

Inria International Program - Associate Team Final report 2022

Associate Team acronym: UNIFY (associated to the JLESC International Lab)

Period of activity: 01/01/2019 – 31/12/2022

Principal investigator (Inria): Gabriel Antoniu, KerData project-team, Rennes

Principal investigator (partner): Tom Peterka, Argonne National Laboratory (USA)

Other participants: DataMove Project-Team, Inria, Grenoble

1. List of participants

All participants are involved in the JLESC international lab and regularly take part in JLESC workshops.

Argonne National Lab

- **Tom Peterka** (ANL PI), Senior Researcher, <http://www.mcs.anl.gov/person/thomas-peterka/>.
- **Rob Ross** Senior Researcher, <http://www.mcs.anl.gov/person/rob-ross/>.
- **Swann Perarnau**, Junior Permanent Researcher, <http://www.mcs.anl.gov/person/swann-perarnau/>.
- **Matthieu Dorier**, Software Development Scientist (permanent): <https://dorier.github.io>.
- **Franck Cappello**, Senior Researcher, <https://www.anl.gov/profile/franck-cappello>.
- **Bogdan Nicolae**, Computer Scientist, <https://www.anl.gov/profile/bogdan-nicolae>.
- **Main junior researchers** (with contractual funding) initially involved as postdocs (NB: later, some of them obtained permanent positions at ANL or elsewhere): **Pierre Matri**, Postdoc Fellow, **Orçun Yildiz**, Postdoc Fellow, **Valentin Reis**, Postdoc Fellow.

Inria – KerData Team

- **Gabriel Antoniu** (Inria PI), Senior Researcher, <https://team.inria.fr/kerdata/gabriel-antoniu/>.
- **Alexandru Costan**, Associate Professor, INSA Rennes, <https://team.inria.fr/kerdata/alexandru-costan/>.
- **Main junior researchers involved:** **Nathanaël Cherièr**, PhD student, **Daniel Rosendo**, PhD student, **Thomas Bouvier**, PhD student, **Hugo Chaugier**, Master student.

Inria – DataMove Team

- **Bruno Raffin**, Senior Researcher, <https://team.inria.fr/datamove/team-members/bruno-raffin/>.
- **Olivier Richard**, Associate Professor, Grenoble Alpes, <https://team.inria.fr/datamove/team-members/olivier-richard>.
- **Main junior researchers involved:** **Salah Zrigui**, PhD Student, **Amal Gueroudji**, PhD Student, **Luc Angelelli**, PhD Student, **Mathilde Jay**, PhD Student.

Key words (according to Inria scientific cartography) – scientific themes and application domains

Computer science and Digital Science - A1.1.4. High performance computing; A1.1.5. Exascale; A1.3. Distributed Systems; A1.3.5. Cloud; A1.3.6. Fog, Edge; A2.6.2. Middleware; A3.1.2. Data management, querying and storage; A3.1.3. Distributed data; A3.1.8. Big data (production, storage, transfer); A6.2.7. High performance computing; A6.3. Computation-data interaction; A7.1.1. Distributed algorithms; A9.2. Machine learning. Other Research Topics and Application Domains: B9.5.6. Data science; B9.8. Reproducibility.

Associate Team website: <https://team.inria.fr/unify/>

1. Overview of the activities

1.1. Short visits and physical meetings (< 1 month)

Participants – Location – Dates - Purpose

- **Nathanaël Cherièr** (Inria), **Tom Peterka** (ANL), **Gabriel Antoniu** (Inria), **Bruno Raffin** (Inria) - Knoxville, USA, 9th JLESC workshop - 15-17 Apr 2019 - Presentations on joint work and discussions on the collaboration.
- **Gabriel Antoniu** (Inria), **Bruno Raffin** (Inria), **Matthieu Dorier** (ANL), Denver, SC19 conference - 17-22 Nov 2019 - Meetings and discussions on ongoing and future collaborative work.

