

4.3.1 Programming Model and Software Architecture

In situ creates a tighter loop between the scientist and her/his simulation. As such, an in situ framework needs to be flexible to let the user define and deploy its own set of analysis. A manageable flexibility requires to favor simplicity and understandability, while still enabling an efficient use of parallel resources. Visualization libraries like VTK or Visit, as well as domain specific environments like VMD have initially been developed for traditional post-mortem data analysis. They have been extended to support in situ processing with some simple resource allocation strategies but the level of performance, flexibility and ease of use that is expected requires to rethink new environments. There is a need to develop a middleware and programming environment taking into account in its fundations this specific context of high performance scientific analytics.

Similar needs for new data processing architectures occurred for the emerging area of Big Data Analytics, mainly targeted to web data on cloud-based infrastructures. Google Map/Reduce and its successors like Spark or Stratosphere/Flink have been designed to match the specific context of efficient analytics for large volumes of data produced on the web, on social networks, or generated by business applications. These systems have mainly been developed for cloud infrastructures based on commodity architectures. They do not leverage the specifics of HPC infrastructures. Some preliminary adaptations have been proposed for handling scientific data in a HPC context. However, these approaches do not support in situ processing.

Following the initial development of FlowVR, our middleware for in situ processing, we will pursue our effort to develop a programming environment and software architecture for high performance scientific data analytics. Like FlowVR, the map/reduce tools, as well as the machine learning frameworks like TensorFlow, adopted a dataflow graph for expressing analytics pipe-lines. We are convinced that this dataflow approach is both easy to understand and yet expresses enough concurrency to enable efficient executions. The graph description can be compiled towards lower level representations, a mechanism that is intensively used by Stratosphere/Flink for instance. Existing in situ frameworks inherit from the HPC way of programming with a thiner software stack and a programming model close to the machine. Though this approach enables to program high performance applications, this is usually too low level to enable the scientist to write its analysis pipe-line in a short amount of time. The data model, i.e. the data semantics level accessible at the framework level for error check and optimizations, is also a fundamental aspect of such environments. The key/value store has been adopted by all map/reduce tools. Except in some situations, it cannot be adopted as such for scientific data. Results from numerical simulations are often more structured than web data, associated with acceleration data structures to be processed efficiently. We will investigate data models for scientific data building on existing approaches like Adios or DataSpaces.

4.3.2 Resource Sharing

To alleviate the I/O bottleneck, the in situ paradigm proposes to start processing data as soon as made available by the simulation, while still residing in the memory of the compute node. In situ processings include data compression, indexing, computation of various types of descriptors (1D, 2D, images, etc.). Per se, reducing data output to limit I/O related performance drops or keep the output data size manageable is not new. Scientists have relied on solutions as simple as decreasing the frequency of result savings. In situ processing proposes to move one step further, by providing a full fledged processing framework enabling scientists to more easily and thoroughly manage the available I/O budget.

The most direct way to perform in situ analytics is to inline computations directly in the simulation code. In this case, in situ processing is executed in sequence with the simulation that is suspended meanwhile. Though this approach is direct to implement and does not require complex framework environments, it does not enable to overlap analytics related computations and data movements with the simulation execution, preventing to efficiently use the available resources. Instead of relying on this simple time sharing approach, several works propose to rely on space sharing where one or several cores per node, called *helper cores*, are dedicated to analytics. The simulation responsibility is simply to handle a copy of the relevant data to the node-local in situ processes, both codes being executed concurrently. This approach often lead to significantly beter performance than in-simulation analytics.

For a better isolation of the simulation and in situ processes, one solution consists in offloading in situ tasks from the simulation nodes towards extra dedicated nodes, usually called *staging nodes*. These computations are said to be performed *in-transit*. But this approach may not always be beneficial compared to processing on simulation nodes due to the costs of moving the data from the simulation nodes to the staging nodes.

But today the choice of the resource allocation strategy is mostly ad-hoc and defined by the programmer. We will investigate solutions that enable a cooperative use of the resource between the analytics and the simulation with minimal hints from the programmer. In situ processings inherit from the parallelization scale and data distribution adopted by the simulation, and must execute with minimal perturbations on the simulation execution (whose actual resource usage is difficult to know a priori). We need to develop adapted scheduling strategies that operate at compile and run time. Because analysis are often data intensive, such solutions must take into consideration data movements, a point that classical scheduling strategies designed first for compute intensive applications often overlook. We expect to develop new scheduling strategies relying on the methodologies developed in Sec. 4.1.5. Simulations as well as analysis are iterative processes exposing a strong spatial and temporal coherency that we can take benefit of to anticipate their behavior and then take more relevant resources allocation strategies, possibly based on advanced learning algorithms or as developed in Section 4.1.

In situ analytics represent a specific workload that needs to be scheduled very closely to the simulation, but not necessarily active during the full extent of the simulation execution and that may also require to access data from previous runs (stored in the file system or on specific burst-buffers). Several users may also need to run concurrent analytics pipe-lines on shared data. This departs significantly from the traditional batch scheduling model, motivating the need for a more elastic approach to resource provisioning. These issues will be conjointly addressed with research on batch scheduling policies (Sec. 4.1).

4.3.3 Co-Design with Data Scientists

Given the importance of users in this context, it is of primary importance that in situ tools be co-designed with advanced users, even if such multidisciplinary collaborations are challenging and require constant long term investments to learn and understand the specific practices and expectations of the other domain.

We will tightly collaborate with scientists of some application domains, like molecular dynamics or fluid simulation, to design, develop, deploy and assess in situ analytics scenarios.

