2025 *Activity Report*

RESEARCH CENTRE: **Inria Paris Centre**

Project-Team

MIMOVE

Middleware on the Move

Project-Team MIMOVE

Creation of the Project-Team: 2018 February 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

A1.2.1. – Dynamic reconfiguration

A1.2.3. – Routing

A1.2.4. – QoS, performance evaluation

A1.2.6. – Sensor networks

A1.3.2. – Mobile distributed systems

A1.3.5. – Cloud

A1.3.6. – Fog, Edge

A1.5. – Complex systems

A1.5.1. – Systems of systems

A1.5.2. – Communicating systems

A2.5. – Software engineering

A2.6.2. – Middleware

A3.2.4. – Semantic Web

A3.2.5. – Ontologies

A9.2. – Machine learning

A9.4. – Natural language processing

Other research topics and application domains

B6.4. – Internet of things

B8.1. – Smart building/home

B8.2. – Connected city

Contents

1 Team members, visitors, external collaborators

Research Scientist

- Nikolaos Georgantas [Team leader, INRIA, Researcher, HDR]

PhD Students

- Emile Royer [INRIA]
- Haidong Zhao [INRIA]

Technical Staff

- Shahin Abdoul Soukour [INRIA, Engineer, until Nov 2025, Pre-Doc until Sep 2025 / Post-Doc after]
- Siamak Solat [INRIA, Engineer, from Feb 2025, Post-Doc]

Interns and Apprentices

- Akash Balamurugan [INRIA, Intern, from Apr 2025 until Sep 2025]
- Luigi Lizzini [INRIA, Intern, from Apr 2025 until Sep 2025]

Administrative Assistants

- Diana Marino Duarte [INRIA]
- Eugenie-Marie Montagne [INRIA]

External Collaborators

- Shahin Abdoul Soukour [Télécom SudParis, from Dec 2025, Post-Doc Researcher]
- William Aboucaya [Université Paris Dauphine - PSL, Post-Doc researcher]
- Maroua Bahri [Sorbonne Université, Associate Professor]
- Patient Ntumba Wa Ntumba [CNAM Paris, Post-Doc Researcher]

2 Overall objectives

MiMove has historically been positioned as a research team addressing distributed computing systems. Such systems comprise components that span global networking and computing infrastructures, mobile networking environments, powerful hand-held user devices, and physical-world sensing and actuation devices. In particular, the Internet of Things (IoT) has been one of our main focuses. In such rich environments, distributed systems have a number of challenging features, such as dynamicity due to volatile resources and user mobility, heterogeneity due to constituent resources developed and run independently, and context-dependence due to the highly changing characteristics of the execution environment, whether technical, physical or social.

In this context, we have addressed in the past several phases in the lifecycle of distributed systems: system design, deployment and runtime. We have tackled aspects such as: system interoperability & composition, resource allocation & system performance, reliable mobile crowdsensing for environmental monitoring, collaborative participatory processes. In our solutions, we have introduced system models, analyses, algorithms and protocols for capturing and managing the characteristics of the systems under study, as well as designed and developed related middleware tools and architectures.

More recently, while keeping the essential character of the team, we have been gradually aligning our activities to a new orientation. We are focusing our distributed system research on distributed machine learning (ML) systems, including federated ML systems, addressing both the training of ML models and their use for ML inference. We situate distributed ML systems of interest in the resource/compute continuum edge-fog-cloud, combined with the IoT. In this setting, ML systems have to deal with the specificities related to the resource environment, but also to handle continuous IoT data streams. The latter requirement makes us concentrate on online machine learning.

In this context, our central objective is to optimize the essential system trade-off between performance and resource usage, where we take into account the performance of both ML models (accuracy) and ML systems (quality of service). To this aim, we tackle algorithmic aspects of both ML models (e.g., parameter tuning) and ML systems (e.g., scheduling, communication). On the other hand, acknowledging the powerful effect that ML-based methods can have on distributed systems, our second research objective is to employ ML models for supporting important system tasks such as resource allocation and scheduling in the compute continuum.

3 Research program

As part of MiMove's new research orientation introduced in the previous section, we have been developing the following research activities.

3.1 Efficient scheduling of ML inference on GPUs

We address optimal scheduling of ML inference tasks on local GPU devices. We take into account the priorities of different tasks (urgency levels) as well as the interference in resource usage that results from task processing parallelization. Scheduling is decided based on dynamic interference prediction.

