

2025 Activity Report

RESEARCH CENTRE: Inria Centre at Rennes University


IN PARTNERSHIP WITH: Université de Rennes


Project-Team

WIDE

the World Is Distributed Exploring the tension
between scale and coordination

In collaboration with Institut de recherche en informatique et systèmes aléatoires
(IRISA)



Project-Team WIDE

Creation of the Project-Team: 2018 June 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

- A1.3.2. – Mobile distributed systems
- A1.3.3. – Blockchain
- A1.3.4. – Peer to peer
- A1.3.5. – Cloud
- A1.3.6. – Fog, Edge
- A2.1.7. – Distributed programming
- A2.6.1. – Operating systems
- A2.6.2. – Middleware
- A2.6.3. – Virtual machines
- A3.4. – Machine learning and statistics
- A4. – Security and privacy
- A4.8. – Privacy-enhancing technologies
- A7.1.1. – Distributed algorithms
- A7.1.2. – Parallel algorithms
- A7.1.3. – Graph algorithms
- A9. – Artificial intelligence
- A9.2. – Machine learning
- A9.9. – Distributed AI, Multi-agent

Other research topics and application domains

- B6.3.1. – Web
- B6.3.5. – Search engines
- B6.4. – Internet of things
- B9.5.1. – Computer science
- B9.5.6. – Data science

Contents

Project-Team WIDE	1
1 Team members, visitors, external collaborators	5
2 Overall objectives	6
2.1 Overview	6
2.2 Planetary-Scale Geo-Distributed Systems	6
2.3 Highly Personalized On-Line Services	7
2.4 Social Collaboration Platforms	7
3 Research program	8
3.1 Overview	8
3.2 Hybrid Scalable Architectures	9
3.3 Personalizable Privacy-Aware Distributed Systems	10
3.4 Network Diffusion Processes	11
3.5 Systemizing Modular Distributed Computability and Efficiency	12
3.6 Evolution of our research program (2022-2026)	14
4 Application domains	15
5 Social and environmental responsibility	15
6 Highlights of the year	15
7 Latest software developments, platforms, open data	16
7.1 Latest software developments	16
7.1.1 DecentralizedFlower	16
7.1.2 nodemanager	16
7.1.3 DecentralizedDeclearn	16
7.1.4 decentralised-data-wallet	16
7.1.5 CAC	17
7.1.6 QAAT	17
7.1.7 Splitschain	17
7.1.8 oversim-ipfs	17
7.1.9 PPFDIMs	18
8 New results	18
8.1 Distributed Algorithms and Blockchain	18
8.1.1 Contention-Aware Cooperation	18
8.1.2 Ethical Risk Analysis of L2 Rollups	18
8.1.3 Communication abstractions in systems prone to malicious attacks	19
8.1.4 Discreet: Distributed delivery service with context-aware cooperation	19
8.1.5 Luby's MIS algorithms made self-stabilizing	20
8.1.6 On the h-majority dynamics with many opinions	20
8.2 Large scale Cloud environments	21
8.2.1 Off-the shelf network traffic analysis	21
8.2.2 Towards efficient kernel network processing for VPNs	21
8.2.3 Efficient Load balancing for multi-tier applications	21
8.2.4 Containers as alternatives for Cloud gaming	21
8.2.5 Disconnecting Users from Virtual Worlds with a Single Packet: an Unreal Untold Story	22
8.3 Artificial Intelligence and Machine Learning	22
8.3.1 Strengthening malware analysis against obfuscation and packing	22
8.3.2 Low-Cost Privacy-Preserving Decentralized Learning	22
8.4 Load Balancing	23

8.4.1	An asymptotically optimal algorithm for generating bin cardinalities	23
9	Bilateral contracts and grants with industry	23
9.1	Bilateral contracts with industry	23
9.1.1	CIFRE with Broadpeak	23
9.1.2	CIFRE with Blacknut: Efficient Containerized Cloud-Gaming Platforms	23
9.1.3	Flexible Virtualization for Processing In Memory YUMPIM (ANR - PRCE)	24
10	Partnerships and cooperations	24
10.1	International initiatives	24
10.1.1	Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program	24
10.1.2	Inria associate team not involved in an IIL or an international program	25
10.2	National initiatives	25
11	Dissemination	28
11.1	Promoting scientific activities	28
11.1.1	Scientific events: organisation	28
11.1.2	Scientific events: selection	28
11.1.3	Journal	29
11.1.4	Invited talks	29
11.1.5	Leadership within the scientific community	30
11.1.6	Scientific expertise	30
11.1.7	Research administration	30
11.2	Teaching - Supervision - Juries - Educational and pedagogical outreach	30
11.2.1	Teaching	31
11.2.2	Supervision	31
11.2.3	Juries	32
11.3	Popularization	33
11.3.1	Productions (articles, videos, podcasts, serious games, ...)	33
12	Scientific production	33
12.1	Major publications	33
12.2	Publications of the year	34
12.3	Cited publications	37

1 Team members, visitors, external collaborators

Research Scientists

- Davide Frey [INRIA, Researcher, HDR]
- Georgios Giakkoupis [INRIA, Researcher]

Faculty Members

- François Taiani [Team leader, Université de Rennes, Professor, HDR]
- Yerom David Bromberg [Université de Rennes, Professor, HDR]
- Brice Ekane Apah [Université de Rennes, Associate Professor]
- Achour Mostefaoui [Université de Rennes, Professor, HDR]
- Barbe Mvondo Djob [Université de Rennes, Associate Professor]
- Michel Raynal [Université de Rennes, Emeritus, HDR]

Post-Doctoral Fellows

- Georgy Ishmaev [Université de Rennes, Post-Doctoral Fellow, until Oct 2025]
- Dimitrios Los [INRIA, Post-Doctoral Fellow, until Feb 2025]

PhD Students

- Timothé Albouy [Université de Rennes, until Jan 2025]
- Hugo Bertin [Université de Rennes, from Feb 2025]
- Fonyuy-Asheri Caleb [INRIA]
- Opale Duvivier [Université de Rennes, CIFRE]
- Adrien Gegout [Université de Rennes, CIFRE]
- Amelie Gonzalez [Université de Rennes]
- Junrui Hua [HIVE COMPUTING SERVICES SAS, CIFRE]
- Dimitri Lereverend [INRIA]
- Honore Cesaire Mounah [INRIA puis Université de Rennes]
- Victoire Nganfang [Université de Rennes]
- Elie Raspaud [INRIA, from Sep 2025]
- Manon Sourisseau [Université de Rennes]
- Stella Tchoutcha [Université de Rennes, from Dec 2025]

Technical Staff

- Olivier Deloubriere [INRIA, Engineer]
- Patricio Inzaghi [INRIA, Engineer]
- Cyrille Kenfack [INRIA, Engineer]
- Elie Raspaud [INRIA, Engineer, until Mar 2025]
- Harvey Williams [INRIA, Engineer, from May 2025]
- Harvey Williams [INRIA, Engineer, until Mar 2025]

Interns and Apprentices

- Iadine Brenda Atangana Wamsa [Université de Rennes, Intern, from Aug 2025 until Sep 2025]
- Zakaria El Maachi [Université de Rennes, Intern, from Jun 2025 until Sep 2025]
- Julien Houget [Université de Rennes, Intern, from Oct 2025]
- Julien Houget [Université de Rennes, Intern, from Jun 2025 until Sep 2025]
- Eugenio Mazzina [INRIA, Intern, until Feb 2025]

Administrative Assistant

- Virginie Desroches [INRIA]

2 Overall objectives

2.1 Overview

The long term goal of the WIDE team is to provide the practical tools and theoretical foundations required to address the scale, dynamicity, and uncertainty that constitute the foundations of modern distributed computer systems. In particular, we would like to **explore the inherent tension between scalability and coordination guarantees**, and develop novel techniques and paradigms that are adapted to the rapid and profound changes impacting today’s distributed systems, both in terms of the application domains they support and the operational constraints they must meet.

These changes are particularly visible in three key areas related to our research: *(i)* planetary-scale information systems, *(ii)* personalized services, and *(iii)* new forms of social applications (e.g. in the field of the sharing economy).

2.2 Planetary-Scale Geo-Distributed Systems

Modern large-scale systems often encompass thousands of server nodes, hosted in tens of datacenters distributed over several continents. To address the challenges posed by such systems, alternative distributed architectures are today emerging that emphasize *decentralized* and *loosely coupled* interactions. This evolution can be observed at multiple levels of an application’s distributed stack: the growing interest, both practical and theoretical, for weak consistency models is such an example. In spite of their potential counter-intuitive behaviors, weakly consistent data-structures allow developers to trade strict coordination guarantees for the ability to deliver a reactive and scalable service even when hit by arbitrary network delays or system partitions. At a higher, more architectural level, similar motivations explain the push for *micro-services* on the server side of on-line applications and the growth of rich *browser-based programming technologies* on their client side. Micro services help development teams decompose complex applications into a set of simpler and loosely-connected distributed services. In a parallel evolution, modern browsers embark increasingly powerful networking APIs such as WebRTC. These APIs are prompting a fresh

rethink of the typical distribution of capabilities between servers and clients. This is likely to lead to more services and computations being offloaded to browsers, in particular within hybrid architectures. The above evolutions, away from tightly synchronized and monolithic deployments towards heterogeneous, composite and loosely coordinated distributed systems, raise a number of difficult challenges at the crossroad of theoretical distributed algorithms, system architecture, and programming frameworks. One of these challenges pertains to the growing complexity arising from these systems: as richer and more diverse services are being composed to construct whole applications, individual developers can only hope to grasp parts of the resulting systems. Similarly, weak consistency models and loose coordination mechanisms tend to lead to counter-intuitive behaviors, while only providing weak overall guarantees. This lack of systematic guarantees and understandability make it harder for practitioners to design, deploy, and validate the distributed systems they produce, leading to rising costs and high entry barriers.

In order to address these challenges, we argue that modern-day distributed systems require new principled algorithms, approaches, and architectural patterns able to provide sound foundations to their development while guaranteeing robust service guarantees, thus lowering the cost of their development and maintenance, increasing their reliability, and rendering them technically approachable to a wider audience.

2.3 Highly Personalized On-Line Services

Ever increasing volumes of data are being produced and made available from a growing number of sources (Internet of Things sensors, open data repositories, user-generated content services).

As a result, digital users find it increasingly difficult to face the data deluge they are subjected to without additional help. This difficulty has fueled the rise of notification solutions over traditional search, in order to push few but relevant information items to users rather than leave them to sieve through a large mass of non-curated data. To provide such personalized services, most companies rely today on centralized or tightly coupled systems hosted in data centers or in the cloud. These systems use advanced data-mining and machine learning techniques to deliver enhanced, personalized, services to users and companies, and often exploit highly parallelized data analytics frameworks such as Spark, and Flink.

Selecting the best information for a user in order to provide a personalized experience requires however to gather enough information about this user, which raises a number of important technical challenges and privacy protection issues. More precisely, this concentration poses strong risks to the privacy of users, and limits the scope of personalization to tightly integrated datasets.

The use of large monolithic infrastructures also limits the use of machine learning and personalization to situations in which data is fully available to the organization managing the underlying computing infrastructure. This set-up prevents for instance cases in which sensitive data may not be shared freely, but might be of mutual interest to several independent participants in order to construct common machine learning models usable by all. Such situations occur for instance in the context of the mining of health-records by independent health-organizations, or in the collective harnessing of individual on-line profiles for personalization purpose by private users.

Alternative decentralized approaches that eschew the need for a central all-encompassing authority holds the promise of delivering knowledge while protecting individual participants. Constructing such systems requires however to address the inherent tension between the need to limit sensitive individual leaks, while maximizing collectively gained insights. Answering this tension calls on techniques and approaches from distributed systems, information theory, security, and randomized processes, making it a rich and dense research area, with a high impact potential. The problem of distributed privacy in a digital interconnected age further touches on interdisciplinary questions of Law, Sociology and Public Policy, which we think can only be explored in collaboration with colleagues from these fields.

2.4 Social Collaboration Platforms

On-line social networks have had a fundamental and lasting impact on the Internet. In recent years, numerous applications have appeared that go beyond the services originally provided by “pure” on-line social networks, such as posting messages or maintaining on-line “friendship” links. These new applications seek to organize and coordinate users, often in the context of the sharing economy, for instance in order to facilitate car-sharing (e.g. BlaBla car, www.blablacar.com), short-term renting (e.g. AirBnB, www.airbnb.com), and peer-to-peer financial services (e.g. Lending Club, www.lendingclub.com). Some systems, such as Bitcoin or

Ethereum, have given rise to new distributed protocols combining elements of cryptography and distribution that are now largely discussed in the research community, and have attracted the attention of policy makers and leading financial actors.

The challenges faced by such social applications blend in many ways issues already discussed in the two previous subsections and cast them in an application-driven context. These social collaboration platforms require mechanisms that go beyond pure message propagation, with stricter consistency and robustness guarantees. Because they involve connected users, these applications must provide usable solutions, in particular in terms of latency and availability. At the same time, because they manipulate real-world transactions and objects (money, cars, accommodations) they must also provide a high level of consistency and guarantees. Many of these applications further operate at a planetary scale, and therefore also face stark scalability issues, that make them highly interesting case studies to investigate innovative architectures combining decentralized and centralized elements.

Formalizing and characterizing the needs and behaviors of these new applications seems particularly interesting in order to provide the fertile ground for new systems and novel theoretical work. The area of social applications also offers avenues for knowledge transfer and societal impact, along two dimensions. First, practical and usable approaches, back by a deep understanding of the foundation of distribution and coordination, are likely to find applications in future systems. Second, developers of complex social applications are often faced with a lack of robust scalable services¹ that can be easily exploited to harness the latest understanding of large-scale distributed coordination. We therefore think these applications offer an opportunity to design and deliver modular reusable bricks that can be easily appropriated by a large population of innovative developers without requiring the level of deep understanding usually necessary to implement these solutions from scratch. Providing such reusable bricks is however difficult, as many interesting formal properties are not composable, and a unified composable theory of distributed systems still need to be fully articulated.

3 Research program

3.1 Overview

In order to progress in the three fields described above, the WIDE team is developing a research program which aims to **help developers control and master the inherent uncertainties and performance challenges brought by scale and distribution**.

More specifically, our program revolves around four key challenges.

- Challenge 1: Designing Hybrid Scalable Architectures,
- Challenge 2: Constructing Personalizable Privacy-Aware Distributed Systems,
- Challenge 3: Understanding Controllable Network Diffusion Processes,
- Challenge 4: Systemizing Modular Distributed Computability and Efficiency.

