

Maxime Gonthier

POST-DOCTORAL FELLOW

✉ maxjp.gr@gmail.com | 🏠 <https://maximegonthier.github.io/#>

Education

ENS de Lyon

Lyon, France

DOCTORATE IN COMPUTER SCIENCE

2020 - 2023

- Supervised by Loris Marchal (senior research scientist at CNRS) and Samuel Thibault (full professor at the University of Bordeaux) on the topic "Scheduling Under Memory Constraints in Task-Based Runtime Systems": <https://theses.hal.science/tel-04260094>.
- Designed schedulers that distribute and order data-sharing tasks across multiple GPUs to minimize data transfers over the memory bus, together with a custom eviction policy, all implemented in the StarPU runtime.
- Evaluated the algorithms both in simulation (SimGrid) and in real experiments on linear algebra applications, and studied their complexity, data-transfer lower bounds, and optimal behavior.

Swiss National Supercomputing Centre

Lugano, Switzerland

CSCS SUMMER SCHOOL

2023

- Training on developing scientific applications in CUDA.

HiPEAC

Fiuggi, Italy

INTERNATIONAL SUMMER SCHOOL ON ADVANCED COMPUTER ARCHITECTURE AND COMPILATION FOR HIGH-PERFORMANCE EMBEDDED SYSTEMS

2021 - 2022

- Training on programming models, high-performance computing, compilation, and security.

University of Versailles St-Quentin

Versailles, France

MASTER'S DEGREE IN ALGORITHMS AND MODELING AT THE INTERFACE OF SCIENCES

2018 - 2020

- Studied graph algorithms, game theory, distributed algorithms, databases, ranking methods, simulation, and networks.

University of Versailles St-Quentin

Versailles, France

BACHELOR'S DEGREE IN COMPUTER SCIENCE

2015 - 2018

Experiences

INRIA - TADaM

Bordeaux, France

PostDoc

2026 - Present

I am currently working as a post-doctoral fellow at the French National Institute for Research in Digital Science and Technology (INRIA) of Bordeaux, supervised by Francieli Boito (<https://www.labri.fr/perso/fzanonboito/>). Here is the main project of this postdoc:

- **Leveraging trace information for better I/O management** With the growing scale of HPC applications and the rise of AI and digital-twin workloads, I/O has become a major bottleneck due to the historical gap between processing and storage speeds. The goal of this project is to design a smarter I/O infrastructure that adapts to applications' characteristics by leveraging traces from production HPC systems such as Santos Dumont and PlaFRIM. My work covers four sub-problems: extending the trace-analysis classification of temporal I/O patterns to more years of data, correlating these patterns with application domains (e.g., whether AI applications are more likely to perform periodic reads), automating the clustering so it can run at execution time, and using the resulting classifications to predict future I/O activity for I/O scheduling, data placement, and burst buffer management.

University of Chicago and Argonne National Laboratory - Globus Labs

Chicago, USA

PostDoc

2024 - 2025

I worked as a post-doctoral fellow at the University of Chicago, under the supervision of Ian Foster and Kyle Chard. Here are two projects I worked on:

- **Dynamic data replication for heterogeneous storage nodes:** Current storage systems use erasure coding to ensure data resilience, but they rely on static parameters that limit optimization for encoding and decoding durations. These static settings are not well-suited to heterogeneous storage nodes with varying sizes and failure rates. To address these challenges, we are developing new dynamic algorithms tailored to heterogeneous storage systems. This research is a collaboration with Bogdan Nicolae from Argonne National Laboratory, Dante D. Sanchez-Gallegos and Jesus Carretero from Carlos III University in Madrid, and researchers and students from UChicago: Valerie Hayot-Sasson, J. Gregory Pauloski, Haochen Pan, Kyle Chard, Hai Duc Nguyen and Ian Foster.
- **Scheduling for mixed criticality tasks:** This project is part of a FACCTS travel grant awarded to foster collaboration with Professors Loris Marchal and Frédéric Vivien from France. The objective is to advance theoretical research on scheduling tasks with mixed criticality across homogeneous processors. To accomplish this, we developed approximations based on heuristics that achieve near-optimal performance while satisfying deadlines when using mixed workloads. We are now focusing on solving the same problem under a scenario with task failure. The question then becomes: How can we effectively duplicate task computations to ensure timely completion without overcrowding the processors?

Uppsala University - Ångström Laboratory

Uppsala, Sweden

RESEARCH STAY

2022

Collaboration with Elisabeth Larsson and Carl Nettelblad during a three-month stay in Uppsala

- Cluster job schedulers are generally unaware of jobs sharing large input files, which can occur in data-intensive scenarios. We investigated how to design an input file-aware job scheduler capable of distributing the load between compute nodes while reusing as much as possible of an input file already loaded into a node's memory.

Inria Bordeaux - STORM

Bordeaux, France

RESEARCH INTERNSHIP

2020

Scheduling for task-based runtime systems under memory constraints - Supervised by Loris Marchal and Samuel Thibault.

- Create and implement an offline scheduler capable of reducing the amount of data transfer in the StarPU runtime

Skills

Programming languages C, C++, Python, Bash, some Java.