5 Social and environmental responsibility

DataMove is environmentally involved at different levels:

- Pursuing research on energy optimization of large scale distributed compute infrastructures
- Intend to include in publications the total amount of compute hours required for running all associated experiments, especially when using supercomputers, to, in a first step, get a measure of the impact of our experimentation activity.
- Lead and participate to different local LIG and INRIA groups in charge of evaluating, proposing and implementing solutions to limit our environmental impact in the lab.
- Take actions for lowering our carbon impact (extend laptop, smart phones, servers life to 6-8 years, favor fixing equipment rather then replacing them, put priority on train rather than plane)

- Bicycle is just our favorite, very low carbon, way for commuting.

6 Highlights of the year

- Bertrand Simon, CNRS junior researcher, joined the DataMove Team in September 2025.
- DataMove again lead the organisation of 2025 edition of the Journées sur la Recherche en Apprentissage Frugal, 26-27 November 2025, Grenoble.
- DataMove participated to the AFNOR "Frugal AI Framework" spec document.
- Carlos Barrios, long term visiting senior scientist at DataMove, defended is HDR "MultiScale-HPC Hybrid Architectures: Developing Computing Continuum Towards Sustainable Advanced Computing", June 6th, 2025.

7 Latest software developments, platforms, open data

7.1 Latest software developments

7.1.1 OAR

Keywords: HPC, Cloud, Clusters, Resource manager, Light grid

Scientific Description: This batch system is based on a database (PostgreSQL (preferred) or MySQL), a script language (Perl) and an optional scalable administrative tool (e.g. Taktuk). It is composed of modules which interact mainly via the database and are executed as independent programs. Therefore, formally, there is no API, the system interaction is completely defined by the database schema. This approach eases the development of specific modules. Indeed, each module (such as schedulers) may be developed in any language having a database access library.

Functional Description: OAR is a versatile resource and task manager (also called a batch scheduler) for HPC clusters, and other computing infrastructures (like distributed computing experimental testbeds where versatility is a key).

URL: `http://oar.imag.fr`

Contact: Olivier Richard

Participant: 3 anonymous participants

Partners: LIG, CNRS, Grid'5000, CIMENT, UAR GRICAD

7.1.2 MELISSA

Keywords: Sensitivity Analysis, HPC, Data assimilation, Exascale, AI4Science

Functional Description: Melissa is a middleware framework for on-line processing of data produced from large scale ensemble runs (parameter sweep data analysis) for sensibility analysis, data assimilation and deep surrogate training. Largest runs so far involved up to 30k core, executed 80 000 parallel simulations, and generated 288 TB of intermediate data that did not need to be stored on the file system. For deep surrogate training Melissa demonstrated it can significantly speed-up training on multiple GPUs by maintaining a very high GPU usage.

URL: `https://gitlab.inria.fr/melissa`

Publications: hal-04145897, hal-04213978, hal-04102400, hal-01383860, hal-01607479, hal-03017033, hal-03927612, hal-03842106

Contact: Bruno Raffin

Partner: Edf

7.1.3 NixOS-Compose

Keywords: Infrastructure software, Deployment, High performance computing, Distributed computing

Functional Description: NixOS-Compose simplifies the process of setting up ephemeral distributed systems by utilizing Nix's functional package management and NixOS's declarative configuration management. The tool facilitates testing, development, infrastructure prototyping, benchmarking, and advanced experiments in high-performance computing by providing easy and reproducible software stack deployment.

URL: https://gitlab.inria.fr/nixos-compose/nixos-compose

Publication: hal-03723771

Contact: Olivier Richard

Partners: LIG, CNRS, UGA

7.1.4 Batsim

Functional Description: BatSim is a Resource and Job Management System (RJMS) framework simulator based on SimGrid. It aims at taking into account platform's hardware capabilities and impacts in simulations. Also, schedulers parts are plugable through a comprehensive API and they are seen as external component of the framework.

Release Contributions: see https://batsim.readthedocs.io/en/latest/changelog.html

URL: https://batsim.readthedocs.io/en/latest/

Contact: Olivier Richard

7.1.5 Kameleon

Keyword: Engineering software systems

Functional Description: Kameleon is a simple but powerful tool to generate customized appliances. With Kameleon, you make your recipe that describes how to create step by step your own distribution. At start Kameleon is used to create custom kvm, docker, VirtualBox, ..., but as it is designed to be very generic you can probably do a lot more than that.

URL: http://kameleon.imag.fr/

Contact: Olivier Richard

Participant: an anonymous participant

Partner: Grid'5000

7.1.6 alumet

Name: ALUMET: unified measurement software

Keywords: Energy, Rust, Power monitoring, High performance computing, Performance measure

Functional Description: Alumet provides a generic measurement pipeline with three steps: poll measurement sources, transform the data, and write the result. It is designed to be able to ingest metrics from various sources without redundant work. Supported sources include RAPL domains, Nvidia's NVML, and Jetson INA sensors. The list of supported devices will quickly grow over time, thanks to the next feature of Alumet.

URL: https://alumet.dev/

Contact: Guillaume Raffin

Partner: Bull - Atos Technologies

7.2 New platforms

7.2.1 Slices-fr/Grid'5000 and Meso Center Gricad

We are very active in promoting the factorization of compute resources at a regional and national level. We have a three level implication, locally to maintain a pool of very flexible experimental machines (hundreds of cores), regionally through the GRICAD meso center, and nationally by contributing to the Slices-fr/Grid'5000 platform, our local resources being included in this platform. Olivier Richard is member of Slices-fr/Grid'5000 scientific committee. The OAR scheduler in particular is deployed on both infrastructures. DataMove is hosting several engineers dedicated to Grid'5000 support.