These four challenges have in common **the inherent tension between coordination and scalability in large-scale distributed systems**: strong coordination mechanisms can deliver strong guarantees (in terms of consistency, agreement, fault-tolerance, and privacy protection), but are generally extremely costly and inherently non-scalable if applied indiscriminately. By contrast, highly scalable coordination approaches (such as epidemic protocols, eventual consistency, or self-organizing overlays) perform much better when the size of a system increases, but do not, in most cases, provide any strong guarantees in terms of consistency or agreement.

The above four challenges explore these tensions from *four complementary angles*: from an architectural perspective (Challenge 1), from the point of view of a fundamental system-wide guarantee (privacy protection, Challenge 2), looking at one universal scalable mechanism (network diffusion, Challenge 3), and considering the interplay between modularity and computability in large-scale systems (Challenge 4). These four challenges range from practical concerns (Challenges 1 and 2) to more theoretical questions (Challenges

¹The repeated debugging of MongoDB's replication algorithm (e.g. see <https://aphyr.com/posts/338-jepsen-mongodb-3-4-0-rc3>) is a telling illustration of the difficulties encountered by development teams when building such platforms.

3 and 4), yet present *strong synergies* and *fertile interaction points*. E.g. better understanding network diffusion (Challenge 3) is a key enabler to develop more private decentralized systems (Challenge 2), while the development of a theoretically sound modular computability hierarchy (Challenge 4) has a direct impact on our work on hybrid architectures (Challenge 1).

3.2 Hybrid Scalable Architectures

The rise of planetary-scale distributed systems calls for novel software and system architectures that can support user-facing applications while scaling to large numbers of devices, and leveraging established and emerging technologies. The members of WIDE are particularly well positioned to explore this avenue of research thanks to their experience on de-concentrated architectures combining principles from both decentralized peer-to-peer [47, 58] systems and hybrid infrastructures (i.e. architectures that combines centralized or hierarchical elements, often hosted in well-provisioned data-centers, and a decentralized part, often hosted in a peer-to-peer overlay) [51]. In the short term, we aim to explore two axes in this direction: browser-based communication, and micro-services.

Browser-based fog computing The dramatic increase in the amount of data being produced and processed by connected devices has led to paradigms that seek to decentralize the traditional cloud model. In 2011 Cisco [48] introduced the vision of *fog computing* that combines the cloud with resources located at the edge of the network and in between. More generally, the term *edge computing* has been associated with the idea of adding edge-of-the-network storage and computation to traditional cloud infrastructures [42].

A number of efforts in this directions focus on specific hardware, e.g. fog nodes that are responsible for connected IoT devices [49]. However, many of today's applications run within web browsers or mobile phones. In this context, the recent introduction of the WebRTC API, makes it possible for browsers and smartphones to exchange directly between each other, enabling mobile, or browser-based decentralized applications.

Maygh [77], for example, uses the WebRTC API to build a decentralized Content Delivery Network that runs solely on web browsers. The fact that the application is hosted completely on a web server and downloaded with enabled websites means that webmasters can adopt the Content Delivery Network (CDN) without requiring users to install any specific software.

For us, the ability of browsers to communicate with each other using the WebRTC paradigm provides a novel playground for new programming models, and for a *browser-based fog architecture* combining both a centralized, cloud-based part, and a decentralized, browser-supported part.

This model offers tremendous potential by making edge-of-the-network resources available through the interconnection of web-browsers, and offers new opportunities for the protection of the personal data of end users. But consistently engineering browser-based components requires novel tools and methodologies.

In particular, WebRTC was primarily designed for exchanging media and data between two browsers in the presence of a coordinating server. Its complex mechanisms for connection establishment make many of the existing peer-to-peer protocols inefficient. To address this challenge, we plan to consider two angles of attack. First, we plan to design novel protocols that take into account the specific requirements set by this new technology. Second, we envisage to investigate variants of the current WebRTC model with cheaper connection-establishment protocols, in order to provide lower delays and bandwidth consumption in large-scale browser-based applications.

We also plan to address the trade-offs associated with hybrid browser-cloud models. For example, when should computation be delegated to browsers and when should it be executed on the cloud in order to maximize the quality of service? Or, how can a decentralized analytics algorithms operating on browser-based data complement or exploit the knowledge built by cloud-based data analytics solutions?

Emergent micro-service deployment and management Micro-services tend to produce fine-grained applications in which many small services interact in a loosely coupled manner to produce a wide range of services within an organization. Individual services need to evolve independently of each other over time without compromising the availability of the overall application. Lightweight isolation solutions such as containers (Docker, ...), and their associated tooling ecosystem (e.g. Google's Borg [76], Kubernetes [46])

have emerged to facilitate the deployment of large-scale micro-service-based applications, but only provide preliminary solutions for key concerns in these systems, which we would like to investigate and extend.

Most of today's on-line computer systems are now too large to evolve in monolithic, entirely pre-planned ways. This applies to very large data centres, for example, where the placement of virtual machines to reduce heating and power consumption can no longer be treated using top-down exhaustive optimisation approaches beyond a critical size. This is also true of social networking applications, where different mechanisms—e.g. to spread news notifications, or to recommend new contacts—must be adapted to the different sub-communities present in the system.

To cope with the inherent complexity of building complex loosely-coupled distributed systems while fostering and increasing efficiency, maintainability, and scalability, we plan to study how novel programming techniques based on declarative programming, components and epidemic protocols can help design, deploy, and maintain self-adaptive structures (e.g. placement of VM) and mechanisms (e.g. contact recommendations) that are optimized to the local context of very large distributed systems. To fulfill this vision, we plan to explore a three-pronged strategy to raise the level of programming abstraction offered to developers.

- First, we plan to explore the use of high-level domain-specific languages (DSL) to declare how large-scale topologies should be achieved, deployed, and maintained. Our vision is a declarative approach to describe how to combine, deploy and orchestrate micro-services in an abstract manner thus abstracting away developers from the underlying cloud infrastructures, and from the intricacies involved in writing low-level code to build a large-scale distributed application that scales. With this effort, we plan notably to directly support the twin properties of *emergence* (the adaptation “from within”) and *differentiation* (the possibility from parts of the system to diverge while still forming a whole). Our central objective is to search for principled programming constructs to support these two capabilities using a modular and incremental software development approach.
- On a second strand of work, we plan to investigate how unikernels enable smaller footprints, more optimization options, and faster boot times for micro-services. Isolating micro-services into VMs is not the most adequate approach as it requires the use of hypervisors, or virtual machine monitors (VMMs), to virtualize hardware resources. VMMs are well known to be heavyweight with both boot and run time overheads that may have a strong impact on performances. Unikernels seem to offer the right balance between performance and flexibility to address this challenge. One of the key underlying challenges is to compile directly the aforementioned provided DSL to a dedicated and customized machine image, ready to be deployed directly on top of a large set of bare metal servers.
- Depending on the workload it is subjected to, and the state of its execution environment (network, VMs), a large-scale distributed application may present erratic or degraded performance that is hard to anticipate and plan for. There is therefore a strong need to adapt dynamically the way resources are allocated to a running application. We would like to study how the DSL approach we envisage can be extended to enable developers to express orchestration algorithms based on machine learning algorithms.

3.3 Personalizable Privacy-Aware Distributed Systems

On-line services are increasingly moving towards an in-depth analysis of user data, with the objective of providing ever better personalization. But in doing so, personalized on-line services inevitably pose risks to the privacy of users. Eliminating, or even reducing these risks raises important challenges caused by the inherent trade-off between the level of personalization users wish to achieve, and the amount of information they are willing to reveal about themselves (explicitly or through the many implicit sources of digital information such as smart homes, smart cars, and IoT environments).

At a general level, we would like to address these challenges through protocols that can provide access to unprecedented amounts of data coming from sensors, users, and documents published by users, while protecting the privacy of individuals and data sources. To this end, we plan to rely on our experience in the context of distributed systems, recommender systems, and privacy, as well as in our collaborations with experts in neighboring fields such as machine learning, and security. In particular, we aim to explore different privacy-utility tradeoffs that make it possible to provide differentiated levels of privacy guarantees depending on the context associated with data, on the users that provide the data, and on those that access it. Our

research targets the general goal of privacy-preserving decentralized learning, with applications in different contexts such as user-oriented applications, and the Internet-of-Things (IoT).

Privacy-preserving decentralized learning Personalization and recommendation can be seen as a specific case of general machine learning. Production-grade recommenders and personalizers typically centralize and process the available data in one location (a data-center, a cloud service). This is highly problematic, as it endangers the privacy of users, while hampering the analysis of datasets subject to privacy constraints that are held by multiple independent organizations (such as health records). A decentralized approach to machine learning appears as a promising candidate to overcome these weaknesses: if each user or participating organization keeps its data, while only exchanging gradient or model information, privacy leaks seem less likely to occur.

In some cases, decentralized learning may be achieved through relatively simple adaptations of existing centralized models, for instance by defining alternative learning models that may be more easily decentralized. But in all cases, processing growing amounts of information calls for high-performance algorithms and middleware that can handle diverse storage and computation resources, in the presence of dynamic and privacy-sensitive data. To reach this objective, we will therefore leverage our work in distributed and privacy-preserving algorithms and middleware [50, 52, 53] as well as the results of our work on large-scale hybrid architectures in Objective 1.

Personalization in user-oriented applications As a first application perspective, we plan to design tools that exploit decentralized analytics to enhance user-centric personalized applications. As we observed above, such applications exhibit an inherent trade-off between personalization quality and privacy preservation. The most obvious goal in this direction consists in designing algorithms that can achieve high levels of personalization while protecting sensitive user information. But an equally important one consists in personalizing the trade-off itself by adapting the quality of the personalization provided to a user to his/her willingness to expose information. This, like other desirable behaviors, appears at odds with the way current systems work. For example, a user of a recommender system that does not reveal his/her profile information penalizes other users causing them to receive less accurate recommendations. We would like to mitigate this situation by means of protocols that reward users for sharing information. On the one hand, we plan to take inspiration from protocols for free-riding avoidance in peer-to-peer systems [54, 60]. On the other hand, we will consider blockchains as a tool for tracking and rewarding data contributions. Ultimately, we aim at enabling users to configure the level of privacy and personalization they wish to experience.

Privacy preserving decentralized aggregation As a second setting we would like to consider target applications running on constrained devices like in the Internet-of-Things (IoT). This setting makes it particularly important to operate on decentralized data in a light-weight privacy-preserving manner, and further highlights the synergy between this objective and Objective 1. For example, we plan to provide data subjects with the possibility to store and manage their data locally on their own devices, without having to rely on third-party managers or aggregators, but possibly storing less private information or results in the cloud. Using this strategy, we intend to design protocols that enable users themselves, or third-party companies to query distributed data in aggregate form, or to run data analytics processes on a distributed set of data repositories, thereby gathering knowledge without violating the privacy of other users. For example, we have started working on the problem of computing an aggregate function over a subset of the data in a distributed setting. This involves two major steps: selection and aggregation. With respect to selection, we envision defining a decentralized data-selection operation that can apply a selection predicate without violating privacy constraints. With respect to aggregation, we will continue our investigation of lightweight protocols that can provide privacy with limited computational complexity [43].

3.4 Network Diffusion Processes

Social, biological, and technological networks can serve as conduits for the spread of ideas, trends, diseases, or viruses. In social networks, rumors, trends and behaviors, or the adoption of new products, spread from person to person. In biological networks, diseases spread through contact between individuals, and mutations spread from an individual to its offsprings. In technological networks, such as the Internet

and the power grid, viruses and worms spread from computer to computer, and power failures often lead to cascading failures. The common theme in all the examples above is that the rumor, disease, or failure starts out with a single or a few individual nodes, and propagates through the network, from node to node, to reach a potentially much larger number of nodes.

These types of *network diffusion processes* have long been a topic of study in various disciplines, including sociology, biology, physics, mathematics, and more recently, computer science. A main goal has been to devise mathematical models for these processes, describing how the state of an individual node can change as a function of the state of its neighbors in the network, and then analyse the role of the network structure in the outcome of the process. Based on our previous work, we would like to study to what extent one can affect the outcome of the diffusion process by controlling a small, possibly carefully selected fraction of the network.

For example, we plan to explore how we may increase the spread or speed of diffusion by choosing an appropriate set of seed nodes (a standard goal in viral marketing by word-of-mouth), or achieve the opposite effect either by choosing a small set of nodes to remove (a goal in immunization against diseases), or by seeding a competing diffusion (e.g., to limit the spread of misinformation in a social network).

Our goal is to provide a framework for a systematic and rigorous study of these problems. We will consider several standard diffusion models and extensions of them, including models from mathematical sociology, mathematical epidemiology, and interacting particle systems. We will consider existing and new variants of spread maximization/limitation problems, and will provide (approximation) algorithms or show negative (inapproximability) results. In case of negative results, we will investigate general conditions that make the problem tractable. We will consider both general network topologies and specific network models, and will relate the efficiency of solutions to structural properties of the topology. Finally, we will use these insights to engineer new network diffusion processes for efficient data dissemination.

Spread maximization Our goal is in particular to study spread maximization in a broader class of diffusion processes than the basic independent cascade (IC) and linear threshold (LT) models of influence [68, 66, 67] that have been studied in this context so far. This includes the *randomized rumor spreading (RS)* model for information dissemination [57], *biased* versions of the *voter model* [62] modelling influence, and the (graph-based) *Moran processes* [70] modelling the spread of mutations. We would like to consider several natural versions of the spread maximization problem, and the relationships between them. For these problems we will use the greedy algorithm and the submodularity-based analytical framework of [68], and will also explore new approaches.

Immunization optimization Conversely we would also like to explore immunization optimization problems. Existing works on these types of problem assume a *perfect-contagion* model, i.e., once a node gets infected, it deterministically infects all its non-immunized neighbors. We plan to consider various diffusion processes, including the standard *susceptible–infected (SI)*, *susceptible–infected–recovered (SIR)* and *susceptible–infected–susceptible (SIS)* epidemic models, and explore the extent to which results and techniques for the perfect-contagion model carry over to these probabilistic models. We will also investigate whether techniques for spread maximization could be applied to immunization problems.

Some immunization problems are known to be hard to approximate in general graphs, even for the perfect-contagion model, e.g., the fixed-budget version of the fire-fighter problem cannot be approximated to any $n^{1-\epsilon}$ factor [45]. This strand of work will consider restricted graph families, such as trees or graphs of small treewidth, for such problems. In addition, for some immunization problems, there is a large gap between the best known approximation algorithm and the best known inapproximability result, and we would like to make progress in reducing these gaps.

3.5 Systemizing Modular Distributed Computability and Efficiency

The applications and services envisaged in Objectives 1 and 2 will lead to increasingly complex and multifaceted systems. Constructing these novel hybrid and decentralized systems will naturally push our need to understand distributed computing beyond the current state of the art. These trends therefore demand research efforts in establishing sound theoretical foundations to allow everyday developers to master the design, properties and implementation of these systems.

