RESEARCH CENTRE

Inria Centre at Rennes University

IN PARTNERSHIP WITH:

Université de Rennes

2024

ACTIVITY REPORT

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)

DOMAIN

Networks, Systems and Services, Distributed Computing

THEME

Distributed Systems and middleware

Inria

Contents

Project-Team WIDE

Creation of the Project-Team: 2018 June 01

Keywords

Computer sciences and digital sciences

A1.2.5. – Internet of things

A1.2.9. – Social Networks

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

A3.5.1. – Analysis of large graphs

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

1 Team members, visitors, external collaborators

Research Scientists

- Davide Frey [INRIA, Researcher]
- George Giakkoupis [INRIA, Researcher]
- Erwan Le Merrer [INRIA, Senior Researcher, until Sep 2024]

Faculty Members

- François Taiani [Team leader, UNIV RENNES, Professor Delegation, until Aug 2024]
- Yerom David Bromberg [UNIV RENNES, Professor]
- Brice Ekane Apah [UNIV RENNES, Associate Professor, from Sep 2024]
- Achour Mostefaoui [UNIV RENNES, Professor, from Sep 2024]
- Barbe Mvondo Djob [UNIV RENNES, Associate Professor]
- Michel Raynal [UNIV RENNES, Emeritus]

Post-Doctoral Fellows

- Georgy Ishmaev [UNIV RENNES, Post-Doctoral Fellow, from May 2024]
- Dimitrios Los [INRIA, Post-Doctoral Fellow, from Sep 2024]
- Luis Santiago Luevano Garcia [INRIA, Post-Doctoral Fellow, until May 2024]

PhD Students

- Timothé Albouy [UNIV RENNES]
- Fonyuy-Asheri Caleb [INRIA, from Feb 2024]
- Opale Duvivier [UNIV RENNES, CIFRE]
- Jade Garcia Bourree [INRIA, until Sep 2024]
- Adrien Gegout [UNIV RENNES, CIFRE]
- Mathieu Gestin [INRIA]
- Augustin Godinot [UNIV RENNES, until Sep 2024]
- Amelie Gonzalez [UNIV RENNES]
- Junrui Hua [HIVE COMPUTING SERVICES SAS, CIFRE, from Sep 2024]
- Dimitri Lereverend [INRIA]
- Honore Cesaire Mounah [INRIA]
- Victoire Nganfang [UNIV RENNES, from Oct 2024]
- Rémy Raes [Inria]
- Arthur Rauch [INRIA]
- Manon Sourisseau [UNIV RENNES]

Technical Staff

- Olivier Deloubriere [INRIA, Engineer, from Jun 2024]
- Patricio Inzaghi [INRIA, Engineer, from Jul 2024]
- Cyrille Kenfack [INRIA, Engineer]
- Elie Raspaud [INRIA, Engineer, from Apr 2024]
- Harvey Williams [INRIA, Engineer, from Apr 2024]

Interns and Apprentices

- Ivan Bouh [UNIV RENNES, Intern, from Mar 2024 until Aug 2024]
- Rahab Epse Kaldjonbe [UNIV RENNES, Intern, from Jun 2024 until Sep 2024]
- Audrey Fongue [UNIV RENNES, Intern, from Mar 2024 until Aug 2024]
- Armand Ledoux [ENS RENNES, Intern, from May 2024 until Jul 2024]
- Eugenio Mazzina [INRIA, Intern, from Sep 2024]
- Victoire Nganfang [INRIA, Intern, until Feb 2024]
- Stella Tchoutcha [UNIV RENNES, Intern, from Mar 2024 until Aug 2024]
- Wilson Waha Lindjeck [UNIV RENNES, Intern, until Jun 2024]

Administrative Assistant

- Virginie Desroches [INRIA]

Visiting Scientist

- Dimitrios Los [UNIV CAMBRIDGE, from Apr 2024 until Aug 2024]

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 services[1] 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,

[1] The 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.

- 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 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 [61, 73] 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) [65]. 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 [62] 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 [56].