HPC & parallel computing

StarPU runtime, OpenMP, CUDA, SimGrid.

Scheduling & systems

Task-based runtime systems, batch schedulers, I/O scheduling, data placement, erasure coding.

Tools

Git, Linux, LaTeX, GDB, Valgrind, Slurm.

Languages

French (native), English (fluent).

Publications

INTERNATIONAL PEER-REVIEWED JOURNALS

J1 *Throughput Estimation of Data Transport Networks From Digital Twin Measurements*

IEEE Transactions on Machine Learning in Communications and Networking (TMLCN), vol. 4, pp. 846–863, 2026

Nageswara S. V. Rao, Anees Al-Najjar, Rajkumar Kettimuthu, Ian Foster, Kyle Chard, Ryan Chard, Amal Gueroudji, Matthieu Dorier, Valerie Hayot-Sasson, Hai Duc Nguyen, Maxime Gonthier, Tekin Bicer, Haochen Pan, Eliu Huerta, Bogdan Nicolae, Parth Patel, and Justin M. Wozniak
<https://doi.org/10.1109/TMLCN.2026.3692126>

- J2 *Radio Afterglow Detection and AI-driven Response (RADAR): A Federated Framework for Gravitational-wave Event Follow-up*

The Astrophysical Journal Supplement Series (ApJS), vol. 280, no. 2, art. 71, 2025

Parth Patel, Alessandra Corsi, E. A. Huerta, Kara Merfeld, Victoria Tiki, Zilinghan Li, Tekin Bicer, Kyle Chard, Ryan Chard, Ian Foster, Maxime Gonthier, Valerie Hayot-Sasson, Hai Duc Nguyen, Haochen Pan
<https://doi.org/10.3847/1538-4365/adfbea>

- J3 *Toward a persistent event-streaming system for high-performance computing applications*

Frontiers in High Performance Computing (Front. High Perform. Comput.), 2025

Matthieu Dorier, Amal Gueroudji, Valerie Hayot-Sasson, Hai Duc Nguyen, Seth Ockerman, Rafael Souza, Tekin Bicer, Haochen Pan, Philip Carns, Kyle Chard, Ryan Chard, Maxime Gonthier, Eliu Huerta, Ben Lenard, Bogdan Nicolae, Parth Patel, Justin Wozniak, Ian Foster, Nageswara S. V. Rao, and Robert B. Ross
<https://doi.org/10.3389/fhpcp.2025.1638203>

- J4 *A scheduler to foster data locality for GPU and out-of-core task-based linear algebra applications*

Journal of Parallel and Distributed Computing (JPDC), vol. 206, p. 105170, 2025

Maxime Gonthier, Loris Marchal, and Samuel Thibault
<https://doi.org/10.1016/j.jpdc.2025.105170>

- J5 *Taming data locality for task scheduling under memory constraint in runtime systems*

Future Generation Computer Systems (FGCS), 2023

Maxime Gonthier, Loris Marchal, and Samuel Thibault
<https://inria.hal.science/hal-03623220v2/document>

INTERNATIONAL PEER-REVIEWED CONFERENCES

- C1 *Icicle: Scalable Metadata Indexing and Real-Time Monitoring for HPC File Systems*

ISC High Performance, 2026

Haochen Pan, Ryan Chard, Song Young Oh, Maxime Gonthier, Valerie Hayot-Sasson, Geoffrey Lentner, Joe Bottiglieri, Rachana Ananthakrishnan, Kyle Chard, Ian Foster
<https://arxiv.org/abs/2604.10295>

- C2 *Core Hours and Carbon Credits: Incentivizing Sustainability in HPC*

Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '25), 2025

Alok Kamatar, Maxime Gonthier, Valerie Hayot-Sasson, André Bauer, Marcin Copik, Raul Castro Fernandez, Torsten Hoefler, Kyle Chard, Ian Foster
<https://doi.org/10.1145/3712285.3759858>

- C3 *Addressing Reproducibility Challenges in HPC with Continuous Integration*

Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '25), 2025

Valerie Hayot-Sasson, Nathaniel Hudson, André Bauer, Maxime Gonthier, Ian Foster, Kyle Chard
<https://doi.org/10.1145/3712285.3759874>

- C4 *Deadline-Aware Scheduling of Mixed-Criticality Tasks*

Proceedings of the 54th International Conference on Parallel Processing (ICPP '25), 2025

Maxime Gonthier, Kyle Chard, Ian Foster, Loris Marchal, Frédéric Vivien
<https://doi.org/10.1145/3754598.3754639>