3.2 Automated machine learning (autoML) on data streams

Automated machine learning aims at automating the optimal tuning of ML algorithms. We address the problem of autoML when applied to IoT data streams in a distributed setting. Distribution is applied as a means for managing both the computational complexity of autoML tasks and the specificities of IoT data collection. In the context of distributed machine learning, we are studying the effect of online federated learning on decision trees.

3.3 Robust ML under resource volatility

We address the problem of scheduling ML training and inference processes over a large-scale network of heterogeneous, volatile and constrained resources.

3.4 NLP in goal-oriented requirements engineering

In Goal-Oriented Requirements Engineering (GORE), application designers capture requirements of the new system-to-be as a goal hierarchy. We leverage domain knowledge graphs (KGs) as sources for inspiring and refining goals. We develop methods and related tools that provide interactive assistance to designers for goal elicitation. To extract relevant knowledge from a KG in an automated way, we use NLP techniques in several innovative ways.

3.5 Reinforcement learning for proactive scheduling of data streams

We address optimal placement of data stream operators in the compute continuum. Extending our previous work on heuristic-based optimal scheduling on fog-cloud resources that minimizes overall resource usage cost, we explore online reinforcement learning (RL) for proactive scheduling, where RL is performed on dynamic simulation data.

3.6 Federated learning on IoT data

We explore federated learning (FL) on IoT data, where we deal in particular with low-volume, sparse, imbalanced, non-IID real data from device Wi-Fi / Bluetooth connectivity logs. We apply our FL method to predict long-term occupancy in industrial buildings.

4 Application domains

Historically, MiMove's research has had a strong focus on the Internet of Things (IoT). This is still currently the case, in particular coupled with the resource/compute continuum. Numerous application domains result from this setting. Connected cities and smart buildings are among those.

5 Latest software developments, platforms, open data

5.1 Latest software developments

5.1.1 KG2GoalModel

Name: Goal model construction from Knowledge Graph

Keywords: Goal-oriented requirements engineering, Knowledge graph, Natural language processing

Functional Description: In Goal-Oriented Requirements Engineering (GORE), application designers capture requirements of the new system-to-be as a goal hierarchy. We leverage domain knowledge graphs (KGs) as sources for inspiring and refining goals. We propose a method and a graphical tool that provide interactive assistance to designers for goal elicitation. To extract relevant knowledge from a KG in an automated way, we use NLP techniques (Semantic Similarity, Natural Language Inference, Sentiment Analysis, Sentence Compression, LLMs) in several innovative ways.

URL: https://github.com/ShahinAbdoulSoukour/KG_for_goal_modeling

Contact: Abdoul Abdoul Soukour

Participants: Abdoul Abdoul Soukour, William Aboucaya, Nikolaos Georgantas

5.1.2 ML inference serving system

Name: ML inference serving system (short name will be disclosed after publishing of this work)

Keywords: Machine learning, Inference serving system, Scheduling

Functional Description: Machine learning (ML) inference serving systems host deep neural network (DNN) models and aim to efficiently schedule incoming inference requests across available GPU resources. However, limited support for task prioritization and insufficient latency estimation under concurrency may restrict their applicability in many real-world scenarios. We present a serving system designed to enhance deadline satisfaction for mixed-priority inference traffic under high GPU utilization. To improve latency estimation, our serving system models potential contention during data transfer and accounts for contention in kernel execution through an adaptive prediction model. By drawing on these predictions, it performs priority-aware scheduling and thereby provides differentiated handling.

URL: https://-will-be-disclosed-after-publishing-of-this-work

Contact: Haidong Zhao

Participants: Haidong Zhao, Nikolaos Georgantas

5.1.3 OSMAC

Name: Online SMAC

Keywords: Machine learning, Automated machine learning, Online Learning, Data stream, Bayesian optimization

Functional Description: Online SMAC is an autoML optimiser for online machine learning that determines the best machine learning model for a given task, and what hyperparameters to use. It uses Bayesian optimisation to find the best model–hyperparameter combination, inspired from the SMAC method.

Release Contributions: First release.

News of the Year: Created the first version.