We plan to investigate these foundations along two directions: first by studying novel approaches to some fundamental problems of *mutual exclusion and distributed coordination*, and second by exploring how we can build a *comprehensive and modular framework* capturing the foundations of *distributed computation*.

Randomized algorithm for mutual exclusion and coordination To exploit the power of massive distributed applications and systems (such as those envisaged in Objectives 1 and 2) or multiple processors, algorithms must cope with the scale and asynchrony of these systems, and their inherent instability, e.g., due to node, link, or processor failures. Our goal is to explore the power and limits of randomized algorithms for large-scale networks of distributed systems, and for shared memory multi-processor systems, in effect providing fundamental building blocks to the work envisioned in Objectives 1 and 2.

For shared memory systems, randomized algorithms have notably proved extremely useful to deal with asynchrony and failures. Sometimes probabilistic algorithms provide the only solution to a problem; sometimes they are more efficient; sometimes they are simply easier to implement. We plan to devise efficient algorithms for some of the fundamental problems of shared memory computing, such as mutual exclusion, renaming, and consensus.

In particular, looking at the problem of *mutual exclusion*, it is desirable that mutual exclusion algorithms be *abortable*. This means that a process that is trying to lock the resource can abort its attempt in case it has to wait too long. Abortability is difficult to achieve for mutual exclusion algorithms. We will try to extend our algorithms for the *cache-coherent* (CC) and the *distributed shared memory* (DSM) model in order to make them abortable, while maintaining expected constant *Remote Memory References* (RMRs) complexity, under optimistic system assumptions. In order to achieve this, the algorithm will use strong synchronization primitives, called compare-and-swap objects. As part of our collaboration with the University of Calgary, we will work on implementing those objects from registers in such a way that they also allow aborts. Our goal is to build on existing non-abortable implementations [59]. We plan then later to use these objects as building blocks in our mutual exclusion algorithm, in order to make them work even if the system does not readily provide such primitives.

We have also started working on blockchains, as these represent a new and interesting trade-off between probabilistic guarantees, scalability, and system dynamics, while revisiting some of the fundamental questions and limitations of consensus in fault-prone asynchronous systems.

Modular theory of distributed computing Practitioners and engineers have proposed a number of reusable frameworks and services to implement specific distributed services (from Remote Procedure Calls with Java RMI or SOAP-RPC, to JGroups for group communication, and Apache Zookeeper for state machine replication). In spite of the high conceptual and practical interest of such frameworks, many of these efforts lack a sound grounding in distributed computation theory (with the notable exceptions of JGroups and Zookeeper), and often provide punctual and partial solutions for a narrow range of services. We argue that this is because we still lack a generic framework that unifies the large body of fundamental knowledge on distributed computation that has been acquired over the last 40 years.

To overcome this gap we would like to develop a systematic model of distributed computation that organizes the functionalities of a distributed computing system into reusable modular constructs assembled via well-defined mechanisms that maintain sound theoretical guarantees on the resulting system. This research vision arises from the strong belief that distributed computing is now mature enough to resolve the tension between the social needs for distributed computing systems, and the lack of a fundamentally sound and systematic way to realize these systems.

To progress on this vision, we plan in the near future to investigate, from a distributed software point of view, the impact due to failures and asynchrony on the layered architecture of distributed computing systems. A first step in this direction will address the notions of *message adversaries* (introduced a long time ago in [75]) and *process adversaries* (investigated in several papers, e.g. [74, 56, 64, 65, 69]). The aim of these notions is to consider failures, not as “bad events”, but as part of the normal behavior of a system. As an example, when considering round-based algorithms, a message adversary is a daemon which, at every round, is allowed to suppress some messages. The aim is then, given a problem P , to find the strongest adversary under which P can be solved (“strongest” means here that giving more power to the adversary makes the problem impossible to solve). This work will allow us to progress in terms of general *layered* theory of distributed computing, and allow us to better *map* distributed computing models and their relations, in the

steps of noticeable early efforts in this direction [74, 41].

3.6 Evolution of our research program (2022-2026)

The overarching goal of WIDE is to provide the practical and theoretical foundations required to address the scale, dynamicity, and uncertainty that characterize modern distributed computer systems. In particular, we would like to explore the inherent tension between scalability and coordination guarantees, by proposing novel techniques and paradigms that facilitate the construction of such systems.

This ultimate goal continues to underpin the team's efforts. On the scientific front, however, distributed systems are undergoing rapid changes, which include the rise of new applications domains, such as Blockchains and cryptocurrencies, and the growth of new technologies, such as distributed Machine Learning and interconnected AI-based decision systems.

The WIDE team is also evolving internally: the arrivals of Barbe Mvondo Djob and Brice Ekane (University of Rennes) has brought new expertise to WIDE, and the opportunity to deepen our understanding of the lower levels of large-scale distributed infrastructures. These novel challenges and opportunities lead us to propose the following four updated objectives.

Objective 1: Large-scale Trustless Sybil-Resistant Systems

We plan to contribute to the theoretical understanding of Blockchain-based and Byzantine-tolerant systems by exploring reusable abstractions that can allow programmers to develop Byzantine-tolerant applications more easily. We plan for example to extend existing work on weak consistency to a BFT setting, building for instance on recent proposals on Byzantine Fault-Tolerant CRDTs [63]. To address scale, we plan to explore novel scalable Byzantine fault-tolerant algorithms, both in the context of closed systems, and then in the more challenging case of open (aka permissionless) systems. Our line of attack is to focus on lightweight BFT primitives that can enable faster and more resource-efficient algorithms [55, 61]. In the case of open systems, we will leverage the expertise of our team in theoretical distributed algorithms and randomized algorithms to address Sybil attacks through novel countermeasures providing (hopefully) cheaper and more equitable alternatives to proof-of-work or proof-of-stake algorithms. One open, yet enticing, question is whether anonymous computing models could provide a path to address this issue. We would also like to investigate how storage can be improved in Blockchains and BFT large-scale systems. Most of these systems are fully replicated, incurring formidable costs (up to 2.6PB of distributed storage in the case of Bitcoin). Coding techniques, that we have used in the past, and adaptable redundancy based on Byzantine quorums [71] are some avenues we would like to explore to address this challenge.

Objective 2: Robustness and Security at Scale

Although WIDE did not focus initially on security issues per se, our historical interest in privacy concerns and Byzantine fault-tolerance has progressively led us to consider a broader range of security properties in distributed and decentralized systems, ranging from anonymity (in anonymity networks, explored in the PhD of Quentin Dufour) to malware protection through large-scale computations.

In terms of malware protection, we would like to harness the power of distribution and collaborative data gathering to help antivirus designers improve and optimize malware detection. We plan in particular to work on the automatic creation of test datasets for antivirus software using automated mutation techniques, building upon our preliminary work in this area. Such a tool is of primary importance in both the academic and industrial fields to be able to quantify the effectiveness of new countermeasures.

On the front of privacy, we plan to investigate the design of a distributed digital data vault able to securely store personal data, leveraging our experience on privacy-preserving decentralized systems [43], and on trusted-execution environments (e.g. SGX). We have started collaborating with the CIDRE team at Inria Rennes, with colleagues at KTH (Sweden), and with the company AriadNext (H2020 Soteria project) on these topics.

At an infrastructure level, and following the recruitment of Djob Mvondo, we plan to explore how progress in virtualization can help advance the team's agenda in terms of large-scale robustness, in particular in a cloud-computing setting [72, 73]. Specifically we would like to investigate how novel heterogeneous

architectures that embed a range of ASICs and specialized units (GPU, FPGA, SMARTNIC, PIM-devices) can be leveraged to provide more robust and more efficient virtualized services.

Objective 3: Fundamentals of distributed randomized algorithms

We plan to continue our theoretical exploration of simple randomized distributed algorithms, where individual entities (nodes or mobile agents) have limited computation and communication power, and are often unreliable. These distributed randomized algorithms are closely related to the mechanisms we plan to explore for Sybil attack protection (Objective 1), and privacy protection (Objective 2).

More concretely, we will investigate three settings: in the first setting, agents perform independent or mildly dependent random walks on a graph, and interact when they meet. In the second (more traditional) setting, the interacting entities are the nodes of graph. Finally, in a third setting, nodes are the computing entities and the goal is to modify the graph edges to achieve certain desirable graph properties (an expander graph [44], or a k-nearest neighbor graph), by means of local decentralized operations (typically adjacent nodes interact by exchanging some of their incident edges). In all three cases, we will strive to derive time- and space- optimal algorithms, with strong robustness guarantees.

4 Application domains

WIDE's research, while primarily focused on the progress of scientific knowledge, has a wide range of potential application domains. Our work on modular algorithmic abstraction has strong links to and is inspired by Software engineering. Our work on graph analysis, and social media practice is of direct relevance to the web, while our work on randomized processes can be applied to track epidemics. Our work on recommenders and kNN graph construction applies to search engines. Finally our work on privacy is of keen interest to Law scholars, as demonstrated by several interdisciplinary projects with colleagues from this discipline.

5 Social and environmental responsibility

- Davide Frey and Francois Taïani participate to the sustainable-development working group at Inria Centre at Rennes University.
- Davide Frey is part of the SENS (science and environment) group at Inria Centre at Rennes University.

6 Highlights of the year

- Yerom David Bromberg and Brice Ekane organized the Annual Symposium of the French Computer Science Society (Congrès Annuel de la Société Informatique de France) on 5 and 6 June 2025.
- Timothé Albouy, Davide Frey, Manon Sourisseau, and François Taïani participated in the filming of the TV show *Esprit Sorcier: Les Surprises de la Recherche* (Wizard Spirit: The Surprises of Research), filmed in Rennes in April 2025 with the participation of 340 middle school students from the Rennes school district.
- WIDE has two new associate teams, one with the University of Cambridge, UK (DAME) coordinated by George Giakkoupis, and one with the Université de Yaoundé I, Cameroun (EASy-AI) coordinated by Yerom David Bromberg.
- In 2025, WIDE's research results continued to be published in some of the most visible and prestigious conferences of its field (AAAI, PoPETS, NSDI, DISC, PerCom, DSN).

7 Latest software developments, platforms, open data

7.1 Latest software developments

7.1.1 DecentralizedFlower

Name: DecentralizedFlower

Keyword: Decentralized Learning

Functional Description: DecentralizedFlower is a framework to test decentralized machine learning algorithms in a cluster environment, in a production environment, and in a combination of the two. The framework enables developers to test algorithms on a testing environment and then seamlessly deploy them into a production setting. The software is based on the Flower federated-learning library developed by the University of Cambridge and the German Company Adap.

Contact: Davide Frey

7.1.2 nodemanager

Keywords: Peer-to-peer, Peer-sampling, Distributed, Distributed Applications

Functional Description: Nodemanager is a solution for setting up peer-to-peer applications. It is essentially written in Rust, but provides interfaces for use in Python.

Contact: Davide Frey

7.1.3 DecentralizedDeclearn

Keyword: Decentralized Learning

Functional Description: DecentralizedDeclearn is a Python library for testing decentralized machine learning algorithms. This library provides developers with a simulation framework for testing their applications before deploying them. This library is based on the Declearn library, which is a Python package providing a framework for federated learning. It was developed by the Magnet team at Inria.

Contact: Davide Frey

7.1.4 decentralised-data-wallet

Name: SOTERIA Data Wallet Prototype

Keywords: Privacy, Data management, Data analytics, Distributed systems, Cryptography, Decentralized Learning

Functional Description: data-wallet-prototype is a Rust library that provides the basic functionality of a digital data wallet, with the particular constraint that distributed computations and interaction with outside parties occur in a fully decentralised manner. This contrasts with existing solutions that rely on centralised systems, such as cloud providers, which are typically used to store encrypted personal information and carry out computations.

In short, this allows: - Users to securely store their own personal data on their own devices without relying on external service providers (using suitable encryption and hardware security measures) - Third parties, such as research bodies, to perform computations across a network of digital wallets in a decentralized manner without compromising user privacy (using protocols from the literature designed to make this possible).

This software is designed to be flexible with respect to the hardware limitations of its environment, enabling it to run on a range of personal devices, including Android systems.

This was funded as part of the SOTERIA Digital Security and Privacy project.

Contact: Davide Frey

7.1.5 CAC

Name: Context Adaptative Cooperation

Keyword: Distributed systems

Functional Description: Context-Adaptive Cooperation (CAC) is a novel cooperation abstraction that allows an arbitrary set of processes to propose values while multiple value acceptances are triggered. Furthermore, each acceptance comes with information about other acceptances that can possibly occur. This code simulates an instance of CAC. Let n be the number of processes, t the number of Byzantine processes whose behaviour may diverge from the initial one. Let m be the total number of proposals made by m different processes. At the end of the simulation, each of the $n-t$ non-Byzantine processes will have accepted one or more of the initial proposals. But if we look at the intersection of all these proposals, only one remains.

Contact: Davide Frey

7.1.6 QAAT

Name: Quasi-Anonymous Asset Transfer

Keywords: Asset transfer, Distributed computing

Functional Description: QAAT is the first asset transfer system that achieves anonymity, and consensus-freedom while incurring as-low-as-possible storage and communication costs. QAAT provides the following three properties. - Quasi-anonymity: QAAT hides the amount and the receiver's identity of every asset transfer. - Lightness: QAAT uses only succinct cryptographic schemes, i.e. with at most polylogarithmic proof size and verification time. Moreover, the storage cost incurred by each process is linear in its number of transfers for a fixed security parameter, and the associated communication cost remains as low as possible. - Consensus-Freedom: QAAT is a deterministic algorithm that can operate in an asynchronous setting prone to failures, thereby supporting responsive applications. This software artifact, currently under development, provides the first working implementation of the QAAT algorithm

Contact: Davide Frey

7.1.7 Splitchain

Name: Splitchain Protocol

Keywords: Blockchain, Rust, Distributed systems

Functional Description: This software is a node in the distributed Splitchain system. It allows to join the system, submit new transactions, participate in consensus to create new blocks, and manages automatically the split and merge of the shards, and the routing of data.

Contact: Davide Frey

7.1.8 oversim-ipfs

Name: OverSim for the InterPlanetary File System (IPFS)

Keywords: C++, Distributed systems, Distributed Storage Systems, IPFS, Network simulator

Functional Description: A fork of the event-driven simulation framework OverSim (2007, Baumgart et al)*.

This project aims to provide a realistic model of IPFS nodes interacting over a network. This includes generating file chunks, publishing provider records, querying the overlay and providing methods to analyse file availability over time.

It also aims to study a novel re-publishment algorithm under development in conjunction with the COAST team and Hivenet.

* <http://www.oversim.org/>

Contact: Davide Frey

Partner: Hivenet

7.1.9 PPFDIMIS

Name: Privacy-Preserving and fully-Distributed Identity Management System

Keywords: Distributed systems, Privacy, Identity management

Functional Description: The system consists of certificate issuers, verifiers, and users.