- C5 *D-Rex: Heterogeneity-Aware Reliability Framework and Adaptive Algorithms for Distributed Storage*
Proceedings of the 39th ACM International Conference on Supercomputing (ICS '25), 2025
Maxime Gonthier, Dante D. Sanchez-Gallegos, Haochen Pan, Bogdan Nicolae, Sicheng Zhou, Hai Duc Nguyen, Valerie Hayot-Sasson, J. Gregory Pauloski, Jesus Carretero, Kyle Chard, Ian Foster
<https://doi.org/10.1145/3721145.3730412>
- C6 *Optimizing Fine-Grained Parallelism Through Dynamic Load Balancing on Multi-Socket Many-Core Systems*
IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2025
Wenyi Wang, Maxime Gonthier, Poornima Nookala, Haochen Pan, Ian Foster, Ioan Raicu, Kyle Chard
<https://ieeexplore.ieee.org/document/11078401>
- C7 *DynoStore: A wide-area distribution system for the management of data over heterogeneous storage*
25th IEEE International Symposium on Cluster, Cloud, and Internet Computing (CCGrid), 2025
Dante D. Sanchez-Gallegos, J. L. Gonzalez-Compean, Maxime Gonthier, Valerie Hayot-Sasson, J. Gregory Pauloski, Haochen Pan, Kyle Chard, Jesus Carretero, Ian Foster
<https://ieeexplore.ieee.org/document/11044834>
- C8 *WRATH: Workload Resilience Across Task Hierarchies in Task-based Parallel Programming Frameworks*
25th IEEE International Symposium on Cluster, Cloud, and Internet Computing (CCGrid), 2025
Sicheng Zhou, Zhuozhao Li, Valerie Hayot-Sasson, Haochen Pan, Maxime Gonthier, J. Gregory Pauloski, Ryan Chard, Ian Foster, Kyle Chard
<https://ieeexplore.ieee.org/document/11044847>
- C9 *TaPS: A Performance Evaluation Suite for Task-based Execution Frameworks*
20th IEEE International Conference on e-Science, 2024
J. Gregory Pauloski, Valerie Hayot-Sasson, Maxime Gonthier, Nathaniel Hudson, Haochen Pan, Sicheng Zhou, Ian Foster, Kyle Chard
<https://ieeexplore.ieee.org/document/10678702>
Best paper award.
- C10 *An Empirical Investigation of Container Building Strategies and Warm Times to Reduce Cold Starts in Scientific Computing Serverless Functions*
20th IEEE International Conference on e-Science, 2024
André Bauer, Maxime Gonthier, Ryan Chard, Haochen Pan, Daniel Grzenda, Martin Straesser, J. Gregory Pauloski, Alok Kamatar, Matthew E. Baughman, Nathaniel Hudson, Ian Foster, Kyle Chard
<https://ieeexplore.ieee.org/abstract/document/10678668>
- C11 *Memory-Aware Scheduling of Tasks Sharing Data on Multiple GPUs with Dynamic Runtime Systems*
36th IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2022
Maxime Gonthier, Loris Marchal, Samuel Thibault
<https://inria.hal.science/hal-03552243v1/document>

INTERNATIONAL PEER-REVIEWED WORKSHOPS

- W1 *Octopus: Experiences with a Hybrid Event-Driven Architecture for Distributed Scientific Computing*
Fault Tolerance for HPC at eXtreme Scales Workshop (FTXS) at SuperComputing, 2024
Haochen Pan, Ryan Chard, Sicheng Zhou, Alok Kamatar, Rafael Vescovi, Valerie Hayot-Sasson, André Bauer, Maxime Gonthier, Kyle Chard, Ian Foster
<https://arxiv.org/pdf/2407.11432>
- W2 *Data-Driven Locality-Aware Batch Scheduling*
26th Workshop on Advances in Parallel and Distributed Computational Models (APDCM) at the 38th IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2024
Maxime Gonthier, Elisabeth Larsson, Loris Marchal, Carl Nettelblad, Samuel Thibault
<https://ieeexplore.ieee.org/document/10596393>

W3 *Diaspora: Resilience-Enabling Services for Real-Time Distributed Workflows*

1st international workshop on Near Real-time Data Processing for Interconnected Scientific Instruments (NRDPISI-1) at the 20th IEEE International Conference on e-Science, 2024

Bogdan Nicolae, Justin M. Wozniak, Tekin Bicer, Hai Nguyen, Parth Patel, Haochen Pan, Amal Gueroudji, Maxime Gonthier, Valerie Hayot-Sasson, Eliu Huerta, Kyle Chard, Ryan Chard, Matthieu Dorier, Nageswara S. V. Rao, Anees Al-Najjar, Alessandra Corsi, Ian Foster

<https://ieeexplore.ieee.org/abstract/document/10678669>

W4 *Locality-Aware Scheduling of Independent Tasks for Runtime Systems*

COLOC: 5th workshop on data locality at the 27th International European Conference on Parallel and Distributed Computing (EuroPar), 2021

Maxime Gonthier, Loris Marchal, Samuel Thibault

<https://hal.science/hal-03290998v1/document>

POSTERS

A1 *Icicle: Scalable Metadata Indexing and Real-Time Monitoring for HPC File Systems*

Research poster at the 13th Greater Chicago Area Systems Research Workshop (GCASR), 2026

Haochen Pan, Ryan Chard, Song Young Oh, Maxime Gonthier, Valérie Hayot-Sasson, Geoffrey Lentner, Joe Bottiglieri, Rachana Ananthakrishnan, Kyle Chard, Ian Foster

https://sc25.supercomputing.org/proceedings/posters/poster_files/post147s2-file2.pdf