8 New results

Our research team has been actively contributing to multiple areas of computer science, with a particular focus on sustainable computing, high-performance computing, and artificial intelligence applications. Below is a summary of our recent scientific publications:

8.1 Multimodal Vision and Attention-Based Detection

The DataMove team has produced several contributions at the intersection of computer vision, multimodal sensing, and attention-based neural architectures, with a particular focus on robust pedestrian and vehicle detection in adverse conditions. These works explore early-fusion strategies across heterogeneous sensors and propose new encoder–decoder models that jointly optimize accuracy and inference efficiency [14, 23].

8.2 Energy, Carbon Footprint, and Sustainability in HPC and AI

8.2.1 Carbon Footprint of High-Performance Computing

A central research theme concerns the environmental impact of high-performance computing (HPC), ranging from system-level carbon emissions to device lifetimes and power-aware scheduling. One study analyzes the evolution of the carbon footprint of large-scale HPC systems by combining performance data with information on energy mixes and projected trajectories toward 2030 [20]. Moving beyond the traditional Top500 and Green500 perspectives, the work considers the entire life span of several major systems and derives a predictive model to estimate the contribution of HPC to global carbon emissions over the next five years. By incorporating the carbon intensity of electricity and long-term deployment patterns, this analysis provides a forward-looking view on how the HPC community may need to adapt architectures, locations, and operational practices.

The environmental footprint is further refined at smaller scales, for example in the context of networked sensor infrastructures embedded in electrical distribution boards [19]. In this line of work, an empirical study compares three scenarios: a baseline board without energy measurements, a board with wired Modbus RS485-based metering, and a board with IEEE 802.15.4 wireless metering. Using Product Environmental Profiles and comparative life-cycle assessment, the authors show that instrumented boards inevitably increase carbon emissions compared to the non-instrumented baseline, but also that the wireless solution can reduce the environmental impact by nearly 45% relative to the wired configuration. The analysis also underlines that current models for wireless devices may overestimate operational consumption by not fully accounting for duty-cycling capabilities, thereby motivating more accurate modeling of connected devices in future work.

Another contribution addresses the lifetime of processors and accelerators and its relationship with the environmental footprint of supercomputers [21]. The work emphasizes that the increasing demand for GPUs, particularly from AI workloads, has created strong pressure on hardware availability and replacement cycles. By modeling aging as a function of operating temperature and frequency, the authors propose node frequency reconfiguration and dedicated scheduling algorithms that aim to increase the total number of floating-point operations delivered by a machine before component failure. Simulation results indicate that appropriate frequency decisions can substantially raise the cumulative computational output of a system, at the cost of controlled performance trade-offs on individual jobs, and that such strategies remain effective under different, imperfect aging models.

Complementing these system-level approaches, a separate study introduces Alumet, a modular framework that standardizes the measurement of energy consumption across hardware and software stacks [16]. Alumet provides a generic pipeline to collect, transform, and export a wide variety of measurements, and is designed with a plugin system to support new environments and energy models without requiring major changes to the core framework. Experimental deployments on heterogeneous platforms show that Alumet can operate at higher acquisition frequencies while limiting monitoring overhead, and that it facilitates the development of energy estimation models in diverse contexts. By improving the accuracy and extensibility of energy measurements, Alumet underpins the broader objective of making energy-aware decisions in HPC and distributed systems.

8.2.2 Power-Aware Scheduling and Dynamic Resource Management

Energy and power constraints are also addressed from a scheduling and runtime management perspective. One article focuses on power-constrained HPC platforms and investigates how to predict workload power consumption and exploit these predictions in power-aware scheduling algorithms [10]. The proposed method combines lightweight, history-based prediction schemes with a scheduler inspired by EASY backfilling, and models power capping as a greedy knapsack-like optimization problem. Using logs from Marconi 100, a 980-node supercomputer, simulation results show that relatively simple prediction models can achieve sufficiently accurate workload power forecasts to reduce overall power consumption without degrading scheduling performance or quality of service.

Dynamic Resource Management (DRM) forms another major axis of research. One study examines how to bridge genericity and programmability for dynamic resources in HPC by interfacing the Dynamic Management of Resources API (DMR-API) with the Dynamic Processes with PSets (DPP) design principles [15]. The DMR-API provides an application-level abstraction that simplifies the integration of dynamic resources into iterative HPC applications, while DPP offers a generic, programming-model-agnostic approach to resource control at the system level. By combining both, the authors propose a methodology that retains the flexibility of DPP while reducing the coding effort required to exploit dynamic resource allocation. Experimental evaluations indicate that DRM can be effectively leveraged in realistic HPC environments with limited software changes, improving job throughput and system utilization.

8.2.3 Environmental Impact of Generative AI

Beyond HPC in a narrow sense, the DataMove team also investigates the sustainability of emerging digital services such as generative AI (Gen-AI). One article addresses the environmental impact of Gen-AI services through a life-cycle and measurement-based study of a Stable Diffusion image generation service [9]. The methodology explicitly differentiates between embodied impacts, related to the manufacturing and deployment of models and hardware, and operational impacts, associated with runtime energy use across data centers, networks, and user devices. The analysis demonstrates that, when Gen-AI is offered as a service, the cumulative impact of numerous terminals and communication networks becomes a significant component of its footprint, and that decarbonizing electricity alone is insufficient to render such services sustainable in the long term. By emphasizing constraints related to energy consumption and rare metals in a finite-resource world, the study argues for early and comprehensive impact assessments of Gen-AI solutions and provides tools to support more informed decisions about their deployment.

8.3 Computing Continuum: Architectures, Testbeds, and Complexity Management

8.3.1 Edge–Cloud and Serverless Scheduling

Several 2025 publications focus on the computing continuum, spanning edge to cloud resources and embracing serverless and container-based paradigms. One contribution, FOA-Energy, proposes a multi-objective scheduling policy for serverless platforms deployed across an edge–cloud continuum [13]. Recognizing that data-centric applications often require large software environments and handle massive data volumes, the authors extend an existing methodology to study serverless infrastructures via simulation. They design a scheduling algorithm that simultaneously considers platform heterogeneity, cold start delays, energy consumption, data transfers, makespan, and resource utilization. Using a standard greedy Kubernetes-inspired algorithm as a baseline, the study shows that the proposed multi-objective scheduler can outperform the baseline by up to three orders of magnitude on several metrics, highlighting the importance of tailored scheduling strategies in heterogeneous, serverless environments.