- A user requests a certificate to an issuer in order to certify personal data (identity, age, education level, medical information, etc.).
- The issuer sends the user a certificate, proving that the user's personal data is true.
- The user sends their certificate to a verifier to access a service (authorisation, purchase, entry, etc.).
- The verifier grants or denies the user access to the service based on the certificate's validity. Verifiers and issuers perform these operations without knowing each other's identities, thus preserving everyone's anonymity. Verifiers and issuers can also manage certificate revocation via a pseudo-consensus mechanism not based on a blockchain.

URL: <https://gitlab.inria.fr/WIDE/dims/>

Contact: Davide Frey

8 New results

8.1 Distributed Algorithms and Blockchain

8.1.1 Contention-Aware Cooperation

Participants: Michel Raynal, Davide Frey, François Taïani.

As shown by Reliable Broadcast and Consensus, cooperation among a set of independent computing entities (sequential processes) is crucial in fault-tolerant distributed computing. Considering n -process asynchronous message-passing systems where some processes may be Byzantine, this work [21] introduces a novel cooperation abstraction, Contention-Aware Cooperation (CAC). While Reliable Broadcast is a one-to- n cooperation abstraction and Consensus is an n -to- n cooperation abstraction, CAC is a d -to- n cooperation abstraction where d ($1 \leq d \leq n$) varies with each run and remains unknown to the processes. Correct processes accept the same set of ℓ pairs $\langle v, i \rangle$ (v is the value proposed by p_i) from the d proposer processes, where $1 \leq \ell \leq d$ and (as d) ℓ remains unknown to the processes (except in specific cases). Those ℓ values are accepted one at a time, potentially in different orders at each process. In addition, CAC provides each process with an imperfect oracle that provides insights into the values that they may accept in the future. Interestingly, the CAC abstraction is particularly efficient in favorable circumstances, when the oracle becomes accurate, which processes can detect. To illustrate its practical utility, the work details two applications leveraging CAC: a fast consensus implementation optimized for low contention (named Cascading Consensus), and a novel naming problem that can be solved under full asynchrony. All algorithms presented require signatures.

8.1.2 Ethical Risk Analysis of L2 Rollups

Participants: Davide Frey, François Taïani.

Layer 2 rollups improve throughput and fees, but can reintroduce risk through operator discretion and information asymmetry. In this work [18], we ask which operator and governance designs produce ethically problematic user risk. We adapt Ethical Risk Analysis to rollup architectures, build a role-based taxonomy of decision authority and exposure, and pair the framework with two empirical signals, a cross sectional snapshot of 129 projects from L2BEAT and a hand curated incident set covering 2022 to 2025. We analyze mechanisms that affect risks to users' funds, including upgrade timing and exit windows, proposer liveness and whitelisting, forced inclusion usability, and data availability choices. We find that ethical hazards rooted in L2 components control arrangements are widespread: instant upgrades without exit windows appear in about 86 percent of projects, and proposer controls that can freeze withdrawals in about 50 percent. Reported incidents concentrate in sequencer liveness and inclusion, consistent with these dependencies. We translate these findings into ethically grounded suggestions on mitigation strategies including technical components and governance mechanisms.

8.1.3 Communication abstractions in systems prone to malicious attacks

Participants: Achour Mostefaoui.

Abstractions play a central role in distributed computing, as they capture essential synchronization properties. Equivalence results among abstractions clarify their relative computational power. While many such equivalences are well established in crash-prone systems, far less is known about Byzantine-prone environments, where faulty processes may behave arbitrarily.

This year, we revisited the equivalence landscape in Byzantine systems, with a particular focus on communications between processes, namely, shared registers and broadcast abstractions. We establish three new reductions. Specifically, we prove that the broadcast abstraction called Byzantine Set-Constrained Delivery Broadcast (BSCD-Broadcast) can implement a Snapshot/Append object and vice versa (i.e., a two-way reduction), and that the FIFO variant of the renowned Byzantine Reliable Broadcast (BRB-Broadcast) can implement BSCD-Broadcast under a majority of correct processes. Interestingly, this assumption mirrors the one required in crash-prone systems for a similar transformation.

On another axis, we worked on the construction of a privacy-preserving single-writer multi-reader (SWMR) atomic register in a Byzantine-prone distributed model. Specifically, we consider a closed model, in which one process can write values in the register and only a subset of the other processes are allowed to read the value. The aim is to ensure that processes that do not have the requisite reading right are unable to read the content of the register, even when they are Byzantine. This makes the content of the register private. We ensure this privacy by encoding the value written by the writer, using secret sharing, into multiple shards and disseminating them among the participating reader processes. The technical challenge is then to organize the coordination between the correct reading processes to achieve Byzantine linearizability, without knowing the content of the register. The main contribution of this work is a linearizable read-write (R/W) privacy-preserving register for $t < \frac{n}{7}$, where t is the number of Byzantine processes and n denotes the total number of processes in the system.

8.1.4 Discreet: Distributed delivery service with context-aware cooperation

Participants: Davide Frey.

End-to-end encrypted messaging applications such as Signal became widely popular thanks to their capability to ensure the confidentiality and integrity of online communication. While the highest security guarantees were long reserved to two-party communication, solutions for n -party communication remained either

inefficient or less secure until the standardization of the MLS Protocol (Messaging Layer Security). This new protocol offers an efficient way to provide end-to-end secure communication with the same guarantees originally offered by the Signal Protocol for two-party communication. However, both solutions still rely on a centralized component for message delivery, called the Delivery Service in the MLS Protocol. The centralization of the Delivery Service makes it an ideal target for attackers and threatens the availability of any protocol relying on MLS. In order to overcome this issue, we proposed DiSCreet (Distributed delIVery Service with Context-awaRE coopEraTion), a design that allows clients to exchange protocol messages efficiently and without any intermediary. It uses a Probabilistic Reliable-Broadcast mechanism to efficiently deliver messages and the Cascade Consensus Protocol to handle messages requiring an agreement. Our solution strengthens the availability of the MLS Protocol without compromising its security. We compare the theoretical performance of DiSCreet with another distributed solution, the DCGKA protocol, and detail the implementation of our solution. We published this work in the Annals of Telecommunications [19]. This work was done in the context of the Alvearium Inria Challenge and involved a collaboration with Ludovic Paillat and Amine Ismail from Hive Computing as well as with Claudia Lavigna Ignat from the LORELEY Inria team and Mathieu Turuani from the PESTO Inria team.

8.1.5 Luby’s MIS algorithms made self-stabilizing

Participants: George Giakkoupis.

In [17], we reconsider two well-known distributed randomized algorithms computing a maximal independent set, proposed in the seminal work of Luby (1986). We enhance these algorithms such that they become self-stabilizing without sacrificing their run-time, i.e., both stabilize in $O(\log n)$ synchronous rounds with high probability on any n -node graph. The first algorithm gets along with three states, but needs to know an upper bound on the maximum degree. The second does not need any information about the graph, but uses a number of states that is linear in the node degree. Both algorithms use messages of logarithmic size.

This work was done in collaboration with Volker Turau (Hamburg University of Technology), and Isabella Ziccardi (IRIF, Paris).

8.1.6 On the h -majority dynamics with many opinions

Participants: George Giakkoupis.

In [25], we present the first upper bound on the convergence time to consensus of the well-known h -majority dynamics with k opinions, in the synchronous setting, for h and k that are both non-constant values. We suppose that, at the beginning of the process, there is some initial additive bias towards some plurality opinion, that is, there is an opinion that is supported by x nodes while any other opinion is supported by strictly fewer nodes. We prove that, with high probability, if the bias is $\omega(\sqrt{x})$ and the initial plurality opinion is supported by at least $x = \omega(\log n)$ nodes, then the process converges to plurality consensus in $O(\log n)$ rounds whenever $h = \omega(n \log n / x)$. A main corollary is the following: if $k = o(n / \log n)$ and the process starts from an almost-balanced configuration with an initial bias of magnitude $\omega(\sqrt{n/k})$ towards the initial plurality opinion, then any function $h = \omega(k \log n)$ suffices to guarantee convergence to consensus in $O(\log n)$ rounds, with high probability. Our upper bound shows that the lower bound of $\Omega(k/h^2)$ rounds to reach consensus by Becchetti et al. (2017) cannot be pushed further than $\tilde{\Omega}(k/h)$. Moreover, the bias we require is asymptotically smaller than the $\Omega(\sqrt{n \log n})$ bias that guarantees plurality consensus in the 3-majority dynamics: in our case, the required bias is at most any (arbitrarily small) function in $\omega(\sqrt{x})$ for any $k \geq 2$.

This work was done in collaboration with Francesco d’Amore (Gran Sasso Science Institute, Italy), Niccolò d’Archivio (COATI Inria team, Sophia Antipolis), and Emanuele Natale (CNRS, COATI Inria team, Sophia Antipolis).

8.2 Large scale Cloud environments

8.2.1 Off-the shelf network traffic analysis

Participants: Barbe Mvondo Djob, Yerom David Bromberg.

Offloading malware detection to large scale platforms poses serious threats in terms on information flow control for mobile devices. Especially for stalkerwares that require the analysis of users network activity. We show in [32] that an on-device approach running at the kernel-level can achieve the same level of protection without leaking any user network traffic data and with a minimal overhead on applications performance. Furthermore, our approach is more secure and performant than on-device VPNs which were the standard approaches to dealing with these issues.

8.2.2 Towards efficient kernel network processing for VPNs

Participants: Honore Cesaire Mounah, Barbe Mvondo Djob, Yerom David Bromberg.

VPNs are software programs essential for both several use cases such as securing remote accesses or providing privacy online. However, at scale, in our work [30], we uncover that VPNs suffer from several bottlenecks due to their underlying designs. Concretely, their performance can drop by more than 65% when hitting 80% of the network capacity of the network card, with different behaviours in terms of input network traffic. We propose several design changes leveraging existing routines in the Linux kernel, to improve performance and reach optimal performance. Concretely, the best open-source VPNs, Wireguard, benefits from our design almost reaching the full duplex capacity of the underlying network card capacity, independently of the input traffic variations.

This work was done in collaboration with Julia Lawall from Inria Paris (Whisper team).

8.2.3 Efficient Load balancing for multi-tier applications

Participants: Brice Ekane Apah, Barbe Mvondo Djob, Yerom David Bromberg.

We introduce DISC [26], a system that tackles the backpressure problem in multi-tier and microservice applications, where large response payloads (often produced by backend services like databases) are redundantly relayed through intermediate and frontend tiers, hurting scalability. DISC lets multiple tiers safely share the same TCP connection so that final response data can bypass unnecessary tiers, while lightweight metadata (headers/footers) still follows the normal path. Unlike prior approaches, DISC works with arbitrary multi-tier depths, heterogeneous protocols (e.g., HTTP, IMAP), and TLS, and requires only modest, localized application changes. Evaluations on both classic benchmarks and modern microservices show substantial gains: up to 41.5% lower cumulative CPU usage, up to 45% higher throughput, and dramatic tail-latency reductions (up to 5.7×), effectively restoring performance independence between tiers and making systems scale where it actually matters—at the backend.

This work was done in collaboration with Alain Tchana and Renaud Lachaize from the University of Grenoble Alpes (KraKos team), and Daniel Hagimont from University of Toulouse III (Sepia team).

8.2.4 Containers as alternatives for Cloud gaming

Participants: Adrien Gegout, Barbe Mvondo Djob, Davide Frey.

VMs are the standard isolation environments for large scale cloud gaming infrastructures. In [28], we explore an alternative, containers. Concretely, we show why VMs are rigid for the fast changing and flexible cloud gaming environments and perform early evaluations to show that containers can not only be on par in terms of gaming experience but can lead to energy savings opportunities. We present an early design showing how VMs can be replaced and discussed several potential issues especially regarding GPU sharing.

This work was done with Pascal Manchon from Blacknut.

8.2.5 Disconnecting Users from Virtual Worlds with a Single Packet: an Unreal Untold Story

Participants: Hugo Bertin, Yerom David Bromberg.

Online worlds, with online games at the forefront, have become ubiquitous in our lives. In 2024, the gaming industry's worldwide revenue was estimated at US\$455 billion. Yet, this industry is facing a growing number of cheating actors and techniques. In this paper [22], we introduce new attacks targeting multiplayer games based on Unreal Engine (UE), such as Fortnite, PUBG, and Valorant. These attacks disconnect players from ongoing game sessions against their will. Cheaters can launch them as a Denial-of-Service against opponents with very few packets (sometimes only one) to steal the victory from the target without exposing themselves as fraudsters. This paper shows how such issues present in a single game engine can spread widely across several games produced by different editors, focusing on Unreal Engine, whose source code is publicly available. UE is also commonly found in digital twins, virtual reality, and other Metaverse solutions. We present our analysis of the design and implementation choices made within Unreal Engine. We cover how to exploit UE networking protocols and discuss how to defeat some common countermeasures used on the Internet against IP spoofing. We propose some mitigation strategies for video game developers.

This work was done in collaboration with Ilies Benhabbour (KAUST) and Marc Dacier (KAUST).

8.3 Artificial Intelligence and Machine Learning

8.3.1 Strengthening malware analysis against obfuscation and packing

Participants: Victoire Nganfang, Barbe Mvondo Djob, Yerom David Bromberg.

Existing malware detection techniques fall short against obfuscation and packing techniques. In this work [33], we introduce a new vision-based malware detector designed to remain effective and robust against the aforementioned techniques. Concretely, we introduce DroidHunter that disassembles apps to extract low-level Smali instructions, encodes each instruction as a single RGB pixel, and produces semantic-rich images that preserve opcode and operand information. Our large scale evaluation on approximately 500K APKs, including obfuscated malwares, show up to 99.9% detection accuracy, and outclass nine state-of-the-art detectors. Furthermore, our work achieves stronger resistance to concept drift which is essential for robustness over time.

This work was done with Simon Queyruat, Valerio Schiovani (University of Neuchatel), and Vianney Kengne Tchendji (University of Dschang).

8.3.2 Low-Cost Privacy-Preserving Decentralized Learning

Participants: Dimitri Lereverend, Davide Frey, François Taiani.

Decentralized learning (DL) is an emerging paradigm of collaborative machine learning that enables nodes in a network to train models collectively without sharing their raw data or relying on a central server. In this work, we introduced Zip-DL [23], a privacy-aware DL algorithm that leverages correlated noise to achieve robust privacy against local adversaries while ensuring efficient convergence at low communication

costs. By progressively neutralizing the noise added during distributed averaging, Zip-DL combines strong privacy guarantees with high model accuracy. Its design requires only one communication round per gradient descent iteration, significantly reducing communication overhead compared to competitors. Our work established theoretical bounds on both convergence speed and privacy guarantees. Moreover, it demonstrated Zip-DL’s practical applicability with extensive experiments that show it outperforms state-of-the-art methods in the accuracy vs. vulnerability trade-off. Specifically, Zip-DL (i) reduces membership-inference attack success rates by up to 35% compared to baseline DL, (ii) decreases attack efficacy by up to 13% compared to competitors offering similar utility, and (iii) achieves up to 59% higher accuracy to completely nullify a basic attack scenario, compared to a state-of-the-art privacy-preserving approach under the same threat model. These results position Zip-DL as a practical and efficient solution for privacy-preserving decentralized learning in real-world applications.