URL: https://inria.hal.science/hal-05290014

Publication: hal-05290014

Contact: Emile Royer

Participants: Emile Royer, Maroua Bahri, Nikolaos Georgantas

5.1.4 OccupFL

Name: Federated Learning for Occupancy Prediction

Keywords: Federated learning, Occupancy prediction, Smart building

Functional Description: Connected-device logs are a common proxy for estimating human occupancy in smart buildings, but they pose several challenges for machine-learning practitioners. As part of the BPI France 2030 CP4SC research and innovation project, data were collected over 511 days at 15-minute intervals across eight zones. Accurate forecasting of device-connectivity counts often requires training data that are long enough to capture seasonal patterns. These data were ill-suited to federated learning, especially for long-term prediction, due to (i) low volume, (ii) extreme sparsity and class imbalance, and (iii) demonstrably non-IID distributions. We address these challenges with a proof-of-concept federated learning pipeline that includes: (1) statistical tests confirming non-IID distributions, (2) a synthetic data generator that preserves seasonal patterns while filling gaps, (3) a dynamic FedProx-style server for stable aggregation, and (4) focal-MSE loss functions calibrated to each zone's imbalance ratio.

URL: https://-will-be-disclosed-after-publishing-of-this-work

Contact: Siamak Solat

Participants: Siamak Solat, Nikolaos Georgantas

5.1.5 AI-based adaptive scheduler

Name: AI-based adaptive scheduler

Keywords: Internet of things, Data stream, Fog computing, Cloud computing, Machine learning, Reinforcement learning

Functional Description: AI-based proactive scheduler optimizes DSPA (Data Stream Processing and Analytics) operator placement across the hierarchical Edge-Fog-Cloud architecture. Unlike traditional reactive heuristics, our scheduler leverages Reinforcement Learning (RL), powered by the Proximal Policy Optimization (PPO) algorithm. We implement a state persistence mechanism that enables cumulative online learning across training episodes, which distinguishes our method from conventional supervised learning approaches that rely on pre-collected datasets. Unlike offline training paradigms, our PPO agent learns directly through continuous interaction with the live YAFS simulation environment, adapting its scheduling policies in real-time based on immediate system feedback. This online learning approach allows the agent to make proactive placement decisions by observing comprehensive state representations encompassing service placements, device resource utilization (CPU, RAM), network link conditions (bandwidth, latency), and application-level metrics.

URL: https://gitlab.inria.fr/abalamur/ai-based-adaptive-scheduler

Contact: Akash Balamurugan

Participants: Akash Balamurugan, Maroua Bahri, Nikolaos Georgantas, Patient Ntumba Wa Ntumba

5.1.6 Interoperability Enabler

Name: Interoperability Enabler for Data Spaces

Keywords: Data spaces, Data marketplace, Interoperability

Functional Description: The Horizon Europe SEDIMARK project designed and prototyped a secure decentralised and intelligent data and services marketplace that bridges remote data platforms and allows the efficient and privacy-preserving sharing of vast amounts of heterogeneous, high quality, certified data and services supporting the common EU data spaces. Interoperability Enabler was designed to facilitate seamless integration and interaction among various artefacts within the SEDIMARK ecosystem, including data, AI models, and service offerings. Interoperability Enabler comprises the following components: (i) Data Formatter – Convert JSON data (time-series data) into the SEDIMARK internal processing format (pandas DataFrames), (ii) Data Mapper – Convert data from pandas DataFrames into JSON, (iii) Data Extractor – Extract relevant data from a pandas DataFrame, (iv) Metadata Restorer – Restore metadata to a pandas DataFrame, (v) Data Merger – Merge two DataFrames by matching column names.

URL: https://github.com/Sedimark/InteroperabilityEnabler

Contact: Abdoul Abdoul Soukour

Participants: Abdoul Abdoul Soukour, Maroua Bahri, Nikolaos Georgantas

6 New results

6.1 Leveraging domain knowledge in software system goal models

Participants: Shahin Abdoul Soukour, Nikolaos Georgantas.