8.3.2 Operational Technology Platforms and Orchestration

Another strand of research addresses the integration of operational technology (OT) with platform-as-a-service (PaaS) models in the continuum. The OTPaaS initiative is presented as a structured framework for managing and storing industrial data with strong requirements on response times, security, reliability, technological and data sovereignty, robustness, and energy efficiency [31]. The associated publication discusses successful deployments, adaptive application management, and integration components for both edge and cloud environments, emphasizing how a PaaS-style abstraction can encapsulate complexity while preserving stringent industrial constraints.

Complementing this, two closely related publications introduce the concept of User-Friendly Orchestration Management (UFOM) for containerized services in the computing continuum [18, 17]. UFOM targets non-expert users by offering an intuitive interface, automated workflows, and contextual assistance to simplify the deployment, monitoring, and maintenance of distributed applications. The approach integrates with osmotic computing principles to support seamless interactions between edge and cloud resources, and evaluates the impact on user-perceived Quality of Experience. A smart home automation case study illustrates how UFOM can democratize orchestration by reducing technical barriers while maintaining system reliability and efficiency in real-world scenarios.

8.3.3 Testbeds and Complexity Analysis for the Continuum

To fully exploit the computing continuum, appropriate research testbeds are necessary. One paper proposes a conceptual testbed for network operating systems in continuum environments, aiming to improve the replicability, scalability, and robustness of experiments [12]. The testbed allows experimenters to modify the operating systems of network devices and dynamically reconfigure network topologies, thereby supporting studies spanning multi-operator settings and internal architectures of telecommunication providers. The authors also investigate mechanisms for virtual topology management, OS deployment, and service orchestration, underlining the need for flexible and programmable infrastructures.

Beyond testbed design, another publication proposes a holistic multidimensional analysis framework to manage complexity in computing continuum systems [11]. The framework combines Quality of Service, Service Level Agreement specifications, and Quality of Experience metrics across multiple levels of the continuum, enabling the characterization of system behavior along several axes simultaneously. By applying this approach to two tiers of a continuum system, the authors show how interlinked metrics can reveal critical properties and bottlenecks that might be missed by single-dimension analyses. This multidimensional perspective supports more informed design and optimization of continuum services, particularly when performance, energy, and user satisfaction must be co-optimized.

9 Bilateral contracts and grants with industry

The amount for CIFRE PhD grants cumulates the support contract DataMove receives and the salary paid directly to the student by the employer.

- **Berger-Levrault (2022-2025)**. CIFRE PhD grant (Halmza Safri). 170K euros
- **ATOS (2022-2026)**. CIFRE PhD grants (Abdessalam Benharii and Guillaume Raffin). 340K euros
- **Orange (2023-2026)**. CIFRE Phd grant (Yoann Dupas). 170K euros.
- **IFPEN (2024-2027)**. Support contract for PhD of Wenke Du. 40K euros
- **SAVOYE (2024-2027)**. CIFRE PhD grant (Gentjan Gjinalaj).

10 Partnerships and cooperations

10.1 European initiatives

10.1.1 Horizon Europe

SEANERGYS

Duration: From June 1, 2025 to May 31, 2029

Partners: 14 EU Partners

Coordinator: FORSCHUNGSZENTRUM JULICH GMBH (FZJ)

Summary: DataMove contributes to the SEANERGYS by developing scheduling policies that maximize resource utilization and energy efficiency, and supports jobs/applications with dynamic and adaptable resource profiles, in particular through the OAR batch scheduler.

LIGHTAIDGE LIGHTAIDGE project on cordis.europa.eu

Title: Light-weight, emissions aware, simulation and orchestration of Edge Computing and Edge Intelligence

Duration: From May 1, 2023 to April 30, 2025

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France

Inria contact: Denis Trystram

Coordinator:

Summary: The annual growth of the global energy consumption of digital technologies is 9%, hindering the EU Green Deal objective of

reducing 55% greenhouse gas (GHG) emission reduction by 2030. With the ever-increasing deployment of Internet of Things (IoT) devices, Edge Computing (EC), and more specifically Edge Intelligence (EI) which seeks to exploit these IoT (Edge) devices to process Artificial Intelligence algorithms has risen as a technology with booming demand potential, but which can also negatively contribute to the global energy consumption and GHG emissions of digital technologies.

Regarding EC and EI, emissions-aware (in CO2 equivalent) simulation and orchestration solutions are still under-explored.

The LIGHTAIDGE project therefore focuses on light-weight, CO2 emissions-aware EI simulation and orchestration. It proposes significant advances by (i) creating a bridge between High-Performance Computing (HPC) and EC communities through the development of a novel, fast and scalable, CO2 emissions aware simulation framework for EC, and (ii) by producing light-weight, CO2 emissions aware Edge Intelligence orchestrators for low-CO2 EI model training.

Foreseen impacts are, at scientific level: the project will establish a bridge between HPC and EC/EI scientific communities, and will pave the path to future, CO2 emissions aware EC and EI research. At technological, economical and societal levels: the project will reduce R&D costs by enabling an economically viable EC and EI prototyping through simulations, will help to drive EI companies in the climate transition by reducing the EI's CO2 emissions through better orchestration, and will contribute to reduce the CO2 emissions due to digital

technologies, participating in the European Union Green Deal's objective. The project also proposes training, transfer of knowledge,

and dissemination/communication activities for the researcher, constituting a solid path to develop his skills and experience.