A2 *Exploring Fine-Grained Parallelism in Dataflow Runtime Systems on Multi-Socket Many-Core Systems*

Research poster (with short paper) at SC, 2025

Wenyi Wang, Maxime Gonthier, Haibin Lai, Poornima Nookala, Haochen Pan, Ian Foster, Ioan Raicu, Kyle Chard

https://sc25.supercomputing.org/proceedings/posters/poster_files/post147s2-file2.pdf

A3 *Continuous Reproducibility in HPC: Addressing Challenges Through Automation*

Poster at the Greater Chicago Area Systems Research Workshop (GCASR), 2025

Valerie Hayot-Sasson, Nathaniel Hudson, André Bauer, Maxime Gonthier, Ian Foster, Kyle Chard

A4 *Memory-Aware Scheduling Of Tasks Sharing Data On Multiple GPUs*

Peer-reviewed poster at ISC High Performance, 2023,

Maxime Gonthier, Loris Marchal, and Samuel Thibault

<https://inria.hal.science/hal-04090595v1/document>

A5 *Memory-Aware Scheduling Of Tasks Sharing Data On Multiple GPUs*

Invited poster at HiPEAC ACACES - 18th International Summer School on Advanced Computer Architecture and Compilation for High-performance Embedded Systems, 2022

Maxime Gonthier, Loris Marchal, and Samuel Thibault

<https://inria.hal.science/hal-04090607v1/document>

A6 *Locality-Aware Scheduling Of Independent Tasks For Runtime Systems*

Invited poster at HiPEAC ACACES - 17th International Summer School on Advanced Computer Architecture and Compilation for High-performance Embedded Systems, 2021,

Maxime Gonthier, Loris Marchal, and Samuel Thibault

<https://inria.hal.science/hal-04090604v1/document>

NATIONAL PEER-REVIEWED CONFERENCES

N1 *Exploiting data locality to maximize the performance of data-sharing tasksets,*

Maxime Gonthier,

CompAS - Conférence francophone d'informatique en Parallélisme, Architecture et Système, 2023

<https://inria.hal.science/hal-04090634v1/document>

SUBMITTED

S1 *To See or to Hear First: The Energy Cost of Prompting Order*

Submitted article at IEEE International eScience Conference (eScience), 2026

Laya Fakher, [Maxime Gonthier](#), Valerie Hayot-Sasson, Thomas Prantl, SeongHo Parker, and André Bauer

S2 *AERO: An autonomous platform for continuous research*

Submitted article at IEEE International eScience Conference (eScience), 2026

Valerie Hayot-Sasson, Abby Stevens, Nicholson Collier, Sudershan Sridhar, Kyle Conroy, J. Gregory Pauloski, Yadu Babuji, [Maxime Gonthier](#), Nathaniel Hudson, Dante D. Sanchez-Gallegos, Ian Foster, Jonathan Ozik, Kyle Chard

<https://arxiv.org/abs/2505.18408>

HPC Wire 2024 prize: <https://www.hpcwire.com/2024-readers-editors-choice-best-hpc-response-to-social-media/>

S3 *Taskinfer: Exploring Task-based Parallelism in LLM Inference on Multi-Socket Many-Core Systems*

Submitted article at SuperComputing, 2026

Haibin Lai, Wenyi Wang, Poornima Nookala, [Maxime Gonthier](#), Ioan Raicu, Zhuozhao Li, Kyle Chard, Ian Foster

RESEARCH REPORTS

R1 *Experiences with Model Context Protocol Servers for Science and High Performance Computing*

Preprint (arXiv:2508.18489), 2025

Haochen Pan, Ryan Chard, Reid Mello, Christopher Grams, Tanjin He, Alexander Brace, Owen Price Skelly, Will Engler, Hayden Holbrook, Song Young Oh, [Maxime Gonthier](#), Michael Papka, Ben Blaiszik, Kyle Chard, Ian Foster

<https://arxiv.org/abs/2508.18489>

R2 *Locality-aware batch scheduling of I/O intensive workloads*

Research report, extended version of [W2], 2023

[Maxime Gonthier](#), Elisabeth Larsson, Loris Marchal, Carl Nettelblad, and Samuel Thibault,

<https://inria.hal.science/hal-03993118v1/document>

R3 *A generic scheduler to foster data locality for GPU and out-of-core task-based applications*

Research report, 2023

[Maxime Gonthier](#), Loris Marchal, and Samuel Thibault,

<https://inria.hal.science/hal-04146714v1/document>

R4 *Locality-Aware Scheduling of Independent Tasks for Runtime Systems*

Research report, 2022

[Maxime Gonthier](#), Loris Marchal, and Samuel Thibault,

<https://inria.hal.science/hal-03144290v7/document>

Awards, Grants, & Funding

2024 - **France and Chicago Collaborating in the Sciences**, France Chicago Center

\$ 14,000

Travel grant (France and Chicago Collaborating in the Sciences or "FACCTS") aimed at fostering collaboration between researchers from France and Chicago. I wrote the funding request for a research project involving Professors Ian Foster (University of Chicago and Argonne National Laboratory), Loris Marchal (CNRS and ENS de Lyon), Kyle Chard (University of Chicago and Argonne), and Frédéric Vivien (INRIA and ENS de Lyon). More information at <https://fcc.uchicago.edu/faccts/>