In Software Engineering (SE), system design is an important phase in the software development lifecycle. During this phase, the architecture and internal components of a software system are defined to meet specific requirements. Effective system design ensures that the final product is robust, scalable and aligned with stakeholder needs. One of the major challenges in system design is effectively capturing and structuring domain knowledge to guide the design process. Requirements Engineering (RE) plays a pivotal role in addressing this challenge.RE is the initial and fundamental step in the design process of any information or software system, focusing on establishing, documenting, analyzing, validating, and managing the requirements of a software system. It takes into account all the activities related to eliciting, specifying, and validating needs and constraints of the stakeholders, ensuring that the final product closely corresponds to their expectations. Traditionally, domain knowledge in RE serves as contextual support, helping to clarify and refine requirements that stem mainly from stakeholder input, regulations or business needs. However, in this research, we explore an alternative approach where domain knowledge is not just an auxiliary resource but an active source of inspiration for generating requirements. Goal-Oriented Requirements Engineering (GORE) is a specific approach within RE that emphasizes identifying and modelling the high-level goals of stakeholders. It focuses on understanding why certain requirements are needed and how they contribute to the overall goals of the system. The goal model, an essential component of GORE, describes the system's goals using a hierarchical structure in which high-level goals are refined (or decomposed) into more specific ones. Despite various automation or semi-automation attempts, building goal models for software system design remains time-consuming and quite complex, often requiring significant manual effort. To respond to these challenges, this PhD research focuses on leveraging domain knowledge in the form of a Knowledge Graph (KG). The KG will assist application designers in creating goals that are inspired from this knowledge, thereby facilitating the construction of goal models. By combining with the integration of Natural Language Processing (NLP) techniques, relevant information from the KG can be captured and suggested, aiding the application designer in building a goal model more efficiently for software system design. This thesis makes several key contributions: the design and implementation of methods for exploring KG by using NLP techniques for semi-automatic goal modelling; the development of a technique to make the formulated goals more abstract to facilitate KG exploration and to extract relevant information; and the development of a prototype, which is an interactive graphical tool, that demonstrates and validates the proposed approach.

6.2 An Interactive Tool for Goal Model Construction using a Knowledge Graph

Participants: Shahin Abdoul Soukour, William Aboucaya, Nikolaos Georgantas.

The goal model is an essential model in Goal-Oriented Requirements Engineering. It is used to describe the system's goals using a hierarchical structure in which high-level goals are refined into more specific ones. Constructing a goal model for a new application can present challenges, demanding considerable time and effort. Although there have been attempts to automate or semi-automate the construction of goal models, these tasks remain complex and manual. This paper presents an interactive graphical tool that leverages a domain Knowledge Graph (KG) to assist the application designer in creating goals derived from this knowledge, thereby facilitating the creation of goal models. We use semantic similarity measurement and Natural Language Inference (NLI) to effectively extract and align triples from the KG with the high-level initial goals formulated by the application designer. The extracted triples undergo sentiment analysis and Graph-to-Text (G2T) generation to build meaningful subgoals. Nevertheless, processing KGs with Natural Language Processing (NLP) techniques can be a lengthy process. We introduce a restriction based approach to bound the exploration of the KG to the most promising nodes. By tuning KG exploration bounds while using our tool in a case study, we analyze the trade-off between the quality of the resulting goal model and time performance, which is a key factor for an interactive approach. Our paper highlights the relevance of our restriction based approach to information retrieval in KGs to facilitate goal model generation.

6.3 OSMAC: A Dynamic SMAC for Data Streams

Participants: Émile Royer, Maroua Bahri, Nikolaos Georgantas.

Automated machine learning (autoML) methods often require multiple passes over data and are computationally intensive, rendering them unsuitable for streaming scenarios where data is continuously generated and distributions evolve over time. The few existing autoML solutions for stream learning mainly rely on random search or genetic algorithms, which struggle to maintain high performance in dynamic environments. By contrast, leading methods in batch learning such as the Sequential Model-based Algorithm Configuration (SMAC) leverage modelbased approaches, suggesting opportunities for improvement in stream settings. To address these challenges and meet the requirements of stream scenarios, we introduce OnlineSMAC, a model-based optimizer for data streams. OnlineSMAC combines Bayesian optimization with an extension of the SMAC optimizer to dynamically select optimal processing pipelines and hyperparameters. Our results show that this approach is highly competitive, achieving performance on par with state-of-the-art stream autoML methods. This highlights the promising potential of using Bayesian optimization for data streams.

6.4 ML Inference Scheduling with Predictable Latency

Participants: Haidong Zhao, Nikolaos Georgantas.

