

# Jérôme Gurhem

92800 Puteaux, France

[jgurhem@aneo.fr](mailto:jgurhem@aneo.fr)

[LinkedIn](#)

[GitHub](#)

## ABOUT ME

HPC, Cloud and SaaS Expert Consultant with a Ph.D. in Computer Science and a strong background in scientific computing, parallel programming, distributed systems, and AI-driven platform design.

Currently Tech & Research Lead for [ArmoniK](#) at ANEO, an open-source multi-cloud compute orchestrator, with full ownership of the multi-year product roadmap, budget and FTE resource allocation, leading a technical team of ~10 people. Also leading the SaaS architecture and agentic AI orchestration for a GPU-native CFD simulation platform.

Proven track record spanning the full product lifecycle: architecture, development, R&D, strategic planning, pre-sales, training, customer support, and people management. Passionate about performance challenges, the industrialization of HPC solutions, and the integration of AI into complex engineering workflows.

## WORK EXPERIENCE

[August 2021 - Present](#) **ANEO, Boulogne-Billancourt, France** Product, Engineering & Research Lead, HPC/Cloud

**ArmoniK** — Open-Source HTC Orchestrator [August 2021 - Present](#)

### • Engineering

- Led the full lifecycle of ArmoniK, an open-source task-based high-throughput computing (HTC) framework for orchestrating large-scale, fault-tolerant workloads across distributed and elastic infrastructures, processing tens of millions of tasks per month in production for major financial institutions
- Architected and implemented its core on Kubernetes and Terraform, driving its evolution from initial design to production-grade deployments scaling to 12,000+ concurrent workers on a single instance within a single cloud region in hybrid multi-cloud environments

### • Product

- Owns ArmoniK's multi-year strategic product roadmap and oversees a portfolio of ArmoniK projects representing €700K–€1.2M in annual revenue, defining technical direction, feature prioritization, and delivering major releases on a 6-month cadence in alignment with client needs and market trends
- Manages ArmoniK's budget and FTE resource allocation, making investment decisions across R&D, engineering, and support — spanning headcount planning, contractor engagement, and cost optimization

### • Research & R&D

- Led R&D for ArmoniK, contributing to scientific publications, and leading Aneo's involvement in the FPT4 project (successor to InterFLOP)
- Co-supervised a PhD thesis on the formal modeling, verification, and optimization of ArmoniK, bridging academic research and industrial application

### • People

- Managed and mentored a cross-functional team of ~10 PhD researchers, engineers, and interns, fostering technical growth and delivering high-quality software components

### • Operations

- Scaled and structured Armonik's customer support operation for Crédit Agricole CIB and established it from the ground up for Natixis — across 30+ deployments (dev, integration, pre-prod, production) for 6 client teams; defined and delivered Level 3 support services, built issue-tracking and client follow-up workflows, and coordinated bug triage and corrective patches
- Actively participated in pre-sales engagements, contributing to technical proposals, architecture workshops with prospective clients, and commercial positioning of Armonik for regulated industries (finance, insurance)
- Created and launched Armonik's training offering, designing the curriculum and delivering 5+ multi-day sessions (presentations, hands-on labs, hackathons) to 5–25 participants per session, both internally and to external clients as a commercial service

### Agentic CFD Simulation Platform March 2026 – May 2026

- Co-designed the SaaS platform layer for an Agentic CFD Simulation Platform combining a GPU-native CFD solver (AWS ParallelCluster/PCS, NVIDIA H100 GPUs) and a vector database (PostgreSQL with pgvector); designed the platform architecture including identity and access management, usage tracking, platform governance, web-based user interface, and future Bring-Your-Own-Cloud (BYOC) deployment model with usage-based pricing
- Co-designed the agentic AI orchestration architecture using LLM-based agents (AWS Bedrock with Mistral models) and multi-agent orchestration protocols (MCP, A2A), enabling natural language-driven simulation workflows with human-in-the-loop control across the full simulation lifecycle (configuration, execution, monitoring, analysis); integrated observability via Langfuse and SigNoz
- Coordinated cross-team delivery across a 7-person team spanning the SaaS platform, the GPU/HPC infrastructure team, and the AI agent development team

### People & Organization August 2021 - Present

- Acted as Manager Référent, supporting consultants' career development through roadmap planning, feedback collection, and alignment with Aneo's culture and HR processes, complemented by hands-on technical coaching and long-term mentoring of junior engineers, PhD researchers, and interns
- Actively involved in strategic recruitment: defined job descriptions, conducted ~20 technical interviews per year, and contributed to onboarding and long-term hiring decisions

### March 2021 - July 2021 **CNRS USR3441 Maison de la simulation, Saclay, France** Postdoctoral researcher

- Contributed to the development and parallelization of a Graph Convolutional Neural Network (GCNN) training and inference library, optimizing for large-scale graph data processing
- Ported a hybrid MPI + OpenMP implementation of Sparse Matrix-Vector Multiplication (SpMV) to the A64FX architecture, leveraging SIMD vectorization for performance gains
- Conducted performance benchmarking and scalability analysis on the Fugaku supercomputer, evaluating efficiency across thousands of cores