This work was done in collaboration with Romaric Gaudel from the MALTE Inria team, and with Sayan Biswas, Anne-Marie Kermarrec, Rafael Pires, and Rishi Sharma from EPFL, Lausanne.

8.4 Load Balancing

8.4.1 An asymptotically optimal algorithm for generating bin cardinalities

Participants: Dimitrios Los.

In the balls-into-bins setting, n balls are thrown uniformly at random into n bins. The naïve way to generate the final load vector takes $\Theta(n)$ time. However, it is well-known that this load vector has with high probability bin cardinalities of size $\Theta(\log n / \log \log n)$. In [15], we present an algorithm in the RAM model that generates the bin cardinalities of the final load vector in the optimal $O(\log n / \log \log n)$ time in expectation and with high probability. Further, the algorithm that we present is still optimal for any $m \in [n, n \log n]$ balls and can also be used as a building block to efficiently simulate more involved load balancing algorithms. In particular, for the Two-Choice algorithm, which samples two bins in each step and allocates to the least-loaded of the two, we obtain roughly a quadratic speed-up over the naïve simulation.

This work was done in collaboration with Luc Devroye (McGill University, Canada).

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

9.1.1 CIFRE with Broadpeak

Participants: Yerom David Bromberg, Barbe Mvondo Djob, Alexandre Duvivier.

The goal of this thesis is to design and implement mechanisms that improve the performance of cache servers and, consequently, improving services that rely on the latter, such as streaming services provided by BroadPeak. This thesis is supervised by Yerom David Bromberg, Barbe Mvondo Djob, and Nicolas Le Scouarnec (Broadpeak). The currently deployed systems at Broadpeak achieve up to 60Gbps and can even reach 150Gbps regarding network throughput. The goal is to achieve 400Gbps on the existing hardware with novel software designs while reducing energy consumption. The thesis will explore ideas that revolve around improving the interaction of user-space applications with kernel network stack subsystems.

9.1.2 CIFRE with Blacknut: Efficient Containerized Cloud-Gaming Platforms

Participants: Davide Frey, Barbe Mvondo Djob, Adrien Gegout.

Cloud gaming enables users without high-end consoles or computers to play video games online on any device with a compatible Internet connection. Users send their commands via a gamepad to a remote server, which applies them and transmits a video stream with game images. Although this paradigm requires few resources on the part of users, it generates a high consumption of resources and energy in the cloud to provide a good quality of service to users with games that perform well, even at start-up. This thesis, supervised by Davide Frey, Barbe Mvondo Djob, Pascal Manchon (Blacknut), and Eric L'Hostis (Blacknut) aims to reduce this resource consumption while improving performance as perceived by users. In particular, we aim on the one hand to enable games to run on containers instead of virtual machines as they do today, and on the other, to predict user demands by pre-allocating resources where it is really useful and necessary.

9.1.3 Flexible Virtualization for Processing In Memory YUMPIM (ANR - PRCE)

Participants: Brice Ekane Apah, Yerom David Bromberg.

Processing In-Memory (PIM) follows near-data processing principles by moving computation closer to where data resides. This project focuses on UPMEM because it is currently the only commercial PIM technology that can be deployed on off-the-shelf servers using standard DDR4/DDR5 memory protocols, it comes with an SDK that eases PIM application development, and it is widely studied in recent research. Prior work shows UPMEM can be a viable alternative to CPUs and GPUs across various workloads (e.g., machine learning, bioinformatics, cryptography), offering high flexibility thanks to its general-purpose compute cores and often delivering better energy efficiency when it outperforms CPUs/GPUs (up to 23.2x vs CPUs and 2.54x on average vs GPUs).

The core objective of the project is UPMEM virtualization for cloud environments. Since virtualization is fundamental to cloud computing for resource sharing and efficient utilization, enabling broad PIM adoption requires supporting time-sharing on UPMEM so that dynamically arriving jobs from uncoordinated users can run. Achieving such time-sharing remains challenging on UPMEM and other current PIM technologies.

In collaboration with Alain Tchana from Krakos Inria team and VATES

10 Partnerships and cooperations

10.1 International initiatives

10.1.1 Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program

EASy-AI

Title: Efficient Systems for Enhancing AI Sustainability

Duration: 2025 -> 2027

Coordinator: Paulin Melatagia (paulinyonta@gmail.com)

Partners:

- Université de Yaoundé I (Cameroun)

Inria contact: Yerom David Bromberg

Summary: The artificial intelligence (AI) market is experiencing rapid growth, with global spending projected to reach \$154 billion in 2023 and surpass \$300 billion by 2026, according to IDC. Cameroon is not exempt from this trend. Many Cameroonian researchers, particularly in the computer science departments at the universities of Dschang and Yaoundé I, are working on specific local challenges that are often absent in Western countries. They are developing AI models for applications such as malaria detection or the processing of local vernacular languages. Paulin Melatagia, recognized across

Africa for his expertise in AI, has been working in these fields long before the global boom in artificial intelligence.

However, training AI models is particularly energy-intensive, requiring vast amounts of data and significant computing power, which poses a major challenge for developing countries like Cameroon, where power outages are frequent. These unpredictable outages, which can last from a few seconds to several days, disrupt the training of AI models, leading to increased energy consumption as operations must be restarted after each interruption.

This project proposes an innovative approach based on operating systems (OS) to support energy-efficient AI adapted to the conditions of power instability. The goal is to optimize resource management, particularly memory and task scheduling, to reduce energy consumption without compromising the performance of the algorithms. In the event of power outages, intelligent resource management would allow the system to operate in a degraded mode, thereby minimizing the impact on performance. Additionally, the integration of technologies such as intelligent sleep modes or the use of low-power cores would enable systems to adjust to AI workloads while remaining energy-efficient and resilient in the face of frequent interruptions.

This project aims to provide viable solutions for resource-limited environments like Cameroon, while addressing the specific challenges posed by frequent power outages.

10.1.2 Inria associate team not involved in an IIL or an international program

DAME

Title: Distributed Algorithms and Markov Chains with Evolving Data

Duration: 2025 -> 2027

Inria PI: George Giakkoupis

Partners:

- University of Cambridge, UK - PI: Thomas Sauerwald

Summary: In today's dynamic data landscape, where information is massive, distributed, and constantly changing, the traditional computational models assuming static data are inadequate. Modern algorithms must be designed to handle frequent data updates, adapting their output in real-time to maintain accuracy, e.g., maintain an approximately sorted permutation as the ranking of the data changes, maintain a small proper vertex coloring of the nodes in a mobile wireless network to facilitate efficient scheduling, or maintain a good balance across servers to which new jobs are allocated and old jobs are completed. This requires the ability to process changes without precise knowledge of the timing or location of changes, a common scenario in large-scale, distributed systems.

In this project, we will explore various new and existing dynamic models, including distributed variants of the evolving data framework, distributed variants of the stochastic graph model, and temporal graph models. We will design and analyse efficient algorithms and processes for fundamental distributed problems, both on fixed and dynamic graph settings. In fixed-graph problems, the graph does not change but the input data, which are distributed among the vertices, may change over time. Examples of such problems we plan to study are sorting with evolving ranks, routing with moving targets, and load balancing on graphs under evolving loads. In dynamic-graph problems, the graph itself changes, in particular, the set of edges may change over time. We will focus on evolving variants of classical local problems, such as maximal independent set (MIS) and maximal matching. We will also consider global problems, including connectivity problems and information diffusion. The main tools we will use are potential function arguments, combined with tools for the analysis of Markov chains.

10.2 National initiatives

ANR JCJC Project sGOV (2023-2027)

Participants: Barbe Mvondo Djob, Yerom David Bromberg.

In this project, we propose to design smart governors (sGOV) to tackle the sub-optimal energy management of idle VMs in the Cloud. In a nutshell, the main objective of sGOV is to identify VMs idle periods, and not account the idle period in the computing of the next CPU state to switch. sGOV design goals are (i) genericity: should be generic enough to be applied to mainstream virtualization systems, and (ii) non-intrusiveness: should not require legacy code to run in user VMs to favor adoption by Cloud providers.

Our core idea with sGOV is that VMs idle periods have specific signatures regarding the interaction between the VM and virtualization system. For example, when a process in a VM stalls waiting for an I/O event (e.g., the arrival of a network packet), no processing is performed on its I/O device interface until the event arises. However, a VM waiting for a hardware event such as the network packet will not behave similarly as a VM waiting for a software interrupt or signal from a process (e.g., SIGALARM signal). Additionally, these behaviors can differ depending on the hardware architecture — a sleep() instruction will not follow the same pattern on an Intel CPU as on AMD or ARM for example.

Partners: IRISA (coordinator, U. Rennes). Budget: 286 814.5€

ANR Second Chance (2023-2027, PRCE)

Participants: Yerom David Bromberg, Barbe Mvondo Djob.

Virtualization is a key technology for datacenters and cloud computing, enabling flexible resource allocation through virtual machines (VMs). Running multiple VMs on the same physical host reduces hardware and management costs while minimizing environmental impact. Central to this process is the hypervisor, a software layer that abstracts physical resources into virtual ones for VMs, each running its own guest operating system to support high-performance applications like web services, databases, and AI tasks. While containers, such as those managed by Docker or Podman, are widely used, they complement rather than replace hypervisors, which offer advanced features like security, performance isolation, persistent storage, and snapshot management. Public cloud platforms often encapsulate containers from different tenants within separate VMs. A critical hypervisor capability is live VM migration, a mature technique that moves a running VM between physical machines without disrupting operations or degrading performance. This feature is essential for cloud and datacenter platforms, supporting administrative tasks while ensuring application availability and performance, with providers like Google performing millions of such migrations monthly.

Given that live migration is commonly used for applications with stringent availability and performance requirements, addressing the problem involves several challenges: determining migration safety without being overly conservative, ensuring acceptable application performance during and after migration, developing extensible techniques to handle new types of CPU feature heterogeneity and emerging application workloads, and maintaining transparency for application developers by avoiding modifications or recompilation of guest code.

Partners: IRISA (coordinator, U. Rennes), LIG (Grenoble INP), Orange Business Services (Eolas).

ANR Project ByBloS (2021-2025)

Participants: George Giakkoupis, Michel Raynal, Davide Frey, Yerom David Bromberg, François Taïani, Timothé Albouy.

Blockchain-based systems have over the last 10 years profoundly impacted society and research. They come however with many inefficiencies, that are inherent to the problem they attempt to solve, Byzantine Tolerant Agreement, one of the most difficult problems of distributed computing. Many Blockchain-based

applications do not require the strong guarantees that an agreement provides. Building on this insight, Byblos seeks to explore the design, analysis, and implementation of lightweight Byzantine decentralized mechanisms for the systematic construction of large-scale Byzantine-tolerant Privacy-Preserving distributed systems.

Partners: IRISA (coordinator, U. Rennes) in Rennes, LIRIS (INSA Lyon) in Lyon, and LS2N (Université de Nantes) in Nantes. Budget: 252 220€

Inria Challenge Project FedMalin

Participants: François Taïani, Davide Frey, Cyrille Kenfack.

FedMalin (project.inria.fr/fedmalin/) is a research project that spans 11 Inria research teams and aims to push FL research and concrete use-cases through a multidisciplinary consortium involving expertise in ML, distributed systems, privacy and security, networks, and medicine. We propose to address a number of challenges that arise when FL is deployed over the Internet, including privacy and fairness, energy consumption, personalization, and location/time dependencies.

FedMalin will also contribute to the development of open-source tools for FL experimentation and real-world deployments, and use them for concrete applications in medicine and crowdsensing.

The FedMalin Inria Challenge is supported by Groupe La Poste, sponsor of the Inria Foundation.

Within Fedmalin, Davide Frey and François Taïani co-supervised the PhD thesis of Rémy Raes, together with Lionel Seinturier and Romain Rouvoy from the Spirals team from Inria Lille. Davide Frey also supervises the work of Cyril Kenfack (Engineer) in order to contribute to a benchmarking environment for the experimentation with federated and decentralized learning platforms and algorithms.

Inria Challenge Project Alvearium

Participants: François Taïani, Davide Frey.

The Alvearium project (project.inria.fr/alvearium/) aims to provide a sovereign alternative peer-to-peer cloud that provides both compute and data storage through a peer-to-peer network rather than from a centralized set of data centers. The company Hive (www.hivenet.com) proposes to exploit the unused capacity of computers and to incentivize users to contribute their computer resources to the network in exchange for similar capacity from the network and/or monetary compensation. By exchanging similar computing resources and network capacity, users can benefit from all cloud services while ensuring the confidentiality of their data as it is fragmented, encrypted and spread across the peer-to-peer network.

The Inria COAST, COATI, MAGELLAN, PESTO and WIDE teams participating in this challenge bring their expertise on aspects of reliable and cost-efficient data placement and repair in the case of node failures, collaboration on shared data, data security and management of malicious nodes in the context of unreliable distributed storage.

Inria Challenge Project Cupseli

Participants: Davide Frey, Yerom David Bromberg.

The Cupseli challenge aims to demonstrate that it is possible to run complex applications (particularly in the field of machine learning) on heterogeneous, distributed, and volatile resources, while achieving strong parallel efficiency and preserving both accuracy and confidentiality. Building on the combined expertise of hive and Inria in storage technologies illustrated in Alvearium (www.inria.fr/en/alvearium), this strategic partnership explores algorithmic and system solutions to optimize computation, memory, and communications, while ensuring security and fault tolerance. The work is organized around three axes: Frugality (adapting training and inference to limited and dynamic resources), Security and Confidentiality

(protecting data and models through encryption, secure enclaves, and defenses against attacks), and Volatility (ensuring robustness and performance despite the unpredictable arrival and departure of resources). The shared goal is to offer a green and sovereign alternative to data centers, by leveraging already-existing resources for the benefit of AI and Big Data applications. This collaboration, anchored in this joint challenge, brings together researchers, PhD students, engineers, and postdocs, with large-scale experiments conducted on hive's infrastructure.

Inria Challenge Project OS

Participants: Yerom David Bromberg, Barbe Mvondo Djob.