Machine learning (ML) inference serving systems can schedule requests to improve GPU utilization and to meet service level objectives (SLOs) or deadlines. However, improving GPU utilization may compromise latency-sensitive scheduling, as concurrent tasks contend for GPU resources and thereby introduce interference. Given that interference effects introduce unpredictability in scheduling, neglecting them may compromise SLO or deadline satisfaction. Nevertheless, existing interference prediction approaches remain limited in several respects, which may restrict their usefulness for scheduling. First, they are often coarse-grained, which ignores runtime co-location dynamics and thus restricts their accuracy in interference prediction. Second, they tend to use a static prediction model, which may not effectively cope with different workload characteristics. In this paper, we evaluate the potential limitations of existing interference prediction approaches, finding that coarse-grained methods can lead to noticeable deviations in prediction accuracy and that static models degrade considerably under changing workloads.

7 Partnerships and cooperations

7.1 European initiatives

7.1.1 Horizon Europe

SEDIMARK

Title: SEcure Decentralised Intelligent Data MARKetplace

Duration: 2022 - 2025

Partner Institutions:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- WINGS ICT SOLUTIONS TECHNOLOGIES PLIROFORIKIS KAI EPIKOINONION ANONYMI ETAIREIA (WINGS ICT SOLUTIONS AE), Greece
- UNIVERSITY COLLEGE DUBLIN, NATIONAL UNIVERSITY OF IRELAND, DUBLIN (NUID UCD), Ireland
- FORUM VIRIUM HELSINKI OY (RADIO- JATELEVISIOTEKNIIKAN TUTKIMUS RTT), Finland
- SIEMENS SRL, Romania
- ATOS SPAIN SA, Spain
- AYUNTAMIENTO DE SANTANDER, Spain
- METLEN ENERGY & METALS AE (METLEN), Greece
- UNIVERSIDAD DE CANTABRIA (UC), Spain
- FONDAZIONE LINKS - LEADING INNOVATION & KNOWLEDGE FOR SOCIETY (FONDAZIONE LINKS), Italy
- ATOS IT SOLUTIONS AND SERVICES IBERIA SL (ATOS IT), Spain
- UNIVERSITY OF SURREY (SURREY), United Kingdom
- EGM (EGM SAS), France

Participants: Shahin Abdoul Soukour, Nikolaos Georgantas.

SEDIMARK aims at designing and prototyping a secure, decentralised and intelligent data and services marketplace, based on Distributed Ledger Technology and Artificial Intelligence, which bridges remote data platforms and allows the efficient and privacy-preserving sharing of vast amounts of heterogeneus, high quality, certified data and services supporting the common EU data spaces. SEDIMARK includes a distributed registry of resources (data/services) stored on edge systems, close to where they are generated and where the data are cleaned, labelled, validated and anonymised. Energy efficient AI techniques will be used for automated data quality management, labelling and classification of data as well as for providing (distributed) analytics and advanced services on top of the data. Semantic interoperability based on common ontologies and data models will allow the easy and efficient discovery, sharing and federation of heterogeneous data from multiple sources.

7.2 National initiatives

BPI France 2030 CP4SC project

Title: Cloud Platform For Smart City

Duration: 2023 - 2025

Partner Institutions:

- ATOS/Eviden
- Ericsson
- INRIA
- INRAE
- IFPEN
- Oslandia
- Vertical M2M

Participants: Siamak Solat, Nikolaos Georgantas.

The goal of the CP4SC platform is to assist governments in implementing ambitious policies towards achieving carbon neutrality by ingesting data from various sources across multiple verticals, such as mobility, energy management, and earth and environmental observation. By placing data analysis at the core of these activities, the CP4SC project provides a comprehensive, adaptable, and secure technological solution that meets the highest requirements of organizations involved in complex projects, with a particular focus on mobility, 5G connectivity, and secure exchanges.

Inria Challenge Cupseli project

Title: Collaborative Unified Platform for a Scalable and Efficient Learning Infrastructure

Duration: 2025 - 2029

Partner Institutions:

- Inria teams: ARGO, MIMOVE, COAST, MAGELLAN, STACK, WIDE, OCKHAM, COATI, NEO, TADAAM, TOPAL
- Hivenet

Participants: Nikolaos Georgantas, Siamak Solat.