EoCoE-III EoCoE-III project on cordis.europa.eu

Title: FOSTERING THE EUROPEAN ENERGY TRANSITION WITH EXASCALE

Duration: From January 1, 2024 to December 31, 2026

Partners:

- DATADIRECT NETWORKS FRANCE, France
- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- UNIVERSITA DEGLI STUDI DI ROMA TOR VERGATA (UNITOV), Italy
- FRIEDRICH-ALEXANDER-UNIVERSITAET ERLANGEN-NUERNBERG (FAU), Germany
- FORSCHUNGSZENTRUM JULICH GMBH (FZJ), Germany
- COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES (CEA), France
- CENTRO DE INVESTIGACIONES ENERGETICAS MEDIOAMBIENTALES Y TECNOLOGICAS (CIEMAT), Spain
- INSTYTUT CHEMII BIOORGANICZNEJ POLSKIEJ AKADEMII NAUK, Poland
- UNIVERSITE LIBRE DE BRUXELLES (ULB), Belgium
- AGENZIA NAZIONALE PER LE NUOVE TECNOLOGIE, L'ENERGIA E LO SVILUPPO ECONOMICO SOSTENIBILE (ENEA), Italy
- CENTRE EUROPEEN DE RECHERCHE ET DEFORMATION AVANCEE EN CALCUL SCIENTIFIQUE (CERFACS), France
- E 4 COMPUTER ENGINEERING SPA (E4), Italy
- CONSIGLIO NAZIONALE DELLE RICERCHE (CNR), Italy
- UNIVERSITA DEGLI STUDI DI TRENTO (UNITN), Italy
- IFP Energies nouvelles (IFPEN), France
- MAX-PLANCK-GESELLSCHAFT ZUR FORDERUNG DER WISSENSCHAFTEN EV (MPG), Germany
- CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE CNRS (CNRS), France
- BARCELONA SUPERCOMPUTING CENTER CENTRO NACIONAL DE SUPERCOMPUTACION (BSC CNS), Spain

Inria contact: Bruno Raffin

Coordinator:

Summary: The Energy-oriented Centre of Excellence for exascale HPC applications (EoCoE-III) applies cutting-edge computational methods in its mission to foster the transition to decarbonized energy in Europe. EoCoE-III is anchored both in the High Performance Computing (HPC) community and in the energy field. It will demonstrate the benefit of HPC for the net-zero energy transition for research institutes and also for key industry in the energy sector. The present project will draw the experience of two successful previous projects EoCoE-I and EoCoE-II, where a set of diverse computer applications from four energy domains achieved significant efficiency gains thanks to its multidisciplinary expertise in applied mathematics and supercomputing. During this 3rd round, EoCoE-III will channel its efforts into 5 exascale lighthouse applications covering the key domains of Energy Materials, Water, Wind and Fusion. A world-class consortium of 18 complementary partners from 6 countries will form a unique network of expertise in energy science, scientific computing and HPC, including 3 leading European supercomputing centres. This multidisciplinary effort will harness innovations in computer science and mathematical algorithms within a tightly integrated co-design approach to overcome performance bottlenecks, to deploy the lighthouse applications on the coming European exascale

infrastructure and to anticipate future HPC hardware developments. New modelling capabilities will be created at unprecedented scale, demonstrating the potential benefits to the energy industry, such as accelerated design of photovoltaic devices, high-resolution wind farm modelling over complex terrains and quantitative understanding of plasma core-edge interactions in ITER-scale tokamaks. These lighthouse applications will provide a high-visibility platform for high-performance computational energy science, cross-fertilized through close working connections to the EERA consortium.

10.2 National initiatives

10.2.1 PEPR NUMPEX

Goals: The main objective of the NumPEx (Numeric for Exascale) program in France is to develop state-of-the-art skills and infrastructures in the field of exascale computing.

Duration: From 2023 to 2030

Web site: NUMPEX

DataMove implication:
- Exa-DoST (Data-oriented Software and Tools for the Exascale): Co-lead WP3.
- Exa-AToW (Architectures and Tools for Large-Scale Workflows): Co-lead WP5.
- Exa-DI (Development and integration): CO-lead WP3.

DataMove budget: 1.295 M euros.

10.2.2 ANR

- **PPR Océan et Climat MEDIATION (2022-2030)**. Methodological developments for a robust and efficient digital twin of the ocean. Pi: INRIA team AIRSEA. Partners: INRIA, CNRS, IFREMER, IRD, Université Aix-Marseille, Institut National Polytechnique de Toulouse, Ecole Nationale Supérieure Mines-Télécom Atlantique Bretagne Pays de la Loire, Service Hygrodgraphique et Océanographique de la Marine, Université Grenoble Alpes, Météo-France-DESR-Centre National de Recherches Météorologiques. Total budget: 2,4 Meuros. DataMove Budget: 110 Keuros. CO-lead of the WP Leveraging AI and HPC for Digital Twins of the Ocean.
- **AAPG2023 PREDICTIONS (2024-2027)**. This project aims to substantially strengthen and expand the foundations of the nascent, but fast-growing area of algorithms with predictions, in a global framework that addresses all aspects of algorithm development: modeling, design, framework of analysis, and performance evaluation. Specifically, we put forward three main objectives. Pi: LIP6. Partners: LIG/DataMove, IRIF,CC-IN2P3, LIRIS. Total budget: 358k euros.Datamova Budget: 128k euros.
- **AAPG2025 SOCLOUD (20252029)**. The aim of the project is to study the human and technical conditions for implementing sobriety in the cloud, and to identify the levers and their consequences. Partenaires: Univ. Besançon, Univ Toulouse, INRIA, Eaton Industries.