- Kate Keahey (ANL) and Matthieu Dorier (ANL) - KerData team, Inria, Rennes - 4-6 Nov 2019 - Visit, meetings and discussions on collaborative topics around the PhD defense of Nathanaël Cheriére, co-advised by Gabriel Antoniu (Inria) and Matthieu Dorier (ANL).
- Franck Cappello (ANL) - DataMove team, Inria, Grenoble - 5-6 Dec 2019 - Talk on data compression, discussion on collaboration and participation to the HDR jury of Thierry Gautier (Inria).
- François Tessier (Inria), Alexandru Costan (Inria) and Gabriel Antoniu (Inria) – Argonne - 27 Sep 2022 - Seminar at ANL and discussions with Tom Peterka (ANL), Franck Capello (ANL), Swann Perarnau (ANL) on collaboration topics.
- Daniel Rosendo (Inria), Thomas Bouvier (Inria), François Tessier (Inria), Alexandru Costan (Inria), Gabriel Antoniu (Inria) met Phil Carns (ANL), Orçun Yildiz (ANL), Bogdan Nicolae (ANL), Tom Peterka (ANL), Franck Cappello (ANL) - UIUC, Urbana-Champaign - 28-30 Sep 2022 - Talks at the 14th JLESC workshop (see Seminars section below) and discussions on collaboration topics.
- Julien Monnot (Inria), Gabriel Antoniu (Inria), François Tessier (Inria), Bogdan Nicolae (ANL), Orçun Yildiz (ANL), Tom Peterka (ANL) - Dallas, SC22 conference - 13-18 Nov 2022 - Meetings and discussions on ongoing and future collaborative work.

1.2. Virtual meetings related to virtual JLESC workshops

2020-2021: Due to the pandemic, no physical visits could take place in 2020 and 2021. Several meetings took place at the 11th (8-10 September 2020), 12th (24-26 February, 2021) and 13th (14-16 December, 2021) virtual JLESC workshops and as a follow-up.

- Follow-up meeting immediately after the 11th JLESC virtual workshop. Participants: Daniel Rosendo, Gabriel Antoniu (Inria) and Orçun Yildiz (ANL). As a result, Daniel started to explore the usage of DeepHyper (<https://deephyper.readthedocs.io/en/latest/>) developed at ANL in conjunction with the E2Clab deployment framework for reproducible experimentation across the digital continuum (<https://team.inria.fr/kerdata/e2clab/>) developed by KerData, Inria. Goal: enable ML-based experimentation across the digital continuum.
- 11th and 12th JLESC virtual workshops: online discussions following the talks given by the various participants (see the list of talks in Section 1.5).

1.3. Weekly remote meetings (Nov 2020 - today)

- **Nov 2020 - June 2021:** Hugo Chaugier (MS student, Inria), Alexandru Costan (Inria), Gabriel Antoniu (Inria) and Bogdan Nicolae (ANL) started weekly meetings for a collaboration on the topic “Study of Catastrophic Forgetting in Deep Learning”.
- **Sep 2021 - today:** Thomas Bouvier (PhD student, Inria), Alexandru Costan (Inria), Gabriel Antoniu (Inria) and Bogdan Nicolae (ANL) keep organizing (almost) weekly meetings to further develop the topic initiated by Hugo Chaugier. Focus: Data parallel continual learning to mitigate catastrophic forgetting in large-scale deep learning workflows.

1.4. Long visits (duration > 1 month)

INRIA International Chair: Ian Foster, ANL & Professor Univ. Chicago, received an INRIA international Chair at DataMove for spending 2.5 months during the 5 years starting in 2022. Unfortunately due to pandemic travel restrictions and personal constraints, the 2022 and 2023 visits have been canceled.