Hivenet offers a highly original data storage architecture, in which data is stored in a distributed and secure manner on the spare storage resources of participants, based on a peer-to-peer structure. This structure ensures scalability, resilience and voluntary sharing of data between users. The aim of this challenge between Hivenet and Inria is to push the limits of distributed AI computing. Its goal is to demonstrate that even the most demanding AI and Big Data applications can run efficiently on heterogeneous, distributed, and volatile resources – while maintaining accuracy, ensuring privacy, and reducing environmental impact.

8 Dissemination

8.1 Promoting scientific activities

8.1.1 Scientific events: selection

Member of the conference program committees

- Nikolaos Georgantas, member of the TPC of the following international conferences: ACM SAC'25, IEEE SOSE'25, IEEE SMARTCOMP'25'26, IEEE WETICE'25, CoopIS'25, MODELSWARD'25'26, ENASE'26.
- Siamak Solat, member of the TPC of the following international conference: IEEE BCCA'25.

Reviewer

- Siamak Solat, reviewer for the following international conference: ACM SAC'26.

8.1.2 Invited talks

- Haidong Zhao, "Strait: Perceiving Priority and Interference in ML Inference Serving", Networks & Systems Workshop (Journées non-thématiques GDR RSD), IRIT – Site ENSEEIHT, Toulouse, March 21, 2025.

8.1.3 Scientific expertise

- Nikolaos Georgantas, member of the EDITE Doctoral School's 2025 selection committee "Communications, Networks and Systems" for PhD fellowships.
- Nikolaos Georgantas, member of the PhD monitoring committee of Himadri Chhaya-Shailesh (Sorbonne Université), Victor Laforet (Sorbonne Université).

8.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

8.2.1 Supervision

- PhD thesis defense: Shahin Abdoul Soukour, "Leveraging domain knowledge in software system goal models", Sorbonne Université, Sep 19, 2025, Nikolaos Georgantas.
- Master's degree thesis defense: Akash Balamurugan, "AI-based scheduler for IoT data streams analytics", Cnam Paris, Oct 16, 2025, Maroua Bahri, Nikolaos Georgantas.
- 1st year Master's internship: Luigi Lizzini, AI-based scheduling for ML inference tasks, Cnam Paris, Apr-Sep 2025, Nikolaos Georgantas.
- PhDs in progress:

  - Haidong Zhao (from March 2023): "Efficient ML Inference Scheduling", Sorbonne Université, Nikolaos Georgantas.
  - Emile Royer (from November 2024): "Distributed automated machine learning with application on IoT data", Sorbonne Université, Maroua Bahri, Nikolaos Georgantas.

8.2.2 Juries

- Nikolaos Georgantas, rapporteur for the habilitation thesis of Joyce El Haddad (Université Paris Dauphine - PSL), Dec 2025.

8.2.3 Educational and pedagogical outreach

- Emile Royer, participation in a round table on "Pursuing a research career", EPF Engineering School, Apr 2025.

8.3 Popularization

8.3.1 Participation in Live events

- Emile Royer, presentation of the interactive animation “The illustrator apprentice” aiming at explaining artificial intelligence to children and the general public, Fête de la science, Campus Pierre et Marie Curie, Sorbonne Université, Oct 2025.

9 Scientific production

9.1 Major publications

[1] R. Angarita, B. Lefèvre, S. Ahvar, E. Ahvar, N. Georgantas and V. Issarny. ‘Universal Social Network Bus: Towards the Federation of Heterogeneous Online Social Network Services’. In: *ACM Transactions on Internet Technology* (2019). DOI: 10.1145/3323333. URL: https://hal.inria.fr/hal-02072544.

[2] A. Bennaceur and V. Issarny. ‘Automated Synthesis of Mediators to Support Component Interoperability’. In: *IEEE Transactions on Software Engineering* (2015), p. 22. URL: https://hal.inria.fr/hal-01076176.

[3] B. Billet and V. Issarny. ‘Spinel: An Opportunistic Proxy for Connecting Sensors to the Internet of Things’. In: *ACM Transactions on Internet Technology* 17.2 (Mar. 2017), pp. 1–21. DOI: 10.1145/3041025. URL: https://hal.inria.fr/hal-01505879.

[4] G. Blair, A. Bennaceur, N. Georgantas, P. Grace, V. Issarny, V. Nundloll and M. Paolucci. ‘The Role of Ontologies in Emergent Middleware: Supporting Interoperability in Complex Distributed Systems’. In: *Big Ideas track of ACM/IFIP/USENIX 12th International Middleware Conference*. Lisbon, Portugal, 2011. URL: http://hal.inria.fr/inria-00629059/en.