2024 - **Early Career Professionals Travel Grant**, SIGHPC

\$ 600 + conference registration

Travel grant to help early career researchers attend conferences. Only 6 grants were awarded in 2024. More information at <https://www.sighpc.org/opportunities/travel-grants>

2024 - **Best Paper Award**, 20th IEEE International Conference on e-Science

Out of 80 submitted papers, 30 were accepted, 3 were nominated for the Best Paper Award, and only one received this distinction. More information at <https://www.escience-conference.org/2024/>. Article: <https://ieeexplore.ieee.org/document/10678702>

Invited Talks

2026 - **TADaaM team, INRIA Bordeaux**, "Towards Highly Compatible I/O-Aware Workflow Scheduling on HPC Systems". Bordeaux, France.

2026 - **CONCACE Seminar**, INRIA Bordeaux. Bordeaux, France.

2025 - **Técnico (University of Lisbon)**, INESC-ID team. Lisbon, Portugal.

2025 - **Globus Labs**, "Understanding GPU Memory Corruption at Extreme Scale: The Summit Case Study". Chicago, USA.

2024 - **Illinois Institute of Technology**, Department of Computer Science. Chicago, USA.

2024 - **University of Chicago**, Department of Computer Science. Chicago, USA.

2024 - **Johns Hopkins University**, Members of the US Department of Energy (DOE) Diaspora project. Baltimore, USA.

2024 - **ANEO**, Cloud and HPC consulting company. Boulogne-Billancourt, France. <https://www.aneo.eu/evenement/incontournables/seminaire-recherche-expertise> - https://www.youtube.com/watch?v=EPIE7LK_LgU&t=1s

2024 - **17th Scheduling for Large-Scale Systems Workshop**. Aussois, France. <https://graal.ens-lyon.fr/~abenoit/aussois24/>

2023 - **Institut de Recherche en Informatique de Toulouse (IRIT)**, APO team. Toulouse, France.

2023 - **University of Versailles**, DAVID team. Versailles, France.

2023 - **INRIA Grenoble**, POLARIS and DATAMOVE teams. Grenoble, France.

2023 - **15th JLESC Workshop**. Bordeaux, France. <https://events.hifis.net/event/617/>

2023 - **SOLHARIS Plenary**, Members of the ANR SOLHARIS project <https://www.irit.fr/solharis/>. Bordeaux, France.

2023 - **ICube - University of Strasbourg**, CAMUS team. Strasbourg, France.

2022 - **Uppsala University**, Scientific Computing Division. Uppsala, Sweden.

2022 - **15th Scheduling for Large-Scale Systems Workshop**. Fréjus, France. <https://scheduling2022.sciencesconf.org/>

2022 - **INRIA Bordeaux**, STORM, TOPAL, TADaaM and CONCACE teams. Bordeaux, France.

2021 - **ENS Lyon**, ROMA team. Lyon, France.

Tutorials

Teaching Experience

During my PhD, I worked as a teaching assistant at the University of Lyon and the engineering school Enseirb-Matmeca. The table below summarizes the courses I taught over three years, for a total of 132 hours.

Audience	Course Title	Type	Year	Hours
Bachelor's Year 1 Computer Science, University of Lyon 1	Operating Systems	Lab	2021	18h
2nd year Engineering, Enseirb-Matmeca	Network Programming	Project	2021 and 2022	66h
2nd year Engineering, Enseirb-Matmeca	Internship Tutoring and Defense Committee Member	Other	2022 and 2023	8h
1st year Engineering, Enseirb-Matmeca	Algorithms	Integrated Course	2022	20h
1st year Engineering, Enseirb-Matmeca	Computer Architecture	Lab	2022	20h
Total				132h

Supervision & Mentoring

PhD Students

- **Alok Kamatar**, University of Chicago (Article published at the SC conference [C2], January 2024 - July 2025 - 25%)
Core Hours and Carbon Credits: Incentivizing Sustainability in HPC I am working with Alok to improve energy efficiency in high-performance computing through new user credit allocation models. We explored the impact of new resource allocation systems based on carbon credits to encourage users to adopt more energy-efficient behaviors. A tool was developed to predict carbon consumption and computational times for certain common applications. We designed a new allocation system taking into account CO2 emissions, hardware lifecycle, and node utilization rates by users. We then simulated the behavior of several typical users. Additionally, we conducted a survey to assess user awareness of their consumption, with respondents from four continents, and developed a game prototype. This game demonstrates that with appropriate information, users can reduce their energy impact without significantly sacrificing computational time.
- **Dante D. Sanchez-Gallegos**, University of Madrid Carlos III (Article published at the CCGrid conference [C7], January 2024 - July 2025 - 20%)
Building a wide-area distribution system for the management of data over heterogeneous storage Development of multi-user data transfer tools with erasure coding. Collaboration on algorithmic solutions for encoding and distributing data fragments over heterogeneous storage nodes, aiming to reduce encoding time and optimize load balancing.
- **J. Gregory Pauloski**, University of Chicago (Article published at the eScience conference and best paper [C9], February 2024 - July 2025 - 15%)
Programming the Continuum: Towards Better Techniques for Developing Distributed Science Applications I helped Greg develop TaPS, a task-based application library to evaluate different execution environments and data transfers.