10.3 Public policy support

DataMove engaged in initiatives aimed at civil society, contributing to the specification AFNOR "Frugal AI Framework", and the revision of the Ecoindex for measuring the carbon footprint of web requests).

11 Dissemination

11.1 Scientific events: organisation

- DataMove is the initiator and oraganizer of JRAF (Frugal Training Research Days), Grenoble 2022-2025.

11.2 Scientific expertise

- Bruno Raffin is Reviewer for The Research Council of Norway (RCN).

11.3 Research administration

- Yves Denneulin is the Scientific Director of the Labex Persyval. Mastering the convergence of the physical and digital worlds.
- Thang Nguyen is : Member of the Scientific Board of MIAI Grenoble (Multidisciplinary Institute in Artificial Intelligence) and Education Director EFELIA-MIAI.
- Olivier Richard is member of the steering committee of GDR-RSD (Réseaux et Systèmes Distribués) since 2024
- Denis Trystram is member of the board of directors of GdR RO (Recherche Operationelle). Initiator and responsable of thr transversal action on numerical frugality. Since 2020
- Thang Nguyen is member of the Scientific board, GT Complexity and Algorithms, GDR IFM

11.4 Teaching

Datamove has a strong teaching activity thanks to its many permanent members that are Associate Professors or Professors at UGA and UGA/INPG Grenoble. We only list bellow the teaching activity of Datamove permanent members. Additionaly, most PhD students teach a few tens of hours every year at UGA.

- Denis Trystram. 200 hours per year, ENSIMAG, Grenoble-INP, Master
- Fanny Dufossé. 17 to 90 hours per year, Algorithmic, Licence. Univ. Grenoble-Alpes and Licence Ensimag, Combinatorial scientific computing, Master, ENS Lyon.
- Pierre-François Dutot. 226 hours per year. Licence (first and second year) at IUT2/UPMF (Institut Universitaire Technologique de Univ. Grenoble-Alpes) and 9 hours Master M2R-ISC Informatique-Systèmes-Communication at Univ. Grenoble-Alpes.
- Grégory Mounié is responsible for the first year (M1) of the international Master of Science in Informatics at Grenoble (MOSIG-M1). 317 hours per year. Master (M1/2nd year and M2/3rd year) at Engineering school ENSIMAG, Grenoble-INP, Univ Grenoble Alpes.
- Bruno Raffin. 28 hours per year. Parallel System. International Master of Science in Informatics at Grenoble (MOSIG-M2). Co-organizer of the 20205 summer school *Solving partial differential equations in fields physics faster with physics-based machine learning*.
- Olivier Richard is responsible for the third year of the computer science department of Grenoble INP. 222 hours per year. Master at Engineering school Polytech-Grenoble, Univ. Grenoble-Alpes. Co-organiser of the tutorial *Reproducible distributed environments with NixOS Compose* at ACM REP'24.
- Frédéric Wagner. 220 hours per year. Engineering school ENSIMAG, Grenoble-INP, Master (M1/2nd year and M2/3rd year).
- Yves Denneulin. 192 hours per year. Engineering school ENSIMAG, Grenoble-INP, Master (M1/2nd year and M2/3rd year).
- Nguyen Kim Thang. 250 hours per year. Engineering school (Ensimag), Grenoble INP, UGA, and master MoSiG (1st and 2nd years), UGA.
- Danilo Carastan dos Santos. 144 hours per year (service reduced due to new recruitment). Licence (third year) and Master (first and second year) at IM2AG-UGA (Informatique, mathématiques et mathématiques appliquées of Univ. Grenoble-Alpes) and 12 hours in first year at the ENSIMAG engineering school.

11.5 Popularization

Datamove made compute frugality one of our research items, with activities ranging from research on energy measurement to evaluation of the carbon impact of data centers, organizing the workshop series JRAF on frugal AI, and raising awareness of the environmental impact of digital technologies, particularly AI, among broader audiences. Denis Trystram, in particular, has developed a collaboration with the philosopher Thierry Ménissier, and has given talks and participated in debates for various non-CS audiences. The 2025 talk:

- *Introduction aux coûts environnementaux de l'IA*. Journée CNRS Calcul Sobre, Grenoble – June 2025

12 Scientific production

12.1 Major publications

[1] D. Carastan-Santos and R. Y. de Camargo. 'Obtaining Dynamic Scheduling Policies with Simulation and Machine Learning'. In: SC'17 -2 International Conference for High Performance Computing, Networking, Storage and Analysis (Supercomputing). Denver, United States, 12th Nov. 2017. URL: https://inria.hal.science/hal-01618940.

[2] P.-F. Dutot, M. Mercier, M. Poquet and O. Richard. 'Batsim: a Realistic Language-Independent Resources and Jobs Management Systems Simulator'. In: *20th Workshop on Job Scheduling Strategies for Parallel Processing (JSSPP)*. 20th Workshop on Job Scheduling Strategies for Parallel Processing. Chicago, United States, 27th May 2016. URL: https://hal.archives-ouvertes.fr/hal-01333471.

[3] S. Friedemann and B. Raffin. 'An elastic framework for ensemble-based large-scale data assimilation'. In: *The international journal of high performance computing applications* 36 (28th June 2022), pp. 1–37. DOI: 10.1177/10943420221110507. URL: https://hal.inria.fr/hal-03017033.

[4] Q. Guilloteau, J. Bleuzen, M. Poquet and O. Richard. 'Painless Transposition of Reproducible Distributed Environments with NixOS Compose'. In: CLUSTER 2022 - IEEE International Conference on Cluster Computing. Vol. CLUSTER 2022 - IEEE International Conference on Cluster Computing. Heidelberg, Germany, 6th Sept. 2022, pp. 1–12. URL: https://hal.science/hal-03723771.