[5] G. Bouloukakis, N. Georgantas, P. Ntumba and V. Issarny. ‘Automated synthesis of mediators for middleware-layer protocol interoperability in the IoT’. In: *Future Generation Computer Systems* 101 (Dec. 2019), pp. 1271–1294. DOI: 10.1016/j.future.2019.05.064. URL: https://hal.inria.fr/hal-02304074.

[6] Y. Du, F. Sailhan and V. Issarny. ‘Let Opportunistic Crowdsensors Work Together for Resource-efficient, Quality-aware Observations’. In: *PerCom 2020: IEEE International Conference on Pervasive Computing and Communications*. Austin / Virtual, United States, Mar. 2020. DOI: 10.1109/PerCom45495.2020.9127391. URL: https://hal.archives-ouvertes.fr/hal-02463610.

[7] S. Hachem, A. Pathak and V. Issarny. ‘Service-Oriented Middleware for Large-Scale Mobile Participatory Sensing’. In: *Pervasive and Mobile Computing* (2014). URL: http://hal.inria.fr/hal-00872407.

[8] P. Ntumba, N. Georgantas and V. Christophides. 'Adaptive Scheduling of Continuous Operators for IoT Edge analytics'. In: *Future Generation Computer Systems* (15th Apr. 2024). DOI: 10.1016/j.future.2024.04.029. URL: https://hal.science/hal-04558794.

9.2 Publications of the year

International peer-reviewed conferences

[9] S. Abdoul-Soukour, W. Aboucaya and N. Georgantas. 'An Interactive Tool for Goal Model Construction using a Knowledge Graph'. In: REFSQ 2025 - 31st International Working Conference on Requirement Engineering: Foundation for Software Quality. Barcelona, Spain, 7th Apr. 2025, p. 15. URL: https://inria.hal.science/hal-04907365.

[10] É. Royer, M. Bahri and N. Georgantas. 'OSMAC: A Dynamic SMAC for Data Streams'. In: *2025 IEEE 37th International Conference on Tools with Artificial Intelligence*. 37th International Conference on Tools with Artificial Intelligence (ICTAI 2025). Athens, Greece, 15th Dec. 2025, pp. 73–80. DOI: 10.1109/ICTAI66417.2025.00018. URL: https://inria.hal.science/hal-05383821.

[11] H. Zhao and N. Georgantas. 'ML Inference Scheduling with Predictable Latency'. In: MAIoT '25: Middleware for Autonomous AIoT Systems in the Computing Continuum in conjunction with the 26th ACM/IFIP International Middleware Conference (Middleware 2025). Nashville, United States: ACM, 15th Dec. 2025, pp. 25–30. DOI: 10.1145/3774901.3778066. URL: https://hal.science/hal-05431608.

Doctoral dissertations and habilitation theses

[12] S. Abdoul Soukour. 'Leveraging domain knowledge in software system goal models'. Sorbonne Université, 19th Sept. 2025. URL: https://theses.hal.science/tel-05448495.

Other scientific publications

[13] S. Abdoul-Soukour, W. Aboucaya and N. Georgantas. 'An Interactive Tool for Goal Model Construction using a Knowledge Graph'. In: Journées nationales du GDR Génie de la Programmation et du Logiciel - GDR GPL 2025. Pau, France, 19th June 2025. URL: https://inria.hal.science/hal-05207665.

[14] A. Balamurugan. 'AI-based scheduler for IoT data streams analytics'. Cnam Paris, 16th Oct. 2025. URL: https://inria.hal.science/hal-05485542.

Software

[15] [SW] É. Royer, M. Bahri and N. Georgantas, *Online SMAC* 30th Sept. 2025. Inria. LIC: GNU Affero General Public License v3.0 only. HAL: ⟨hal-05290014⟩, URL: https://inria.hal.science/hal-05290014, SWHID: ⟨swh:1:dir:1c98f9797a6ef6c6a9c9e48fc3c2556dea610276;origin=https://hal.archives-ouvertes.fr/hal-05290014;visit=swh:1:snp:a31222249926baf9ed59d62b1bb446f9274becdf;anchor=swh:1:rel:72c3b01c8aaefbe7791f6a726a07fc03e00ce56c;path=/⟩.