- **Haochen Pan**, University of Chicago (Articles: FTXS workshop [W1], ISC [C1], and preprint [R1]; January 2024 - November 2025 - 20%)
Building a Hybrid Event-Driven Architecture for Distributed Scientific Computing Contribution to a project aimed at enhancing the resilience of applications in the U.S. Department of Energy. My role includes identifying and formalizing use cases. I also work with Haochen on Icicle, a scalable metadata indexing and real-time monitoring system for HPC file systems, and on Model Context Protocol (MCP) servers that expose scientific and HPC capabilities to large language model agents.
- **Wenyi Wang**, University of Chicago (Article published at the IPDPS conference [C6], February 2024 - November 2025 - 25%)
Optimizing Fine-Grained Parallelism Through Dynamic Load Balancing on Multi-Socket Many-Core Systems Development of a lock-free execution system for tasks. Collaboration on the design of a synchronization-free scheduling strategy, with a particular focus on load balancing.
- **Shu Shi**, University of Chicago (November 2024 - November 2025 - 25%)
Multi-LLM Serving for HPC With the rise of Large Language Models (LLMs), supercomputers face increasing demand for their GPUs, which are already in limited supply. I have been assisting Shu in the development of scheduling solutions aimed at optimizing the placement and reuse of machine learning models already loaded into memory on computation nodes. The goal is to reduce model reload costs and maximize the use of available GPUs, even when the GPU quantity is undersized compared to the workload.
- **Haibin Lai**, Southern University of Science and Technology (Article submitted to SuperComputing [S3], 2025)
Taskinfer: Exploring Task-based Parallelism in LLM Inference on Multi-Socket Many-Core Systems I have been co-supervising Haibin on exploiting task-based parallelism to accelerate LLM inference on multi-socket many-core systems.

Undergraduate Students

- **Sicheng Zhou**, Undergraduate Student, University of Chicago (Article published at the CCGrid conference [C8], March 2024 - June 2024 - 20%)
Workload Resilience Across Task Hierarchies in Task-based Parallel Programming Frameworks Joint work on improving failure management in task-based programming systems. The project focuses on failure categorization and developing tailored responses, including a real-time monitoring system.
- **Karaen Senthilkumar**, Undergraduate Student, University of Chicago (Submitted poster at SciFM 2026, September 2025 - Present - 40%)
Fairness-Aware Scheduling for Bursty HPC Workloads via Hierarchical Algorithms I am supervising Karaen on the design of hierarchical scheduling algorithms that provide fairness for bursty HPC workloads.

Engineers

- **Abdraman Mahamat**, Engineer (April 2026 - Present - 20%)
Cross-system prediction of stripe count I am working with Abdraman on predicting the file stripe count across different HPC storage systems, in order to improve I/O performance and data placement.

Mobility

My PhD was conducted at two different locations. In 2020 I worked at ENS Lyon, within the ROMA team. I benefited from the theoretical scheduling expertise of its members. Then, from 2021 to 2023, I continued my work at INRIA Bordeaux. The STORM team, of which I was a part, has extensive experience in performance analysis,

programming, and especially runtimes, thanks to the local tool StarPU. This allowed me to implement, improve, and evaluate the theoretical scheduling strategies we had developed in Lyon. This collaboration led to the production of the following articles (see the "Publications" section): [J5, C11, W4, A4, A5, A6, N1, J4, R3, R4].

In 2022, I also joined the Scientific Computing team at Uppsala University (Sweden) for a research stay. This research stay was an opportunity to collaborate with Elisabeth Larsson and Carl Nettelblad. Initially, my work focused on developing a simulator for a job management system on a supercomputer. Later, my contribution was the creation of new supercomputer job schedulers capable of reducing the amount of data transfers, thus improving both resource utilization and user satisfaction. The scientific output is as follows: [W2, R2].

From 2024 to 2025, I was a postdoc at Globus Labs at the University of Chicago, under the supervision of Ian Foster and Kyle Chard. My work in Chicago provided an opportunity to collaborate with various labs. I collaborated with Bogdan Nicolae and Hai Duc Nguyen from Argonne National Laboratory (see [W3, C5]). In Chicago, I also worked with a researcher from the Illinois Institute of Technology (IIT), associated with publication [C10]. I was actively involved in projects with researchers from the Universidad Carlos III de Madrid, focusing on erasure coding and data replication (see publication [C7]). My postdoctoral research also allowed me to collaborate with researchers from ETH Zurich like Marcin Copik (see [C2]).