[5] G. Lucarelli, B. Moseley, N. K. Thang, A. Srivastav and D. Trystram. 'Online Non-preemptive Scheduling on Unrelated Machines with Rejections'. In: SPAA 2018 - 30th ACM Symposium on Parallelism in Algorithms and Architectures. Vienna, Austria: ACM Press, 2018, pp. 291–300. DOI: 10.1145/3210377.3210402. URL: https://hal.archives-ouvertes.fr/hal-01986312.

[6] L. Meyer, M. Schouler, R. A. Caulk, A. Ribés and B. Raffin. 'High Throughput Training of Deep Surrogates from Large Ensemble Runs'. In: *SC '23: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. SC 2023 - The International Conference for High Performance Computing, Networking, Storage, and Analysis. Denver, CO, United States: ACM, 18th Nov. 2023, pp. 1–14. DOI: 10.1145/3581784.3607083. URL: https://hal.science/hal-04213978.

[7] M. F. Silva Vasconcelos, D. Cordeiro, G. da Costa, F. Dufossé, J.-M. Nicod and V. Rehn-Sonigo. 'Optimal sizing of a globally distributed low carbon cloud federation'. In: CCGrid 2023 - IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing. Bangalore, India, 2023, pp. 1–13. DOI: 10.1109/CCGrid57682.2023.00028. URL: https://hal.science/hal-04032094.

[8] F. Zanon Boito, E. Camilo Inacio, J. Luca Bez, P. O. A. Navaux, M. A. R. Dantas and Y. Denneulin. 'A Checkpoint of Research on Parallel I/O for High Performance Computing'. In: *ACM Computing Surveys* 51.2 (Mar. 2018), 23:1–23:35. DOI: 10.1145/3152891. URL: https://hal.univ-grenoble-alpes.fr/hal-01591755.

12.2 Publications of the year

International journals

[9] A. Berthelot, E. Caron, M. Jay and L. Lefèvre. 'Understanding the environmental impact of generative AI services'. In: *Communications of the ACM* Special Issue on Sustainability and Computing 68.7 (2025), pp. 46–53. DOI: 10.1145/3725984. URL: https://hal.science/hal-04920612 (cit. on p. 20).

[10] D. Carastan-Santos, G. da Costa, I. Fontana de Nardin, M. Poquet, K. Rzadca, P. Stolf and D. Trystram. 'Scheduling with lightweight predictions in power-constrained HPC platforms'. In: *IEEE Transactions on Parallel and Distributed Systems* (8th July 2025), pp. 1–12. DOI: 10.1109/TPDS.2025.3586723. URL: https://hal.science/hal-04747713 (cit. on p. 20).

International peer-reviewed conferences

[11] C. J. Barrios, Y. Denneulin and F. Le Mouël. 'A Holistic Approach to Complexity Management and Multidimensional Analysis in Computing Continuum'. In: WSCC 2025 - 3rd International Workshop on Scalable Compute Continuum. Dresde (Germany), Germany, 2025, pp. 1–12. URL: https://hal.science/hal-05393403 (cit. on p. 21).

[12] J. Caposiena, O. Carrillo, F. Le Mouël, B. Jonglez, P. Neyron and T. Arrabal. 'Towards a flexible Network Operating System Testbed for the Computing Continuum'. In: CCGridW 2025 - 25th IEEE International Symposium on Cluster, Cloud and Internet Computing Workshops. 2025 IEEE 25th International Symposium on Cluster, Cloud and Internet Computing Workshops (CCGridW). Tromsø Norway, Norway, 2025, pp. 148–155. DOI: 10.1109/CCGridW65158.2025.00029. URL: https://hal.science/hal-05154217 (cit. on p. 21).

[13] A. A. Da Silva, Y. Georgiou, M. Mercier, G. Mounié and D. Trystram. 'FOA-Energy: A Multi-objective Energy-Aware Scheduling Policy for Serverless-based Edge-Cloud Continuum'. In: SAC 2025 - 40th ACM/SIGAPP Symposium on Applied Computing. Catania International Airport Catania Italy, Italy: ACM, 31st Mar. 2025, pp. 225–232. DOI: 10.1145/3672608.3707941. URL: https://hal.science/hal-05156857 (cit. on p. 20).

[14] Y. Dupas, O. Hotel, G. Lefebvre and C. Cerin. 'MEFA: Multimodal Image Early Fusion with Attention Module for Pedestrian and Vehicle Detection'. In: VISAPP 2025 - 20th International Conference on Computer Vision Theory and Applications. Porto, Portugal: SCITEPRESS - Science and Technology Publications, 2025, pp. 610–617. DOI: 10.5220/0013236000003912. URL: https://hal.science/hal-05010469 (cit. on p. 19).

[15] D. Huber, S. Iserte, M. Schreiber, A. J. Peña and M. Schulz. 'Bridging the Gap Between Genericity and Programmability of Dynamic Resources in HPC'. In: ISC High Performance 2025 - 40th ISC High Performance International Conference. Hamburg, Germany, 2025, pp. 1–11. URL: https://inria.hal.science/hal-04994828 (cit. on p. 20).

[16] G. Raffin, D. Trystram and O. Richard. 'Alumet: a Modular Framework to Standardize the Measurement of Energy Consumption'. In: PECS 2025 - Workshop on Performance and Energy Efficiency in Concurrent and Distributed Systems. Dresden, Germany: Springer, 2025, pp. 1–12. URL: https://hal.science/hal-05246933 (cit. on p. 20).

[17] P. J. Rojas Yepes, C. J. B. Hernández, O. Carrillo and F. Le Mouël. 'User-Friendly Orchestration Management Proposal'. In: CCGRID 2025 - 25th IEEE International Symposium on Cluster, Cloud, and Internet Computing. Tromso, Norway, Norway: IEEE, 2025, pp. 1–4. URL: https://hal.science/hal-05064624 (cit. on p. 21).