Data centers are today at the heart of all computing, from providing the computing power that supports machine learning, databases, video streaming, etc., down to providing tiny sensors with extra computing power and storage. By centralizing computing, data centers have the potential to deliver massive computing resources while adapting the resource consumption efficiently to changing needs. Nevertheless, data centers have not fully realized their potential of optimizing large-scale computing usage. Instead, studies have consistently shown that, even though new data centers continue to be built, existing data centers are massively underused, typically reaching a usage ratio of only 50

The essential problem of managing a data center is to allocate hardware resources, in an environment in which application requirements are not known a priori and are constantly changing, and where at the same time hardware capabilities are regularly evolving. The Defi OS will attack the problem of data center underusage at the operating system level and hypervisor level, as these are the software components that interact directly with the hardware. The project OS (project.inria.fr/defios/) brings together researchers from the Whisper, WIDE, KrakOS, and Benagil teams and will investigate how virtual machine migration, heterogeneous architectures, rack scale computing, and custom resource management policies can be harnessed to raise the data center usage ratio toward 90

11 Dissemination

11.1 Promoting scientific activities

Participants: Davide Frey, George Giakkoupis, Achour Mostefaoui, Barbe Mvondo Djob, François Taïani.

11.1.1 Scientific events: organisation

Member of the organizing committees

- Davide Frey is a member of the Steering Committee of the PaPoC workshop series.

11.1.2 Scientific events: selection

Chair of conference program committees

- Davide Frey served as PC co-chair the 12th edition of the PaPoC workshop (PaPoC 2025) [35]
- Achour Mostefaoui will serve as the PC co-chair of the ApPLIED@PODC'26 Workshop on Advanced tools, programming languages, and PLatforms for Implementing and Evaluating algorithms for Distributed systems (ApPLIED 2025).

Member of the conference program committees

- François Taïani served on the PC of the 55th Annual IEEE/IFIP International Conference on Dependable Systems and Networks 2025 (DSN 2025).
- François Taïani served on the PC of the 20th European Dependable Computing Conference 2026, (EDCC 2026).
- François Taïani served on the PC of 29th Conference on Principles of Distributed Systems (OPODIS 2025).
- Barbe Mvondo Djob served on the PC of 26th ACM/IFIP International Middleware Conference (Middleware 2025)
- Barbe Mvondo Djob served on the PC of 44th International Symposium on Reliable Distributed Systems (SRDS 2025)
- Barbe Mvondo Djob served on the PC of 20th edition of ACM European Conference on Computer Systems (EUROSYS 2025)
- Achour Mostefaoui served on the PC of ACM PODC, IEEE ICDCS, Opodis, Euro-Par international conferences in 2025.
- Achour Mostefaoui is serving on the PC of DSN 2026.
- Davide Frey served on the PC of the 45th IEEE International Conference on Distributed Computing Systems (ICDCS 2025).
- Davide Frey served on the PC of the 25th International Conference on Distributed Applications and Interoperable Systems (DAIS 2025).
- Davide Frey served on the PC of the 44th International Symposium on Reliable Distributed Systems (SRDS 2025).
- Davide Frey served on the PC of the 26th ACM/IFIP International Middleware Conference (Middleware 2025).
- George Giakkoupis served on the PC of the 52nd EATCS International Colloquium on Automata, Languages, and Programming (ICALP 2025).
- George Giakkoupis served on the PC of the 33rd International Colloquium On Structural Information and Communication Complexity (SIROCCO 2026).

11.1.3 Journal

Reviewer - reviewing activities

- Davide Frey was a reviewer for the Parallel Computing Journal
- George Giakkoupis was a reviewer for Israel Journal of Mathematics (IJMATH), and Autonomous Agents and Multi-Agent Systems Journal (AAMSFJ)

11.1.4 Invited talks

- George Giakkoupis. *Distributed Stochastic Graph Algorithms*. 3rd Bertinoro Workshop on Distributed Geometric Algorithms (DiG@BiCi25), Bertinoro, Italy, Sep. 30 2025
- George Giakkoupis. *Naively sorting evolving data*. ADYN Seminar: Algorithms, Dynamics, and Information Flow in Networks, Virtual, Germany, Mar. 31 2025
- Davide Frey gave a talk at the Workshop on "Distributed Computing: Past, Present, Future" in honor of Maurice Herlihy at LIP6, Paris.

11.1.5 Leadership within the scientific community

- François Taïani serves as co-chair of the scientific committee of GDR RSD (Groupement de Recherche Réseaux & Systèmes Distribués) since 2024.
- George Giakkoupis was a member of the steering committee (member-at-large) of the ACM Symposium on Principles of Distributed Computing (PODC) from 2023-2025.

11.1.6 Scientific expertise

- Achour Mostefaoui served as a member of the evaluation committee CE25 of the ANR (French research funding agency) on Software platforms, systems and networks from 2022 to 2025.
- Davide Frey was an expert reviewer for the National Science Center, Poland.
- George Giakkoupis is a member of the Working Group GT CoA: Complexité et Algorithmes since 2023.

11.1.7 Research administration

- François Taïani served as a member of the thesis committee of the Doctoral School Matisse (ED N. 601).
- François Taïani served as a member of the Bureau du Comité des Projets (BCP) of the Inria Centre at Rennes University.
- François Taïani served as a member of the local Inria secondment committee (Commission des délégations Inria) in Rennes.
- François Taïani served as Career Advice Person, (Référént conseil-parcours professionnel chercheurs) for IRISA/Inria Centre at Rennes University since 2019.
- François Taïani served as vice-chairman of the recruitment committee for a Professorship on "Cybersécurité, réseaux, systèmes" of ESIR (University of Rennes).
- François Taïani served on the recruitment committee for an MCF "Systèmes, IA, HPC" of INSA Lyon.
- François Taïani served on the recruitment committee for a Professorship "Computer Science" of Nantes University.
- Davide Frey served on the recruitment committee for Maître de Conférence at IDMC, Université de Lorraine, Nancy.
- George Giakkoupis is a local correspondent of the Inria Centre at Rennes University for the preparation of the Annual Activity Reports by the project teams since 2023.

11.2 Teaching - Supervision - Juries - Educational and pedagogical outreach

Participants: Yerom David Bromberg, Davide Frey, George Giakkoupis, Achour Mostefaoui, Barbe Mvondo Djob, François Taïani, Brice Ekane Apah.

11.2.1 Teaching

- Engineering School: François Taïani, Operating Systems, 28h, 2nd year of Engineering School (M1), ESIR / U. Rennes, France.
- Engineering School: François Taïani, Distributed Systems, 12h, 3rd year of Engineering School (M2), ESIR / U. Rennes, France.
- Engineering School: François Taïani, Introduction to Operating Systems, 72h, 1st year of Engineering School (L3), ESIR / U. Rennes, France.
- Engineering School: Barbe Mvondo Djob, Network and Security for IOT, 45h, ESIR M1, Rennes, France
- Engineering School: Barbe Mvondo Djob, Cloud Computing for IOT, 45h, ESIR M2, Rennes, France
- Engineering School: Barbe Mvondo Djob, Cybersecurity and Hacking, 30h, Polytechnic 2nd year, Paris, France
- Computer Science Dpt.: Achour Mostefaoui, Graph Theory, 42h, 3rd year of Bachelor of Science (L3), ISTIC / U. Rennes, France.
- Master: Davide Frey, Scalable Distributed Systems, 10h, M1, EIT/ICT Labs Master School, U. Rennes, France.
- ENS L3 : Davide Frey, Distributed Algorithms, 11h, L3 parcours SI, ISTIC, ENS Rennes, France.
- ENS L3: George Giakkoupis, Distributed Algorithms, 9h, L3 parcours SI, ISTIC, ENS Rennes, France.
- Master: Davide Frey, Cloud Computing, 12h, M2-MIAGE, U. Rennes, France.
- Computer Science Dpt: Brice Ekane Apah, Networking, 45h, L2, U. Rennes, France.
- Computer Science Dpt: Brice Ekane Apah, Enterprise network architectures, 36h, M2, U. Rennes, France.
- Computer Science Dpt: Brice Ekane Apah, Next-Generation Network Architecture, 23h, M2, U. Rennes, France.
- Computer Science Dpt: Brice Ekane Apah, Software Techniques for Cloud Computing, 33h, M2, U. Rennes, France.
- Computer Science Dpt: Brice Ekane Apah, Learning Progress Monitoring, 13.5h, M2, U. Rennes, France.

11.2.2 Supervision

- PhD (defended in December 2025): Vincent Kowalski, Useful byzantine abstraction at computability level of shared memory, supervised by Achour Mostefaoui and Matthieu Perrin (Univ. Nantes)
- PhD (defended in December 2024): Timothé Albouy, Towards Lightweight Scalable and Open Byzantine-Fault-Tolerant Distributed Objects, U. Rennes, supervised by François Taïani and Davide Frey.
- PhD in progress: Dimitri Lerévérénd, Privacy-Preserving Decentralized Learning Through Model Fragmentation and Private Aggregation, started in September 2023, supervised by Davide Frey, Romaric Gaudel (MALT team) and François Taïani.
- PhD in progress: Manon Sourisseau, Byzantine-Tolerant Netcodes For Tomorrow's Metaverse, started in October 2023, supervised by François Taïani, Yerom David Bromberg, and Jérémie Découchant (TU Delft).

- PhD in progress: Rémy Raes, Distributed Machine Learning in Ubiquitous Environments using Location-dependent Models, started in 2023, supervised by Davide Frey, François Taïani, Romain Rouvoy and Lionel Seinturier (Spirals team, Inria Lille).
- PhD in progress: Augustin Godinot, Auditing the mutations of AI-models, started in November 2022, supervised by Erwan Le Merrer, Gilles Trédan (LAAS/CNRS), François Taïani, and Camilla Penzo (PEReN).
- PhD in progress: Hua Junrui, Advanced Techniques for Efficient Distributed Hash Tables Management with Fault Tolerance against Byzantine Faults in Large-Scale Distributed Systems, started in November 2024, supervised by François Taïani, Gérald Oster (LORELEY team, Inria Nancy), and Alexandru Dobrila (Hive).
- PhD in progress: Ludovic Paillat, Security for peer-to-peer cloud storage without central authority, started in 2023, supervised by Davide Frey, Claudia Ignat (LORELEY team, Inria Nancy), Alexandru Dobrila (HIVE), Mathieu Turiani (PESTO Team, Inria Nancy).
- PhD in progress: Adrien Gegout, Efficient containerized Cloud Gaming, started in October 2023, supervised by Davide Frey, Barbe Mvondo Djob, Pascal Manchon (Blacknut).
- PhD in progress: Cesaire Honoré, Scheduling in heterogeneous architectures, started in December 2022, supervised by Yerom David Bromberg and Barbe Mvondo Djob
- PhD in progress: Victoire Nganfng, Resisting to Massive proliferation of new Android malware threats, started in November 2024, supervised by Yerom Yerom David Bromberg and Valério Schiavoni (University of Neuchâtel, Switzerland)
- PhD in progress: Stella Tchoutcha, Energy-Efficient Function-as-a-Service (FaaS) for Low-Power Edge and IoT Deployments, started in December 2025, supervised by Barbe Mvondo Djob and Nikos Parlavantzas (Magellan team)
- PhD in progress: Caleb Fonyuy-Asheri, Heterogeneous VM migration , started in February 2024, supervised by Yerom David Bromberg, Alain Tchana (Grenoble INP), Barbe Mvondo Djob, Renaud Lachaise (UGA)
- PhD in progress: Alexandre Duvivier, CDN performance optimization, started in October 2023, supervised by Yerom David Bromberg, Barbe Mvondo Djob, Nicolas Le Scouarnec (Broadpeak)
- PhD in progress: Amelie Gonzalez, Linux network stack optimization, started in September 2023, supervised by Yerom David Bromberg, Barbe Mvondo Djob, Julia Lawal (Whisper team, Inria Paris)
- PhD in progress : Hugo Bertin, Attacking Games to Strengthen their Defenses, started in February 2025, Supervised by Yerom David Bromberg and Marc Dacier (KAUST)
- PhD in progress : Elie Raspaud, Enhancing privacy in distributed machine learning using homomorphic cryptography, started in September 2025, supervised by Davide Frey, Philippe Chartier and Mohammed Lemou (CNRS)
- PostDoc: Georgy Ishmaev, Byzantine Fault-Tolerant Distributed Ledgers, May 2024-October 2025, supervised by François Taïani and Davide Frey, funded by ANR project ByBloS n. ANR-20-CE25-0002.
- PostDoc: Dimitrios Los, Distributed Algorithms on Evolving Data, until Feb 2024, supervised by George Giakkoupis, funded by Action Exploratoire DisEvo: Distributed Algorithms on Evolving Data

11.2.3 Juries

- François Taïani was a reviewer for Yacine Belal's PhD thesis: Trustworthy Collaborative Learning: Personalization, Privacy, and Robustness at the Edge, INSA Lyon (France), 10 June 2025.
- François Taïani was a reviewer for Sara Tucci-Piergiovanni's HDR thesis: Blockchains: from The Wild to Distributed Computing, Université Paris-Saclay (France), 29 September 2025.

- Achour Mostefaoui was a reviewer for Alejandro Naser’s PhD thesis: Fault-Tolerant Computing with Unreliable Channels, IMDEA Research Institute (Madrid, Spain), 10 December 2025.
- Davide Frey was a reviewer for Eugenio Lomurno’s PhD thesis: Adversarial and Generative Deep Learning for Data Privacy in Human-Centered Artificial Intelligence, Politecnico di Milano (Milano, Italy), 9 May 2025.
- Davide Frey was a reviewer for Hassan Nazeer Chaudhry’s PhD thesis: Efficient Processing of Graph-Based Data Streams, Politecnico di Milano (Milano, Italy), 29 August, 2025.
- Davide Frey was a reviewer for Bulat Aydarovich NASRULIN’s PhD thesis: Trustworthy Foundations for Web3, Delft University of Technology (Delft, The Netherlands), 26 September 2025.
- Brice Ekane Apah was an examiner for PAPA Asane FALL PhD thesis: Linux as a micro-kernel : the memory management’s case, University of Grenoble Alpes (Grenoble, France), 05 December 2025.
- Barbe Mvondo Djob served as one of the French Baccalureate jury presidents (président du jury du baccalauréats) for 2025.

11.3 Popularization

Participants: Timothé Albouy, Davide Frey, Manon Sourisseau, François Taïani.

11.3.1 Productions (articles, videos, podcasts, serious games, ...)

- François Taïani, Manon Sourisseau, Timothé Albouy, and Davide Frey took part in the filming of the television programme *Esprit Sorcier: les surprises de la recherche* (Wizard Spirit: the surprises of research), filmed in Rennes in April 2025 with the participation of 340 secondary school pupils from the Rennes region. They presented their current research on Byzantine fault-tolerant protocols as part of the ByBloS project. The programme is now available on Internet streaming services and [YouTube](#).