### March 2018 - March 2021 **CNRS USR3441 Maison de la simulation, Saclay, France** Ph.D. Student

- Conducted a comprehensive survey and classification of task-based programming paradigms for high-performance computing applications
- Designed, implemented, and deployed task-parallel scientific applications on large-scale supercomputing platforms
- Performed experimental evaluations of block-based direct solvers for dense linear systems, sparse matrix-vector multiplication (SpMV), and Kirchhoff pre-stack depth migration in seismic imaging workflows
- Partially supported by TOTAL SA

### March 2017 - August 2017 **CNRS USR3441 Maison de la simulation, Saclay, France** Intern — Parallel Programming

- Developed and tested parallel applications using YML and XMP programming models on high-performance computing platforms
- Conducted performance evaluation and scalability analysis across multiple scientific workloads
- Designed and prototyped a new parallel implementation of Kirchhoff pre-stack depth migration, targeting improved efficiency and scalability on supercomputers

## LECTURER

**December 2022 - Present** **Université de Versailles Saint Quentin - High Performance Master, Guyancourt, France**

- Classes: Numerical calculus
- Topics: Iterative methods to solve linear systems, band matrices, Lapack, BLAS
- 12 hours of practical courses.

**January 2022 - May 2025** **Institut des Sciences et Techniques des Yvelines, Vélizy, France**

- Classes: Advanced Operating Systems
- Topics: Elf format, process scheduling, threading, virtual memory
- Student project: implementing a Shell in C
- Lecture on "Changes in application development brought by Cloud Computing"
- 2 hours lecture + 40 hours of practical courses.

**March 2018 - June 2018** **Institut des Sciences et Techniques des Yvelines, Vélizy, France**

- Classes: Operating Systems
- Topics: Unix basic commands, bash scripts, C processes
- 60 hours of practical courses.

## EDUCATION

**March 2018 - March 2021** **Ph.D. in Computer Science**

Thesis title: Distributed and Parallel Programming Paradigms Using Graphs of Tasks for Post-Petascale Supercomputers  
Advisor: Pr. Serge G. Petiton  
University of Lille (Lille, France)

**September 2014 - September 2017** **Polytech'Lille, Computer Science and Statistics**

Master in Computer Science and Applied Mathematics  
University of Lille (Lille, France)  
Degree project: Distributed and parallel multi-level programming on supercomputers using YML+XMP, a development and execution environment for parallel and distributed applications based on a graph of PGAS components managed at runtime

**September 2016 - February 2017** **School of Computer Engineering**

Master in Informatics Engineering - Erasmus  
Technical University of Madrid (Boadilla del Monte, Spain)

**September 2012 - July 2014** **Undergraduate Studies**

Preparatory Classes in Mathematics and Physics: Specific training for competitive entry examinations for the French "Grandes Ecoles"  
Lycée Clémentineau (Reims, France)

September 2009 - July 2012 **Scientific Baccalauréat**

Major: Mathematics  
Lycée Samuel Beckett (La Ferté-sous-Jouarre, France)

## TECHNICAL SKILLS

**Programming Languages:** C#, Python, Rust, Bash

**Cloud Platforms:** AWS (S3, EC2, SQS, AmazonMQ, DynamoDB, ElastiCache, EKS, CloudWatch, Bedrock, ParallelCluster/PCS), GCP (GCS, PubSub, GKE, Cloud Logging)

**Infrastructure & DevOps:** Kubernetes, Terraform, Docker, GitHub Actions

**Distributed Systems & APIs:** gRPC, Protocol Buffers, multi-language SDK generation (C#, Java)

**SaaS Platform Design:** Identity & access management, usage tracking & metering, platform governance, BYOC architecture, usage-based pricing models

**AI & Agentic Orchestration:** AWS Bedrock (Mistral), MCP, A2A, multi-agent orchestration, agentic workflow architecture, human-in-the-loop design, AI agent integration with SaaS/HPC backends, conversational simulation workflows

**Observability & Reliability:** Grafana, Prometheus, Langfuse, SigNoz, autoscaling (HPA signals), fault-tolerant design, checkpointing, logging/monitoring best practices

**Security & Compliance:** TLS, IAM, least-privilege, image scanning & CVE hygiene

**Product & Strategy:** Multi-year roadmap ownership, budget & FTE resource allocation, pre-sales engineering, training program design, customer support operations (L3)

## ACADEMIC RESEARCH

### Technical reports

Q. Delamea, **J. Gurhem**, D. Brasseur, D. Diakite, E. Soubiran, F. Dewas, F. Lemaitre, J. Fonseca, M. Elflitty, L. Ziane Khodja, A. Louvion, Y. Kharouni, D. Dubuc, N. Gruel, W. Kirschenmann, "ArmoniK, A Fault-Tolerant General-Purpose Distributed Computing Framework For Massive Task-Based Workloads", Technical report, August, 2023. [link](#)