[18] P. J. Rojas Yepes, C. J. B. Hernández, O. Carrillo and F. Le Mouël. 'User-Friendly Orchestration Management Proposal'. In: 2025 - 5th Workshop CATAÏ. Chalon sur Saône, France, 2025, pp. 1–4. URL: https://hal.science/hal-05064641 (cit. on p. 21).

National peer-reviewed Conferences

[19] M. Gradvohl, E. Chargy, E. Dreina, D. Carastan-Santos and F. Rousseau. 'Étude de l'empreinte carbone d'un réseau de capteurs dans un tableau de distribution électrique'. In: *ALGOTEL 2025 – 27èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications*. ALGOTEL 2025 – 27èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications. Saint Valery-sur-Somme, France, 2025, pp. 1–5. URL: https://hal.science/hal-05033384 (cit. on p. 19).

Conferences without proceedings

[20] A. Benhari, Y. Denneulin, F. Desprez, F. Dufossé and D. Trystram. 'Analysis of the carbon footprint of HPC'. In: PECS 2025 - International Workshop on Performance and Energy Efficiency in Concurrent and Distributed Systems. Dresden, Germany, 2025, pp. 1–13. URL: https://hal.science/hal-05248774 (cit. on p. 19).

[21] R. Boëzennec, F. Dufossé, G. Pallez and A. Tremodeux. 'Improving Supercomputer Usage with Aging Awareness'. In: Sustainable Supercomputing (Workshop of SC25). St. Louis, Missouri, United States, 16th Nov. 2025. URL: https://hal.science/hal-05109521 (cit. on p. 19).

[22] J. Caposiena, O. Carrillo, B. Jonglez, P. Neyron and T. Arrabal. 'Vers un banc d'essai flexible pour les systèmes d'exploitation réseau dans le Computing Continuum'. In: ComPAS. Ed. by F. Le Mouël. Bordeaux, France, 24th June 2025. URL: https://hal.science/hal-05128936.

[23] Y. Dupas, O. Hotel, G. Lefebvre and C. Cérin. 'MEFA-MS: Attention-Based U-Net for Pedestrian and Vehicle Detection'. In: ICMV 2025 - Eighteenth International Conference on Machine Vision. Paris, France, 2025, pp. 1–8. URL: https://hal.science/hal-05368360 (cit. on p. 19).

[24] E. Foussard, M. Nattaf, M.-L. Espinouse and G. Mounié. 'Minimisation du délai moyen en présence d'une contrainte de santé de l'équipement : propriétés structurelles'. In: ROADEF 2025 - 26ème édition du congrès annuel de la Société Française de Recherche Opérationnelle et d'Aide à la Décision. Champs-sur-Marne, France, 2025, pp. 1–2. URL: https://hal.science/hal-05023454.

Reports & preprints

[25] M. Bacou, D. Beserra, E. Dedu, L. Desgeorges, D. Donsez, A. Guitton, B. Jonglez, A. Legrand, G. Papadopoulos, O. Richard, S. Si-Mohammed, N. Tamdrari and F. Theoleyre. *Journée thématique du GDR RSD : pratiques expérimentales de la communauté systèmes et réseaux*. 31st Jan. 2025. URL: https://hal.science/hal-04924273.

[26] S. Bouveret, A. Bugeau, F. Emmanuelle, J. Lefevre, L. Lefèvre, A.-L. Ligozat, P. Marquet, A.-C. Orgerie and D. Trystram. *Quiz sur les impacts environnementaux du numérique*. EcoInfo, Feb. 2025, pp. 1–5. URL: https://hal.science/hal-04960328.

[27] E. Foussard, J. Martinovic, M.-L. Espinouse, G. Mounié and M. Nattaf. *Bin Packing with Thresholds: Mathematical Models and Theoretical Results*. 2025. URL: https://hal.univ-grenoble-alpes.fr/hal-05230011.

[28] A. Letizia, C. Cérin and D. Donsez. *WildCount: Embedded Deep Learning for Wildlife Recognition*. LIG : Laboratoire d'informatique de Grenoble; Grenoble INP - UGA, 17th Sept. 2025, pp. 1–17. URL: https://hal.science/hal-05267715.

[29] G. Saraiva, M. Vasconcelos, S. Mazzini Bruschi, D. Carastan-Santos and D. Cordeiro. *Estimating CO2 emissions of distributed applications and platforms with SimGrid/Batsim*. 2025. URL: https://hal.science/hal-05212970.

[30] N. K. Thắng. *Price of Anarchy in Resource Allocation and Auto-Bidding Advertising via Duality in Linear/Convex Programming*. 2025. URL: https://hal.science/hal-05187001.

Scientific popularization

[31] C. J. Barrios and Y. Denneulin. 'Bridding OT and PaaS in Edge-to-Cloud Continuum: The OTPaaS Concept'. In: COMPAS 2025 - Conférence Francophone d'Informatique en Parallélisme, Architecture et Système. BORDEAUX, France, 2025, pp. 1–13. URL: https://hal.science/hal-05128715 (cit. on p. 21).

Software

[32] [SW] G. Raffin, *Alumet* 18th Apr. 2025. LIC: European Union Public License 1.2. HAL: ⟨hal-05036272⟩, URL: https://hal.science/hal-05036272, VCS: https://github.com/alumet-dev/alumet, SWHID: ⟨swh:1:dir:cfe6a77d1cf766a668dc8dd6edd0b2b481281220;origin=https://github.com/alumet-dev/alumet;visit=swh:1:snp:4c238b48d88b6041bf7d5d146607b7f0e4510887;anchor=swh:1:rev:97abae555fa19b7c9029da249786e234fd080275⟩.