Since 2026, I am a postdoctoral fellow at INRIA Bordeaux within the TADaaM team, under the supervision of Francieli Boito. My work focuses on the project "Leveraging trace information for better I/O management", which aims to design a smarter I/O infrastructure for HPC systems by leveraging traces from production platforms such as Santos Dumont and PlaFRIM, together with machine learning and large language models to analyze these traces and predict future I/O activity. This new mobility allows me to continue collaborating with my former colleagues at INRIA Bordeaux while strengthening my ties with researchers in Brazil (LNCC) and Argonne National Laboratory.

Personal Technological Developments ---

1. Data-Aware Reactive Task Scheduling for StarPU (DARTS)

- **Contribution:** I designed and implemented almost the entire code. Samuel Thibault and Nathalie Furmento participated in debugging the code.
- **Development duration:** 2 years
- **URL:** https://gitlab.inria.fr/starpu/starpu/-/blob/master/src/sched_policies/darts.c?ref_type=heads
- **Languages:** C (StarPU runtime)
- **Objective:** The StarPU runtime includes a set of schedulers. The goal of DARTS is to address the needs of applications whose memory requirements exceed the memory available on the computing units used. Thus, DARTS employs data locality techniques to minimize transfers and consequently improve performance. DARTS is accompanied by a custom eviction policy, which required implementing support in StarPU for creating new eviction policies.
- **Impact:** Thanks to the widespread use of StarPU and the generalist approach of DARTS, it can be used not only for all linear algebra applications but also for any applications that can be expressed as tasks sharing data. Furthermore, the strategy of DARTS and its eviction policy are not intrinsically tied to StarPU and could be replicated on other runtimes.
- **Audience:** DARTS is integrated into the main branch of StarPU and is part of the StarPU 1.5.0 release. As a result, DARTS can be selected as a scheduling and eviction policy by all software using StarPU (<https://starpu.gitlabpages.inria.fr/#software>). The ExaGeostat project is a geostatistics computing

tool for large-scale architectures. It already uses StarPU to solve large linear algebra applications. The ability of DARTS to handle limited memory would be ideal for accelerating such applications. Lastly, some companies, such as ANEO, develop runtimes and are interested in incorporating the intuition behind DARTS into their own applications (<https://aneoconsulting.github.io/ArmoniK/armonik>).

- **Associated Publications:** [J5, C11, A4, A5, N1, J4, R3]

2. Hierarchical Fair Packing for StarPU (HFP)

- **Contribution:** I designed and implemented almost the entire code. Samuel Thibault and Nathalie Furmento participated in debugging the code.
- **Development duration:** 2 years
- **URL:** https://gitlab.inria.fr/starpu/starpu/-/blob/master/src/sched_policies/HFP.c?ref_type=heads
- **Languages:** C (StarPU runtime)
- **Objective:** Hierarchical Fair Packing is a scheduling strategy in the StarPU runtime aimed at improving performance under memory constraints in an offline setting. This means that scheduling must be done in advance rather than dynamically during application execution.
- **Impact:** HFP is useful for regular systems with large amounts of data to load. One objective is to identify application cases where it would be beneficial. I am particularly considering the resolution of large sparse matrices.
- **Audience:** Integrated into the StarPU 1.5.0 release.
- **Associated Publications:** [J5, C11, W4, A6, R4].

3. Visualization Tool for StarPU Schedulers

- **Contribution:** I designed and implemented the entire code.
- **Development duration:** 2 years
- **URL:** https://gitlab.inria.fr/starpu/starpu/-/tree/master/tools/darts?ref_type=heads
- **Languages:** C
- **Objective:** Allows visualization of task execution order for StarPU schedulers. Also enables visualization of data loads associated with each task. Supports outer product, GEMM, and Cholesky factorization.
- **Impact:** Personal use. Will help developers and researchers understand and debug scheduler behavior by making task ordering and data movement visible. Also used to analyze the behavior of DARTS, HFP and other state of the art schedulers present in StarPU.
- **Audience:** Integrated into the StarPU 1.5.0 release.
- **Associated Publications:** [J5, C11, J4, R3, R4].

4. Supercomputer Workload Simulator

- **Contribution:** I designed and implemented the entire code.
- **Development duration:** 2 years
- **URL:** <https://github.com/MaximeGonthier/Locality-aware-batch-scheduling> and <https://github.com/MaximeGonthier/batch-simulator>
- **Languages:** C, Python

- **Objective:** A job scheduling simulator for supercomputers. It allows retrieving job submission logs and adding scheduling policies while taking data loads into account.
- **Impact:** Enables rapid prototyping and evaluation of new scheduling policies on real submission logs, without requiring access to a production supercomputer.
- **Audience:** Used to produce the experimental results reported in the associated publications.
- **Associated Publications:** [W2, C2, R2].

5. Data Replication and Distribution Simulator (DynoStore D-Rex)

- **Contribution:** I developed 70% of the project, while Dante Domizzi developed the remaining 30%.
- **Development duration:** 1 year
- **URL:** <https://github.com/dynostore/D-rex>
- **Languages:** C
- **Objective:** Simulates erasure coding to replicate data across storage nodes. We also simulated block mapping on a set of heterogeneous storage nodes and simulated storage node failures.
- **Impact:** Allows evaluating erasure-coding and data-placement strategies under storage heterogeneity and node failures before deploying on real systems.
- **Audience:** Used to produce the experimental results reported in the associated publication.
- **Associated Publications:** [C5].