1.5. Seminars (at JLESC workshops)¹

- [9th JLESC workshop](#), Knoxville, April 15-17, 2019
 - Pierre Matri (PhD student, Inria): project talk entitled “Leveraging Blob-Based Storage for Stream Processing”
 - Nathanaël Cheriére (PhD student, Inria): project talk entitled “Rescaling Transient Distributed Storage Systems”
 - Gabriel Antoniu (Inria) gave a short talk entitled “Scalable Data Ingestion for Stream Processing”
- [11th JLESC workshop](#), virtual, September 8-10, 2020
 - Daniel Rosendo (PhD student, Inria): short talk entitled “E2Clab: Exploring the Computing Continuum through Repeatable, Replicable and Reproducible Edge-to-Cloud Experiments”
 - François Tessier (Inria): Dynamic Provisioning of Storage Resources on Supercomputers
- [12th JLESC workshop](#), virtual, February 24 - 26 2021
 - Gabriel Antoniu (Inria) was invited to give a keynote talk on the 10-year collaboration with Argonne

¹ Due to the large number of talks given workshops by the participants to the UNIFY Associate team at JLESC, we only list here talks by French team members. The full list of talks including talks from ANL collaborators is available on the JLESC site for each edition.

- National Laboratory within the JLESC international Laboratory
- Daniel Rosendo (PhD student, Inria): short talk entitled “E2Clab: Optimizing Complex Workflow Deployments on the Edge-to-Cloud Continuum”
- François Tessier (Inria): Storage allocation over hybrid HPC/Cloud Infrastructures
- [13th JLESC workshop](#), virtual, December 14 – 16, 2021
 - Bruno Raffin, Lucas Nussbaum and Gabriel Antoniu, Inria - Joint talk: INRIA@Exascale: Unifying Software Development Leveraging the E4S Initiative.
 - Panel on AI in HPC. Moderator: Franck Cappello, ANL. Panelists: Prasanna, Balaprakash, ANL, Piotr Luszczek, UTK, Takahide Matsuoka, Riken, Dario Garcia-Gasulla, BSC, Martin Schultz, JSC, Volodymyr Kindratenko, Bruno Raffin (Inria).
 - Daniel Rosendo (PhD student, Inria): short talk entitled “Reproducible Performance Optimization of Complex Applications on the Edge-to-Cloud Continuum”
 - Thomas Bouvier (PhD student, Inria): short talk entitled “Deploying Heterogeneity-aware Deep Learning Workloads on the Computing Continuum”
 - Joshua Bowden (Inria): short talk entitled “Asynchronous I/O using Damaris - Results from integration with Code_Saturne CFD code”
- [14th JLESC workshop](#) in Urbana, IL, USA, September 28 – 30, 2022
 - Daniel Rosendo (PhD student, Inria): Project talk entitled “Towards a Collaborative Environment for the Cost-effective Reproducibility of Edge-to-Cloud Experiments on Multi-platforms” on the results of his 3-month internship at ANL
 - Thomas Bouvier (PhD student, Inria): Project talk entitled “Data Parallel Rehearsal-based Continual Learning” on the results of his 3-month internship at ANL
 - François Tessier (Inria): invited panelist to discuss HPC/Cloud convergence
 - François Tessier (Inria): Short talk entitled “Investigating allocation of heterogeneous storage resources on HPC systems”
 - Gabriel Antoniu (Inria): Short talk entitled “Towards Integrated Hardware/Software Ecosystems for the Edge-Cloud-HPC Continuum: the Transcontinuum Initiative”

1.6. Joint workshops

- **Workshop on the In Situ Co-Execution of High-Performance Computing & Data Analysis** of ISC - July 2, 2021 - <https://hpcda.github.io/> - Bruno Raffin (Inria) was co-organizer and Matthieu Dorier (ANL) was invited for a talk.
- Bruno Raffin (Inria) organized a **joint break-out session on the topic “In Situ Processing at Large”** at the 13th JLESC Workshop, 2021. All participants are contributors to the UNIFY Associate Team.
 - Joshua Bowden (Inria) – “Asynchronous I/O using Damaris: results from integration with Code_Saturne CFD code”
 - Matthieu Dorier (ANL) – “Enabling elastic in situ visualization in ParaView using Mochi”
 - Orcun Yildiz and Tom Peterka (ANL) – “Dynamic Heterogeneous In Situ Workflows with LowFive and Wilkins”
 - Bruno Raffin (INRIA) – “Structuring Data Transfers”

1.7. Submission of joint projects

The participants in the UNIFY Associate Team submitted a follow-up proposal for a new Inria-funded Associate Team called UNIFY 2 (2023-2025). Requested budget: 10,000€/year for 3 years, starting in January 2023. Co-funding from ANL available. Submitted in October 2022, under evaluation.