### Conference Articles

Q. Delamea, J. Burman, **J. Gurhem**, M. Khairallah, W. Kirschenmann, F. Lemaitre and S. Vialle, "Cloud-Agnostic Serverless Platform for Fault-Tolerant Execution of Dynamic Task Graphs", 2025 IEEE Cloud Summit, Washington, DC, USA, June 26-27, 2025. [link](#)

M. Vandromme, N. Hochart, S. G. Petiton, **J. Gurhem**, M. Tsuji and M. Sato, "Sequences of Distributed Matrix-Vector Product for Very Large and Very Sparse Irregular Matrices", SC24-W: Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis, Atlanta, GA, USA, 2024. [link](#)

M. Vandromme, **J. Gurhem**, M. Tsuji, S. G. Petiton, and M. Sato, "Scaling the PageRank Algorithm for Very Large Graphs on the Fugaku Supercomputer", Computational Science - ICCS 2022 - 22nd International Conference, London, UK, June 21-23, 2022. [link](#)

**J. Gurhem**, H. Calandra and S. G. Petiton, "Parallel and Distributed Task-Based Kirchhoff Seismic Pre-Stack Depth Migration Application", in 20th International Symposium on Parallel and Distributed Computing (ISPDC 2021), Cluj-Napoca, Romania. [link](#)

**J. Gurhem**, M. Vandromme, M. Tsuji, S. G. Petiton, and M. Sato, "Sequences of Sparse Matrix-Vector Multiplication on Fugaku's A64FX processors", EAHPC-2021 - Embracing Arm for High Performance Computing Workshop. [link](#)

**J. Gurhem** and S. G. Petiton, "A current task-based programming paradigms analysis", in Computational Science - ICCS 2020, V. V. Krzhizhanovskaya, G. Závodszy, M. H. Lees, J. J. Dongarra, P. M. A. Sloot, S. Brissos, and J. Teixeira, Eds., Cham: Springer International Publishing, 2020, pp. 203-216, ISBN: 978-3-030-50426-7. [link](#)

**J. Gurhem**, M. Tsuji, S. G. Petiton, and M. Sato, "Distributed and Parallel Programming Paradigms on the K computer and a Cluster", in Proceedings of the International Conference on High Performance Computing in Asia-Pacific Region, ser. HPC Asia 2019, Guangzhou, China: ACM, 2019, pp. 9–17, ISBN: 978-1-4503-6632-8. [link](#)

## Extended Abstracts

**J. Gurhem**, H. Calandra and S. G. Petiton, "Task-Based Sparse Matrix Vector Product and Kirchhoff Migration", 2021 Oil & Gas HPC Conference

## Talks

**J. Gurhem**, W.Kirschenmann, "ArmoniK : une solution open source pour l'orchestration et la distribution de calculs", COMPAS 2025 Conférence francophone d'informatique en Parallélisme, Architecture et Système, June 24-27, 2025, Bordeaux, France.

**J. Gurhem**, "Resilient Orchestration of Compute-Intensive Workflows: The ArmoniK Framework for HPC and ML", DP2E-AI 2025 - 1st International Workshop on Distributed and Parallel Programming for Extreme-scale AI, June 16-17, 2025, Paris, France.

**J. Gurhem**, "ArmoniK: un orchestrateur de calculs open source", Workshop CSMA Juniors 2025, May 14-16, 2025, Toulouse, France.

S. G. Petiton, **J. Gurhem**, H. Calandra, "A Taxonomy of Distributed and Parallel Languages for High Performance Tasks-Based Multilevel Computing", SIAM Conference on Parallel Processing for Scientific Computing, February 2020, Seattle, U.S.

**J. Gurhem**, "Distributed and Parallel Programming using Graphs of Tasks with a Scheduler Optimizing Data Migrations", PhD Day 2019 at Maison de la Simulation, June 2019, Saclay, France.

**J. Gurhem**, "Programming Paradigm and Language Candidates for Extreme Scale Distributed and Parallel Computing", MATHIAS 2018 Computational Science Engineering & Data Science by TOTAL, Oct 2018, Serris, Paris, France.

**J. Gurhem**, "Distributed and Parallel Dense Block Linear Algebra using YML and XMP", First French-Japanese-German on Programming and Computing for Exascale and Beyond, Embassy of France in Tokyo, April 2017, Tokyo, Japan.

## LANGUAGES

**French:** Native language

**English:** Fluent (C1)

## EXTRA-CURRICULAR

Active technology watch across software engineering, distributed systems, and AI — with a current focus on agentic AI frameworks and AI-assisted development tools, exploring their application to accelerate product development and expand platform capabilities

Speaker at technical meetups and conferences (upcoming: Rust Paris, June 2026), in addition to academic venues (COMPAS, IEEE Cloud Summit, CSMA)

Reading (Fantasy, Science-Fiction)