Co-supervised Technological Developments

I have co-supervised the following softwares (ranked by software size/development duration in decreasing order):

- **Octopus:** <https://github.com/globus-labs/diaspora-event-sdk>. Developed by Haochen Pan. A cloud-edge communication software that enables sending, receiving, and processing messages between data producers and consumers. It facilitates the automatic triggering of events to process data while ensuring cloud storage. This is particularly useful for applications such as epidemiological models or autonomous laboratories ("self-driving labs," see <https://www.anl.gov/autonomous-discovery>).
- **DynoStore:** <https://github.com/dynostore/dynostore>. Developed by Dante D. Sanchez-Gallegos. DynoStore is a data transfer software between different infrastructures. It uses data containers to orchestrate storage and data movement across multiple sites. These containers have generic interfaces to connect sites using different storage systems (e.g., Amazon S3, Ceph, or Lustre).
- **XQueue:** https://gitlab.com/pnookala/gnu-openmp/-/tree/xtask?ref_type=heads. Developed by Wenyi Wang. XQueue is a lock-free task queue (thus without locks and synchronizations) that replaces the GNU task queue and its lock. XQueue significantly reduces execution time by eliminating this synchronization lock. Additionally, XQueue integrates NUMA-aware load balancing.
- **TaPS:** <https://github.com/proxystore/taps>. Developed by J. Gregory Pauloski. The Task Performance Suite allows the evaluation of task-based execution supports and data management systems through a set of real and synthetic scientific applications.
- **WRATH:** https://github.com/ClaudiaCumberbatch/resilient_compute. Developed by Sicheng Zhou. This software allows the creation of "compute pools," i.e., sets of computing resources in the cloud. These resources have specific properties: minimum storage size, geographical distribution, minimum number of nodes per site, etc. This is useful for multi-site applications (such as Federated Learning) that require a certain amount of memory and computing power across different locations. The goal here is to handle

different types of failures (memory saturation, loss of a compute node, user error, etc.). WRATH enables dynamic reconfiguration of workload partitioning and allows the addition or replacement of computing resources dynamically.

- **HPC Scheduling Game:** <https://game.funcx.org/>. Developed by Alok Kamatar. The objective is to collect information about supercomputer usage by users. In this game, you are a computer scientist with numerous jobs to execute, but you must complete them within a given deadline without exceeding your credit allocation. Your task is to determine which machine to submit each job to.
- **Icicle:** <https://github.com/globus-labs/icicle/>. Developed by Haochen Pan. A scalable metadata indexing and real-time monitoring system for HPC file systems, enabling fast search and live insight into very large file systems.
- **Science MCP servers:** <https://github.com/globus-labs/science-mcps>. Developed by Haochen Pan. A collection of Model Context Protocol (MCP) servers that expose scientific and HPC capabilities (data, compute, and workflow services) to large language model agents.

Collective Responsibilities

VOLUNTEERING AT CONFERENCES AND PARTICIPATION IN PROGRAM COMMITTEES

- **ISC High Performance 2026:** Session chair, Energy-Efficient Systems track.
- **Conférence francophone d'informatique en Parallélisme, Architecture et Système (COMPAS 2026):** Session chair, IO.
- **SuperComputing 2026:** Program committee (Algorithms track).
- **HIPC 2025:** Program committee.
- **ISPDC 2025:** Program committee.
- **SuperComputing 2025:** Member of the reproducibility committee.
- **ICDCS 2025:** Program committee.
- **EuroPar 2025:** Program committee.
- **APDCM 2025:** Program committee.
- **HotCloudPerf 2025:** Program committee.
- **SuperComputing 2024:**
 - Member of the reproducibility committee.
 - Managing an information booth on SIGHPC: ACM's Special Interest Group on High-Performance Computing.
- **IPDPS 2021:** Member of the volunteer team. My responsibilities included setting up sessions and providing technical assistance.

VOLUNTEERING AS A REVIEWER

- **Future Generation Computer Systems Journal 2025**
- **Future Generation Computer Systems Journal 2024**
- **IEEE Cluster Conference 2025**
- **IEEE Cluster Conference 2024**

- **Reproducibility of HPC experiments (repro-hpc) 2025:** <https://repro-hpc.github.io>

COMMITMENT TO INTEGRATION AND INCLUSION

- **Globus Labs:** I was part of the welcome committee, where I helped new students adapt to the lab, the city, and PhD work.
- **HandiManager (2019):** Participation in the HandiManager training program to promote the inclusion of people with disabilities in the professional environment.

SCIENTIFIC OUTREACH

- **My Thesis in 180 Seconds:** I participated in the "My Thesis in 180 Seconds" competition to learn how to share my Ph.D. research with a broader audience.
- **"Fête de la science":** I will take part in the 2026 edition, in which the research organizations (Inria, CNRS, etc.) partner with the academic Rectorat to make science accessible to everyone. We will engage young audiences across four departments of the Nouvelle Aquitaine region to discover science, research and make it exciting for students from geographically remote schools.