1.8. Co-organization of scientific events

Virtual Workshop organization: François Tessier co-organized the 12th JLESC virtual workshop in February 2021.

1.9. Students co-supervision

- **Co-advised PhD thesis of Nathanael Cherié (ENS Rennes)** by Gabriel Antoniu (Inria) and Matthieu Dorier (ANL), with strong collaboration with Rob Ross (ANL) – *Towards Malleable Distributed Storage Systems: from Models to Practice*. Defense date: 5 November 2019.
- Two 3-months undergraduate internships co-advised by Nathanaël Cherié and Gabriel Antoniu (Inria) and Matthieu Dorier (ANL): **Juliette Fournis d’Albiat (ENS Cachan)** : *Models for rescaling elastic storage systems*. **Tom Bordin (ENS Rennes)** : *Models for data transfers in elastic storage systems*.
- **Hugo Chaugier, MS student**, Inria, Rennes, 5 months (Feb - June 2022), co-advised by Alexandru

Costan (Inria), Gabriel Antoniu (Inria) and Bogdan Nicolae (ANL) through weekly online meetings. Topic: Study of Catastrophic Forgetting in Deep Learning. Defended in June 2021.

- **Thomas Bouvier, PhD student (INSA Rennes/Inria, Rennes)**, January 2021 - Today. Co-advised by Alexandru Costan (Inria), Gabriel Antoniu (Inria), with technical advisorship from Bogdan Nicolae (ANL) through weekly online meetings since December 2021. Topic: Data parallel continual learning to mitigate catastrophic forgetting in large-scale deep learning workflows. Planned defense date: May 2024.

1.10. Research internships for Master and PhD students

Research internships at the partner were organized in 2019, were stopped in 2020-2021 due to pandemics (and replaced by remote co-supervision as explained above). Regular internships were resumed in 2022.

2019

- Aleksandr Danilin, Master Student, Inria (DataMove team) joined Swann Perarnau at ANL for a 3-months summer internship to work on the *Development of a GPU memory management module of a portable, open-source, integrated, validated, and scalability tested HPC framework*.

2022

- Daniel Rosendo, PhD student, Inria (KerData), joined Kate Keahey at ANL for a 3-month summer internship to work on the topic: *Advancing Chameleon and Grid'5000 testbeds to support experiments for the Computing Continuum*. Detailed description: https://jlesc.github.io/projects/e2clab_project/.
- Thomas Bouvier, PhD student, Inria (KerData), joined Bogdan Nicolae at ANL for a 3-month summer internship to work on the topic: *Towards Continual Learning at Scale*. Detailed description: https://jlesc.github.io/projects/continual_learning_project/.

2. Scientific achievements

This Associate Team has been exploring innovative approaches to workflow optimization, adaptive data management and processing through hybrid techniques leveraging the strengths of three areas: HPC, Big Data analytics and Machine learning. Here are the main achievements. More details are available on the JLESC sites (see links at the end of each description below).

Axis 1 - Efficient Data Management for hybrid workflows using transient storage systems

As parallel file systems reach their limits in HPC, an alternative approach is emerging: the creation of data services tailored to the needs of the application that uses it ([Dorier et al. 2018](#)). In order to do so, a set of building blocks for data services (e.g. membership service, key-value store, ...) has been developed as part of the Mochi project (<https://www.mcs.anl.gov/research/projects/mochi/>). In particular, some applications (typically workflows) have requirements in number of machines that vary over time, thus being able to rescale efficiently can be needed for data services co-deployed with malleable applications. Preliminary work by Cherié et al. showed that the reconfiguration could be done in a short time when the amount of data per node is balanced. In this joint work we investigated the use of rescaling of distributed storage systems in HPC environments. In particular, we developed PufferScale, a rescaling scheduler interoperable within Mochi microservices. Pufferscale was demonstrated with a set of microservices representative of a storage system designed for High Energy Physics applications, and evaluated on the Grid'5000 plat. This work led to 4 joint papers [1,2,3,4].