12 Scientific production

12.1 Major publications

- [1] T. Albouy, D. Frey, M. Raynal and F. Taïani. ‘Good-case Early-Stopping Latency of Synchronous Byzantine Reliable Broadcast: The Deterministic Case’. In: DISC 2022 - 36th International Symposium on Distributed Computing. Augusta, GA, United States, 25th Oct. 2022. doi: [10.4230/LIPIcs.DISC.2022.4](https://doi.org/10.4230/LIPIcs.DISC.2022.4). URL: <https://inria.hal.science/hal-03791921>.
- [2] A. Auvolet, D. Frey, M. Raynal and F. Taïani. ‘Byzantine-Tolerant Causal Broadcast’. In: *Theoretical Computer Science* 885 (Sept. 2021), pp. 55–68. doi: [10.1016/j.tcs.2021.06.021](https://doi.org/10.1016/j.tcs.2021.06.021). URL: <https://hal.inria.fr/hal-03346710>.
- [3] D. Bosk, D. Frey, M. Gestin and G. Piolle. ‘Hidden Issuer Anonymous Credential’. In: *Proceedings on Privacy Enhancing Technologies* 2022 (June 2022), pp. 571–607. doi: [10.56553/popets-2022-0123](https://doi.org/10.56553/popets-2022-0123). URL: <https://hal.archives-ouvertes.fr/hal-03789485>.
- [4] Y.-D. Bromberg, Q. Dufour and D. Frey. ‘Multisource Rumor Spreading with Network Coding’. In: *INFOCOM 2019 - IEEE International Conference on Computer Communications*. Paris, France: IEEE, Apr. 2019, pp. 1–10. URL: <https://hal.inria.fr/hal-01946632>.
- [5] Y.-D. Bromberg, Q. Dufour, D. Frey and E. Rivière. ‘Donar: Anonymous VoIP over Tor’. In: NSDI 2022 - 19th USENIX Symposium on Networked Systems Design and Implementation. RENTON, WA, United States, 4th Apr. 2022. URL: <https://hal.inria.fr/hal-03923695>.

- [6] G. Damaskinos, R. Guerraoui, A.-M. Kermarrec, V. Nitu, R. Patra and F. Taïani. ‘FLeet: On-line Federated Learning via Staleness Awareness and Performance Prediction’. In: *Middleware ’20: Proceedings of the 21st International Middleware Conference*. 21st International Middleware Conference. Delft (virtual), Netherlands, 7th Dec. 2020. DOI: [10.1145/3423211.3425685](https://doi.org/10.1145/3423211.3425685). URL: <https://hal.archives-ouvertes.fr/hal-03390450>.
- [7] D. Frey, M. Gestin and M. Raynal. ‘The Synchronization Power (Consensus Number) of Access-Control Objects: the Case of AllowList and DenyList’. In: *DISC 2023 - 37th International Symposium on Distributed Computing*. L’aquila, Italy: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, pp. 1–32. DOI: [10.4230/LIPIcs.DISC.2023.21](https://doi.org/10.4230/LIPIcs.DISC.2023.21). URL: <https://inria.hal.science/hal-04399298>.
- [8] G. Giakkoupis. ‘Expanders via local edge flips in quasilinear time’. In: *STOC 2022 - 54th Annual ACM SIGACT Symposium on Theory of Computing*. Rome, Italy: ACM, 25th May 2022, pp. 64–76. DOI: [10.1145/3519935.3520022](https://doi.org/10.1145/3519935.3520022). URL: <https://hal.inria.fr/hal-03792482>.
- [9] G. Giakkoupis, M. Jafari Giv and P. Woelfel. ‘Efficient Randomized DCAS’. In: *STOC 2021 - 53rd Annual ACM SIGACT Symposium on Theory of Computing*. Rome (Virtual), Italy: ACM, 21st June 2021, pp. 1–64. DOI: [10.1145/3406325.3451133](https://doi.org/10.1145/3406325.3451133). URL: <https://hal.inria.fr/hal-03195692>.
- [10] G. Giakkoupis, M. Kiwi and D. Los. ‘Naively Sorting Evolving Data is Optimal and Robust’. In: *FOCS 2024 - IEEE 65th Annual Symposium on Foundations of Computer Science*. Chicago, United States: IEEE Computer Society, 2024, pp. 2217–2242. DOI: [10.1109/FOCS61266.2024.00130](https://doi.org/10.1109/FOCS61266.2024.00130). URL: <https://hal.science/hal-04888959>.
- [11] R. Guerraoui, A.-M. Kermarrec, G. Niot, O. Ruas and F. Taïani. ‘GoldFinger: Fast & Approximate Jaccard for Efficient KNN Graph Constructions’. In: *IEEE Transactions on Knowledge and Data Engineering* 35.11 (1st Nov. 2023), pp. 11461–11475. DOI: [10.1109/TKDE.2022.3232689](https://doi.org/10.1109/TKDE.2022.3232689). URL: <https://inria.hal.science/hal-04394851>.
- [12] R. Guerraoui, A.-M. Kermarrec, O. Ruas and F. Taïani. ‘Smaller, Faster & Lighter KNN Graph Constructions’. In: *WWW ’20 - The Web Conference 2020*. Taipei Taiwan, France: ACM, 20th Apr. 2020, pp. 1060–1070. DOI: [10.1145/3366423.3380184](https://doi.org/10.1145/3366423.3380184). URL: <https://hal.inria.fr/hal-02888286>.
- [13] H. Lakhlef, M. Raynal and F. Taïani. ‘Vertex Coloring with Communication Constraints in Synchronous Broadcast Networks’. In: *IEEE Transactions on Parallel and Distributed Systems* 30.7 (July 2019), pp. 1672–1686. DOI: [10.1109/TPDS.2018.2889688](https://doi.org/10.1109/TPDS.2018.2889688). URL: <https://hal.inria.fr/hal-02376726>.
- [14] T. Maho, T. Furon and E. L. Merrer. ‘SurFree: a fast surrogate-free black-box attack’. In: *CVPR 2021 - Conference on Computer Vision and Pattern Recognition*. Proc. of IEEE Conference on Computer Vision and Pattern Recognition, CVPR. Virtual, France, 19th June 2021, pp. 10430–10439. URL: <https://hal.archives-ouvertes.fr/hal-03177639>.

12.2 Publications of the year

International journals

- [15] L. Devroye and D. Los. ‘An asymptotically optimal algorithm for generating bin cardinalities’. In: *Mathematics and Computers in Simulation* 228 (Feb. 2025), pp. 147–155. DOI: [10.1016/j.matcom.2024.08.034](https://doi.org/10.1016/j.matcom.2024.08.034). URL: <https://hal.science/hal-04889571> (cit. on p. 23).
- [16] R. Duvignau, M. Raynal and E. M. Schiller. ‘Self-stabilizing multivalued consensus in the presence of Byzantine faults and asynchrony’. In: *Theoretical Computer Science* 1039 (June 2025), p. 115184. DOI: [10.1016/J.TCS.2025.115184](https://doi.org/10.1016/J.TCS.2025.115184). URL: <https://inria.hal.science/hal-05211139>.
- [17] G. Giakkoupis, V. Turau and I. Ziccardi. ‘Luby’s MIS Algorithms Made Self-Stabilizing’. In: *Information Processing Letters* 188 (11th Aug. 2025), p. 106531. DOI: [10.1016/j.ipl.2024.106531](https://doi.org/10.1016/j.ipl.2024.106531). URL: <https://inria.hal.science/hal-05207077> (cit. on p. 20).

- [18] G. Ishmaev, E. Anceaume, D. Frey and F. Taïani. ‘Ethical Risk Analysis of L2 Rollups’. In: *ACM Transactions on the Web* (25th Dec. 2025). doi: [10.1145/3786147](https://doi.org/10.1145/3786147). URL: <https://cnrs.hal.science/hal-05446710> (cit. on p. 19).
- [19] L. Paillat, C.-L. Ignat, D. Frey, M. Turuani and A. Ismail. ‘Discreet: distributed delivery service with context-aware cooperation’. In: *Annals of Telecommunications - annales des télécommunications* 80.3-4 (Apr. 2025), pp. 357–374. doi: [10.1007/s12243-024-01053-1](https://doi.org/10.1007/s12243-024-01053-1). URL: <https://inria.hal.science/hal-04829916> (cit. on p. 20).

International peer-reviewed conferences

- [20] T. Albouy, A. Fernández Anta, C. Georgiou, M. Gestin, N. Nicolaou and J. Wang. ‘AMECOS: A Modular Event-based Framework for Concurrent Object Specification’. In: *OPODIS 2024 - 28th International Conference on Principles of Distributed Systems*. Vol. Proceedings of the 28th International Conference on Principles of Distributed Systems (OPODIS 2024). Lucques, Italy: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi: [10.4230/LIPIcs.OPODIS.2024.4](https://doi.org/10.4230/LIPIcs.OPODIS.2024.4). URL: <https://inria.hal.science/hal-04577664>.
- [21] T. Albouy, D. Frey, M. Gestin, M. Raynal and F. Taïani. ‘Contention-Aware Cooperation’. In: *OPODIS 2025 - 29th International Conference On Principles Of Distributed Systems*. 9. Iasi, Romania: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 3rd Dec. 2025, 9:1–9:20. doi: [10.4230/LIPIcs.OPODIS.2025.9](https://doi.org/10.4230/LIPIcs.OPODIS.2025.9). URL: <https://inria.hal.science/hal-05456351> (cit. on p. 18).
- [22] H. Bertin, I. Benhabbour, M. Dacier and Y.-D. Bromberg. ‘Disconnecting Users from Virtual Worlds with a Single Packet: an Unreal Untold Story’. In: *iMETA 2025 - 3rd International Conference on Intelligent Metaverse Technologies & Applications*. Dubrovnik, Croatia, 14th Oct. 2025, pp. 67–74. doi: [10.1109/iMETA66706.2025.11306743](https://doi.org/10.1109/iMETA66706.2025.11306743). URL: <https://hal.science/hal-05202617> (cit. on p. 22).
- [23] S. Biswas, D. Frey, R. Gaudel, A.-M. Kermarrec, D. Lerévérénd, R. Pires, R. Sharma and F. Taïani. ‘Low-Cost Privacy-Preserving Decentralized Learning’. In: *Proceedings on Privacy Enhancing Technologies Symposium*. Privacy Enhancing Technologies Symposium. Vol. 2025. 3. Washington DC, United States, July 2025, pp. 451–474. doi: [10.56553/popets-2025-0108](https://doi.org/10.56553/popets-2025-0108). URL: <https://hal.science/hal-04993586> (cit. on p. 22).
- [24] Y.-D. Bromberg, J. Decouchant, M. Sourisseau and F. Taïani. ‘Formalizing Rollback Netcodes for Robust and Real-Time Client-Server Architectures’. In: *Proceedings of the 29th International Conference on Principles of Distributed Systems*. OPODIS 2025 - 29th International Conference on Principles of Distributed Systems. Iasi, Romania: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2026, pp. 1–17. doi: [10.4230/LIPIcs.OPODIS.2025.11](https://doi.org/10.4230/LIPIcs.OPODIS.2025.11). URL: <https://inria.hal.science/hal-05454044>.
- [25] F. d’Amore, N. d’Archivio, G. Giakkoupis and E. Natale. ‘On the h-majority dynamics with many opinions’. In: *39th International Symposium on Distributed Computing (DISC 2025)*. Berlin, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi: [10.4230/LIPIcs.DISC.2025.27](https://doi.org/10.4230/LIPIcs.DISC.2025.27). URL: <https://hal.science/hal-05198643> (cit. on p. 20).
- [26] B. Ekane, D. Mvondo, R. Lachaize, Y.-D. Bromberg, A. Tchana and D. Hagimont. ‘DISC: Backpressure Mitigation In Multi-tier Applications With Distributed Shared Connection’. In: *Proceedings of the 22nd USENIX Symposium on Networked Systems Design and Implementation (NSDI 25)*. NSDI 2025 - 22nd USENIX Symposium on Networked Systems Design and Implementation. Philadelphia (Pennsylvania), United States, 28th Apr. 2025, pp. 55–70. URL: <https://hal.science/hal-05051628> (cit. on p. 21).
- [27] E. Gafni, G. Losa, M. Raynal and G. Taubenfeld. ‘Brief Announcement: Stranger-Free Tasks’. In: *PODC 2025 - 44th ACM Symposium on Principles of Distributed Computing*. Mexico, Mexico: ACM, 16th June 2025, pp. 203–206. doi: [10.1145/3732772.3733545](https://doi.org/10.1145/3732772.3733545). URL: <https://inria.hal.science/hal-05211143>.

- [28] A. Gegout, D. Mvondo, D. Frey and P. Monchon. ‘Towards an Efficient Containerized Cloud Gaming Platform’. In: *ASHES - IPDPS Workshops*. Vol. 2025 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW). Milano, France, 3rd June 2025, pp. 78–86. DOI: [10.1109/IPDPSW66978.2025.00018](https://doi.org/10.1109/IPDPSW66978.2025.00018). URL: <https://hal.science/hal-05094575> (cit. on p. 22).
- [29] A. Godinot, E. L. Merrer, C. Penzo, F. Taïani and G. Tredan. ‘Queries, Representation & Detection: The Next 100 Model Fingerprinting Schemes’. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. AAAI 2025 - 39th Annual AAAI Conference on Artificial Intelligence. Vol. 39. 16. Philadelphia (Pennsylvania), United States, 11th Apr. 2025, pp. 16817–16825. DOI: [10.1609/aaai.v39i16.33848](https://doi.org/10.1609/aaai.v39i16.33848). URL: <https://inria.hal.science/hal-05093903>.
- [30] H. C. Mounah, D. Mvondo, J. Lawall and Y.-D. Bromberg. ‘The Impact of Kernel Asynchronous APIs on the Performance of a Kernel VPN’. In: *SYSTOR 2025 - 18th ACM International System and Storage Conference*. Virtual conference, Israel, 2025, pp. 167–173. DOI: [10.1145/3757347.3759133](https://doi.org/10.1145/3757347.3759133). URL: <https://hal.science/hal-05211974> (cit. on p. 21).
- [31] D. Mvondo, B. Djongwe Teabe and N. Parlavantzas. ‘Efficient Memory Usage For Edge FaaS Platforms’. In: *PerCom 2026 - 24th IEEE International Conference on Pervasive Computing and Communications*. Pise, Italy, 2026. URL: <https://inria.hal.science/hal-05465369>.
- [32] D. B. T. Mvondo Djob and Y.-D. Bromberg. ‘Secure access to network data for mobile network traffic analysis applications’. In: *55th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2025)*. Naples, Italy, 23rd June 2025. DOI: [10.1109/dsn64029.2025.00063](https://doi.org/10.1109/dsn64029.2025.00063). URL: <https://hal.science/hal-05088761> (cit. on p. 21).
- [33] V. Nganfang, S. Queyrut, D. Bromberg, V. Schiavoni, D. Mvondo and V. Kengne Tchendji. ‘DroidHunter: A Robust Vision-Based Detection Against Hidden Android Malware’. In: *ASIACCS 2026 - 21st ACM ASIA Conference on Computer and Communications Security*. Bangalore, India: ACM, 1st June 2026. DOI: [10.1145/3779208.3785386](https://doi.org/10.1145/3779208.3785386). URL: <https://hal.science/hal-05419946> (cit. on p. 22).
- [34] P. de Rosa, S. Queyrut, Y.-D. Bromberg, P. Felber and V. Schiavoni. ‘PhishingHook: Catching Phishing Ethereum Smart Contracts leveraging EVM Opcodes’. In: *DSN 2025 - 55th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. Naples, France: IEEE, 2025, pp. 265–266. DOI: [10.1109/DSN-S65789.2025.00075](https://doi.org/10.1109/DSN-S65789.2025.00075). URL: <https://inria.hal.science/hal-05211704>.