More details: https://jlesc.github.io/projects/reconf_distr_storage/

Axis 2 – Unified data processing architecture for hybrid workflows including simulations and analytics

Supporting reproducible experiments across the Computing Continuum. Distributed digital infrastructures for computation and analytics are now evolving towards an interconnected ecosystem allowing complex applications to be executed from IoT Edge devices to the HPC Cloud (aka the Computing Continuum) to support complex workflows combining simulation, analytics and AI. Understanding end-to-end performance for data processing in such a complex continuum is challenging. One important challenge is to accurately reproduce relevant behaviors of a given application workflow and representative settings of the physical infrastructure underlying this complex continuum. At Inria we introduced a rigorous methodology for such a process and validated it through E2Clab ([Rosendo et al. 2020](#)). It is the first platform to support the complete analysis cycle of an application on the Computing Continuum (from configuration, to deployment, automated application optimization, metrics collection). In this work we extended the E2Clab approach to enable scientists to effectively reproduce and explore experiments run in a large variety of experimental testbeds with different characteristics. We demonstrated how our approach for reproducibility helps scientists to reproduce complex Edge-to-Cloud workflows across Chameleon/CHI@Edge/G5K. Our ultimate goal is to lower the barrier to reproducing research by combining the reproducible artifacts and the experimental environment. A joint paper is in

preparation to be submitted soon [5].

More details: https://jlesc.github.io/projects/e2clab_project/.

Towards Continual Learning at Scale. Deep Learning (DL) emerged as a more and more widely used data processing technique to extract valuable information from ever-growing volumes of data. However, when trained on sequential tasks, i.e. without full access to the input data at the beginning of the training, typical Deep Neural Networks (DNNs) suffer from catastrophic forgetting, a phenomenon causing them to reinforce new patterns at the expense of previously acquired knowledge. Memory rehearsal methods are effective in mitigating accuracy degradation in such continual settings. In this paper, we discuss how to adopt rehearsal in the context of data-parallel training. We proposed an approach where rehearsal buffers are distributed across nodes to achieve effective scalability on HPC systems. We investigated the trade-offs introduced by parallel and distributed continual learning in terms of training time, accuracy and memory usage. A joint paper is in preparation to be submitted soon [6]. More details: https://jlesc.github.io/projects/continual_learning_project/.

Axis 3 - ML-based resource management for hybrid exascale workflows and adaptive data services

The pandemic severely disrupted the activity on this topic (no visit, no traveling). However, Amal Geroudji, PhD student 2020-2023, co-advised by Bruno Raffin (Inria) and Julien Bigot (CEA), investigated this topic by developing a solution to enable in situ data processing with the framework Dask. Amal Gueroudji will proceed after her defense in mai 2023 with a postdoc at ANL in the team of Rob Ross.

3. Production & Impact

3.1. Joint publications

1. Nathanaël Cheriére, Matthieu Dorier, Gabriel Antoniu. *Is it Worth Relaxing Fault Tolerance to Speed Up Decommission in Distributed Storage Systems?* In CCGrid 2019 – IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing, May 2019, Larnaca, Cyprus. pp.1-10.
2. Nathanaël Cheriére, Matthieu Dorier, Gabriel Antoniu. *How fast can one resize a distributed file system?* In Journal of Parallel and Distributed Computing, Elsevier, 2020, 140, pp.80-98. (10.1016/j.jpdc.2020.02.001).
3. Nathanaël Cheriére, Matthieu Dorier, Gabriel Antoniu, Stefan M. Wild, Sven Leyffer, Robert Ross. *Pufferscale: Rescaling HPC Data Services for High Energy Physics Applications 2020*. In Proc. 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID), May 2020, Melbourne, Australia. pp.182-191, (10.1109/CCGrid49817.2020.00-75)
4. Pierre Matri, Yevhen Alforov, Alvaro Brandon, María Pérez, Alexandru Costan, Gabriel Antoniu, Michael Kuhn, Philip Carns, Thomas Ludwig. *Mission Possible: Unify HPC and Big Data Stacks Towards Application-Defined Blobs at the Storage Layer*. Future Generation Computer Systems, 2020, 109, pp.668-677. (10.1016/j.future.2018.07.035)
5. Daniel Rosendo, Kate Keahey, Alexandru Costan, Gabriel Antoniu. *Supporting reproducible experimentation for complex Edge-to-Cloud workflows across multiple testbeds*. In preparation, 2023, to be submitted.
6. Thomas Bouvier, Hugo Chaugier, Bogdan Nicolae, Alexandru Costan, Gabriel Antoniu. *Leveraging Rehearsal Buffers to Enable Efficient Data-Parallel Continual Learning*. In preparation, 2023, to be submitted.

3.2. Software

Distributed Rehearsal Buffer. As part of Axis 2 (Towards Continual Learning at Scale JLESC project), Thomas Bouvier, PhD student, Inria (KerData) has been developing a software prototype for a Distributed Rehearsal Buffer in close collaboration with Bogdan Nicolae, Computer Scientist, ANL. This experimental code uses Mochi (<https://www.mcs.anl.gov/research/projects/mochi/>), a software component developed at ANL. This work benefits from the support of the maintainers of Mochi: Matthieu Dorier, Software Development Specialist, ANL and Philip Carns, Computer Scientist, ANL. Currently this software is intended for internal use, as a proof of concept. Repository: <https://gitlab.inria.fr/Kerdata/Kerdata-Codes/distributed-rehearsal-buffer>.

3.3. Patents

N/A.

3.4. Demos & videos

Daniel Rosendo – *Instrumentation with Trovi and E2Clab*, G5K User Group Meeting, 2021

<https://drive.google.com/file/d/1ulsqrJKVBQIKfRd69ykpKtEzFVA8HQBq/view?resourcekey>

This demo is related to Axis 2 – topic: Supporting reproducible experiments across the Computing Continuum

3.5. Current position of students and postdocs involved in the associate team

[Nathanaël Cherièr](#) defended his thesis in December 2019. He is now researcher at Microsoft Research Cambridge.

[Orçun Yildiz](#) is now an Assistant Computer Scientist at ANL, Argonne, USA.

[Pierre Matri](#) is now a Data Technologist, working in the New York.

[Valentin Reis](#) is now Software Engineer at Groq, Austin, Texas.

[Hugo Chaugier](#) is now a machine learning engineer at Delfox, France.

The other PhD students involved are continuing their respective PhD theses.

3.6. Partnerships with companies (if any)

N/A

3.7. Other forms of impact

Kate Keahey served as a referee of the PhD thesis of Michael Mercier (Inria, DataMove): *Contribution to the infrastructure convergence between high performance computing and big data processing*. Defended in July 2019.

4. Future of the partnership

Building on the promising progress of UNIFY after the pandemics, a new Associated Team proposal UNIFY 2 has been submitted in 2022 and is under evaluation. Partners and topics are in the continuity of UNIFY.

5. Self-Assessment

UNIFY is a follow-up to a successful series of collaborations initiated between the KerData team and partners from ANL, started as early as 2010, extended to DataMove in 2019. This collaboration was highlighted as an example of a long-term collaboration within JLESC, its activities were the subject of a keynote presentation in 2020. UNIFY has been impacted by the pandemics and its activity were limited to virtual interactions in 2020-2021, but were strongly resumed in 2022 through thanks to 2 3-month internships. We are confident that the dynamics has the potential to grow to our “usual” standards in the framework of the future UNIFY 2 Associate Team.

6. Feed-back on the Associate Team program

The Associate Team Program is an extremely valuable tool for initiating and maintaining international collaborations, and for providing PhD students with the opportunity to get involved in these collaborations through long internships, and thus start building a valuable international network. We consider it as a very important asset for Inria, which should reinforce its financial support for this program. Even if the pandemics may have justified reduced funding in recent years, as traveling was not possible, the end of the pandemics should lead to restoring funding for this program at least to its traditional levels, to avoid missing valuable opportunities for increased synergies at international level.