Edition (books, proceedings, special issue of a journal)

- [35] *PaPoC '25: 12th Workshop on Principles and Practice of Consistency for Distributed Data*. PaPoC '25: 12th Workshop on Principles and Practice of Consistency for Distributed Data. World Trade Center Rotterdam Netherlands, France: ACM, 31st Mar. 2025. DOI: [10.1145/3721473](https://doi.org/10.1145/3721473). URL: <https://inria.hal.science/hal-05474249> (cit. on p. 28).

Reports & preprints

- [36] T. Albouy, E. Anceaume, D. Frey, M. Gestin, A. Rauch, M. Raynal and F. Taïani. *Asynchronous BFT Asset Transfer: Quasi-Anonymous, Light, and Consensus-Free*. 18th Feb. 2025. URL: <https://inria.hal.science/hal-04578985>.
- [37] A. Bellet, E. Cyffers, D. Frey, R. Gaudel, D. Lérévérend and F. Taïani. *Unified Privacy Guarantees for Decentralized Learning via Matrix Factorization*. 20th Oct. 2025. URL: <https://hal.science/hal-05462250>.
- [38] T. Chauvin, E. Le Merrer, F. Taïani and G. Tredan. *Log Probability Tracking of LLM APIs*. 3rd Dec. 2025. URL: <https://hal.science/hal-05421014>.
- [39] G. Ishmaev. *Ethics of Blockchain Technologies*. 2025. DOI: [10.48550/arXiv.2504.02504](https://doi.org/10.48550/arXiv.2504.02504). URL: <https://inria.hal.science/hal-04993192>.
- [40] G. Ishmaev, E. Anceaume, D. Frey and F. Taïani. *Ethical Risk Analysis of L2 Rollups*. 5th Dec. 2025. URL: <https://inria.hal.science/hal-05404607>.

12.3 Cited publications

- [41] Y. Afek and E. Gafni. ‘Asynchrony from synchrony’. In: *ICDCN*. 2013, pp. 225–239 (cit. on p. 14).
- [42] A. Ahmed and E. Ahmed. ‘A survey on mobile edge computing’. In: *2016 10th International Conference on Intelligent Systems and Control (ISCO)*. Jan. 2016, pp. 1–8. doi: [10.1109/ISCO.2016.7727082](https://doi.org/10.1109/ISCO.2016.7727082). URL: <http://dx.doi.org/10.1109/ISCO.2016.7727082> (cit. on p. 9).
- [43] T. Allard, D. Frey, G. Giakkoupis and J. Lepiller. ‘Lightweight Privacy-Preserving Averaging for the Internet of Things’. In: *MAIOT 2016 - 3rd Workshop on Middleware for Context-Aware Applications in the IoT*. Trento, Italy: ACM, Dec. 2016, pp. 19–22. doi: [10.1145/3008631.3008635](https://doi.org/10.1145/3008631.3008635). URL: <https://hal.inria.fr/hal-01421986> (cit. on pp. 11, 14).
- [44] Z. Allen-Zhu, A. Bhaskara, S. Lattanzi, V. Mirrokni and L. Orecchia. ‘Expanders via local edge flips’. In: *Proceedings of the twenty-seventh annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2016, pp. 259–269 (cit. on p. 15).
- [45] E. Anshelevich, D. Chakrabarty, A. Hate and C. Swamy. ‘Approximability of the Firefighter Problem: Computing Cuts over Time’. In: *Algorithmica* 62.1-2 (2012), pp. 520–536 (cit. on p. 12).
- [46] D. Bernstein. ‘Containers and Cloud: From LXC to Docker to Kubernetes’. In: *IEEE Cloud Computing* 1.3 (Sept. 2014), pp. 81–84. doi: [10.1109/MCC.2014.51](https://doi.org/10.1109/MCC.2014.51). URL: <http://dx.doi.org/10.1109/MCC.2014.51> (cit. on p. 9).
- [47] M. Bertier, D. Frey, R. Guerraoui, A.-M. Kermarrec and V. Leroy. ‘The Gossple Anonymous Social Network’. In: *ACM/IFIP/USENIX 11th International Middleware Conference (MIDDLEWARE)*. Ed. by I. Gupta and C. Mascolo. Vol. LNCS-6452. Middleware 2010. Bangalore, India: Springer, Nov. 2010, pp. 191–211. doi: [10.1007/978-3-642-16955-7_10](https://doi.org/10.1007/978-3-642-16955-7_10). URL: <https://hal.inria.fr/inria-00515693> (cit. on p. 9).
- [48] F. Bonomi. *Connected vehicles, the internet of things, and fog computing*. VANET 2011, 2011. Keynote speech at VANET. 2011 (cit. on p. 9).
- [49] F. Bonomi, R. Milito, J. Zhu and S. Addepalli. ‘Fog Computing and Its Role in the Internet of Things’. In: *1st MCC Workshop on Mobile Cloud Computing*. 2012. doi: [10.1145/2342509.2342513](https://doi.org/10.1145/2342509.2342513). URL: <http://doi.acm.org/10.1145/2342509.2342513> (cit. on p. 9).
- [50] A. Boutet, D. Frey, R. Guerraoui, A. Jégou and A.-M. Kermarrec. ‘Privacy-Preserving Distributed Collaborative Filtering’. In: *Computing*. Special Issue on NETYS 2014 98.8 (Aug. 2016). URL: <https://hal.inria.fr/hal-01251314> (cit. on p. 11).
- [51] A. Boutet, D. Frey, R. Guerraoui, A.-M. Kermarrec and R. Patra. ‘HyRec: Leveraging Browsers for Scalable Recommenders’. In: *Middleware 2014*. Bordeaux, France, Dec. 2014. doi: [10.1145/2663165.2663315](https://doi.org/10.1145/2663165.2663315). URL: <https://hal.inria.fr/hal-01080016> (cit. on p. 9).
- [52] A. Boutet, D. Frey, R. Guerraoui, A.-M. Kermarrec, A. Rault, F. Taïani and J. Wang. ‘Hide & Share: Landmark-based Similarity for Private KNN Computation’. In: *DSN*. Rio de Janeiro, Brazil, 2015. doi: [10.1109/DSN.2015.60](https://doi.org/10.1109/DSN.2015.60). URL: <https://hal.archives-ouvertes.fr/hal-01171492> (cit. on p. 11).
- [53] A. Boutet, D. Frey, A. Jégou, A.-M. Kermarrec and H. Ribeiro. ‘FreeRec: an Anonymous and Distributed Personalization Architecture’. In: *Computing* (Dec. 2013). URL: <https://hal.inria.fr/hal-00909127> (cit. on p. 11).
- [54] B. Cohen. *Incentives Build Robustness in BitTorrent*. 2003. URL: <http://citeseer.ist.psu.edu/cohen03incentives.html> (cit. on p. 11).
- [55] D. Collins, R. Guerraoui, J. Komatovic, P. Kuznetsov, M. Monti, M. Pavlovic, Y. Pignolet, D. Seredinschi, A. Tonkikh and A. Xygkis. ‘Online Payments by Merely Broadcasting Messages’. In: *IEEE DSN*. 2020. doi: [10.1109/DSN48063.2020.00023](https://doi.org/10.1109/DSN48063.2020.00023). URL: <https://doi.org/10.1109/DSN48063.2020.00023> (cit. on p. 14).
- [56] C. Delporte-Gallet, H. Fauconnier, R. Guerraoui and A. Tielmann. ‘The disagreement power of an adversary’. In: *Distributed Computing* 24.3-4 (2011), pp. 137–147 (cit. on p. 13).

- [57] A. J. Demers, D. H. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. E. Sturgis, D. C. Swinehart and D. B. Terry. ‘Epidemic Algorithms for Replicated Database Maintenance’. In: *PODC*. 1987, pp. 1–12 (cit. on p. 12).
- [58] D. Frey, R. Guerraoui, A.-M. Kermarrec, M. Monod, K. Boris, M. Martin and V. Quéma. ‘Heterogeneous Gossip’. In: *Middleware 2009*. Urbana-Champaign, IL, United States, Dec. 2009. URL: <https://hal.inria.fr/inria-00436125> (cit. on p. 9).
- [59] W. M. Golab, V. Hadzilacos, D. Hendler and P. Woelfel. ‘RMR-efficient implementations of comparison primitives using read and write operations’. In: *Distributed Computing* 25.2 (2012), pp. 109–162 (cit. on p. 13).
- [60] R. Guerraoui, K. Huguenin, A.-M. Kermarrec, M. Monod and S. Prusty. ‘LiFTinG: Lightweight Freerider-Tracking Protocol in Gossip’. In: *11th ACM/IFIP/USENIX International Middleware Conference (MIDDLEWARE)*. Bangalore, India, Nov. 2010. DOI: [10.1007/978-3-642-16955-7_16](https://doi.org/10.1007/978-3-642-16955-7_16). URL: <https://hal.inria.fr/inria-00505268> (cit. on p. 11).
- [61] R. Guerraoui, P. Kuznetsov, M. Monti, M. Pavlovic and D. Seredinschi. ‘The Consensus Number of a Cryptocurrency’. In: *ACM PODC*. 2019. DOI: [10.1145/3293611.3331589](https://doi.org/10.1145/3293611.3331589). URL: <https://doi.org/10.1145/3293611.3331589> (cit. on p. 14).
- [62] R. A. Holley and T. M. Liggett. ‘Ergodic Theorems for Weakly Interacting Infinite Systems and the Voter Model’. In: *The Annals of Probability* 3.4 (1975), pp. 643–663 (cit. on p. 12).
- [63] K. Huang, H. Wei, Y. Huang, H. Li and A. Pan. ‘Byz-GentleRain: An Efficient Byzantine-tolerant Causal Consistency Protocol’. In: *CoRR* abs/2109.14189 (2021). arXiv: [2109.14189](https://arxiv.org/abs/2109.14189). URL: <https://arxiv.org/abs/2109.14189> (cit. on p. 14).
- [64] D. Imbs and M. Raynal. ‘A liveness condition for concurrent objects: x-wait-freedom’. In: *Concurrency and Computation: Practice and experience* 23.17 (2011), pp. 2154–2166 (cit. on p. 13).
- [65] F. Junqueira and K. Marzullo. ‘A framework for the design of dependent-failure algorithms’. In: *Concurrency and Computation: Practice and Experience* 19.17 (2007), pp. 2255–2269 (cit. on p. 13).
- [66] D. Kempe, J. M. Kleinberg and É. Tardos. ‘Influential Nodes in a Diffusion Model for Social Networks’. In: *ICALP*. 2005, pp. 1127–1138 (cit. on p. 12).
- [67] D. Kempe, J. M. Kleinberg and É. Tardos. ‘Maximizing the Spread of Influence through a Social Network’. In: *Theory of Computing* 11 (2015), pp. 105–147 (cit. on p. 12).
- [68] D. Kempe, J. M. Kleinberg and É. Tardos. ‘Maximizing the spread of influence through a social network’. In: *KDD*. 2003, pp. 137–146 (cit. on p. 12).
- [69] P. Kuznetsov et al. ‘Understanding non-uniform failure models’. In: *Bulletin of the EATCS* 106 (2012), pp. 53–77 (cit. on p. 13).
- [70] E. Lieberman, C. Hauert and M. Nowak. ‘Evolutionary dynamics on graphs’. In: *Nature* 433.7023 (2005), pp. 312–316 (cit. on p. 12).
- [71] D. Malkhi and M. Reiter. ‘Byzantine quorum systems’. In: *Distributed computing* 11.4 (1998), pp. 203–213 (cit. on p. 14).
- [72] D. Mvondo, A. Tchana, R. Lachaize, D. Hagimont and N. D. Palma. ‘Fine-Grained Fault Tolerance for Resilient pVM-Based Virtual Machine Monitors’. In: *50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN*. IEEE, 2020, pp. 197–208. DOI: [10.1109/DSN48063.2020.00037](https://doi.org/10.1109/DSN48063.2020.00037). URL: <https://doi.org/10.1109/DSN48063.2020.00037> (cit. on p. 14).
- [73] D. Mvondo, B. Teabe, A. Tchana, D. Hagimont and N. D. Palma. ‘Memory flipping: a threat to NUMA virtual machines in the Cloud’. In: *2019 IEEE Conference on Computer Communications, INFOCOM 2019*. IEEE, 2019, pp. 325–333. DOI: [10.1109/INFOCOM.2019.8737548](https://doi.org/10.1109/INFOCOM.2019.8737548). URL: <https://doi.org/10.1109/INFOCOM.2019.8737548> (cit. on p. 14).
- [74] M. Raynal and J. Stainer. ‘Synchrony weakened by message adversaries vs asynchrony restricted by failure detectors’. In: *PODC*. Proceedings of the 2013 ACM symposium on Principles of distributed computing. Montréal, Canada: ACM, July 2013, pp. 166–175. DOI: [10.1145/2484239.2484249](https://doi.org/10.1145/2484239.2484249). URL: <https://hal.inria.fr/hal-00920734> (cit. on pp. 13, 14).

-
- [75] N. Santoro and P. Widmayer. ‘Time is not a healer’. In: *Annual Symposium on Theoretical Aspects of Computer Science*. Springer. 1989, pp. 304–313 (cit. on p. 13).
- [76] A. Verma, L. Pedrosa, M. Korupolu, D. Oppenheimer, E. Tune and J. Wilkes. ‘Large-scale cluster management at Google with Borg’. In: *Tenth European Conference on Computer Systems (Eurosys 2015)*. ACM. 2015, p. 18 (cit. on p. 9).
- [77] L. Zhang, F. Zhou, A. Mislove and R. Sundaram. ‘Maygh: Building a CDN from Client Web Browsers’. In: *8th ACM European Conference on Computer Systems*. EuroSys ’13. Prague, Czech Republic: ACM, 2013, pp. 281–294. DOI: [10.1145/2465351.2465379](https://doi.org/10.1145/2465351.2465379). URL: <http://doi.acm.org/10.1145/2465351.2465379> (cit. on p. 9).