A number of efforts in this directions focus on specific hardware, e.g. fog nodes that are responsible for connected IoT devices [63]. 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 [94], 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 [93], Kubernetes [60]) 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 [64, 66, 67] 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 [68, 75]. 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 [57].

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 [83, 81, 82] that have been studied in this context so far. This includes the *randomized rumor spreading (RS)* model for information dissemination [72], *biased* versions of the *voter model* [77] modelling influence, and the (graph-based) *Moran processes* [85] 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 [83], 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 [59]. 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 inaproximability 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 [74]. 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 [92]) and *process adversaries* (investigated in several papers, e.g. [90, 71, 79, 80, 84]). 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 [90, 55].

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 Erwan Le Merrer (Inria) and Djob Mvondo (University of Rennes) has brought new expertise to WIDE, and the opportunity to expand our activities regarding the remote auditing of large-scale black-box AI systems (for Erwan), and to deepen our understanding of the lower levels of large-scale distributed infrastructures (for Djob). 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 [78]. 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 [69, 76]. 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 of proof-of-stake algorithms. One open, yet enticing, questions 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 [86] 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 [57], 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 [87, 88]. 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: Collaborative and stealthy audits of algorithms

This research objective is interested in the possibility of (and the algorithmic means for) auditing algorithms running at third parties (such as classifiers, recommenders or ranking applications) [70]. These algorithms, often coined *black-box algorithms* [89], can only be interacted with by sending inputs and observing the result of their computation through outputs. While their full reverse engineering is either intractable or even undecidable (i.e., retrieving a full map of the outputs depending on all the possible inputs), the coordinated action of several observers (or *auditors*) can help infer important properties of these algorithms, such as bias, stability or security in their decisions.

The challenges are thus 1) to first understand what can or cannot be inferred, given for instance a number of requests as inputs, a set of assumptions for what is running in the black-box, and considering which type of adversary is running and modifying the audited algorithm; 2) to turn initial theoretical results into practical tools. To this end, we must find ways to interface with the audited algorithm in vivo, so that input/output interactions can be performed. This may imply coordinating of various auditors, and sharing their observation results for better efficiency.

Objective 4: 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), privacy protection (Objective 2), and remote auditing (Objective 3).

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 [58], 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 while 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 of the University of Rennes.
- Davide Frey is part of the SENS (science and environment) group at Inria of the University of Rennes

6 Highlights of the year

- WIDE organized a one-day workshop on Reliable Distributed Systems and Blockchain on Dec 17 2024, in Rennes. The workshop featured key international experts who showcased key recent results on distributed systems and blockchain research.
- WIDE organized a one-day workshop on Operating Systems on April 15 2024 in the context of the Inria challenge "Operating System". The workshop showcased the latest research in Operating Systems conducted by the four Inria teams highly active in the field, namely Whisper (Inria Paris), KrakOS (Inria Grenoble), Benagil (Inria Saclay), and WIDE (Inria Rennes).
- WIDE organize a one-day workshop on Virtualization on 20 September 2024 for the kickoff of the ANR project "Second Chance" led by David Bromberg. It showcased the latest research on Operating Systems and virtualization conducted by the core teams involved, KrakOS (Inria Grenoble) and WIDE (Inria Rennes), along with guest teams active in Operating Systems research, such as Stack (Inria Nantes) and INP Toulouse.
- WIDE welcomed two new permanent members in September 2024: Achour Mostéfaoui (PR, U. Rennes) and Brice Ekane (MCF, U. Rennes).
- In 2024, WIDE's research results continue to be published in some of the most visible and prestigious conferences of its field (FOCS, PODC, Middleware, ICDCS).

7 New software, platforms, open data

7.1 New software

7.1.1 DecentralizedFlower

Name: DecentralizedFlower

Keyword: Decentralized Learning

Functional Description: DecentralizedFlower 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 copperation 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

8 New results

8.1 Distributed Algorithms and Systems

8.1.1 Foundations of Reliable Cooperation under Asynchrony, Byzantine Faults, and Message Adversaries

Participants: Timothé Albouy.

This PhD thesis [45] explores fault-tolerant distributed systems. It focuses more specifically on implementing reliable broadcast in asynchronous environments prone to hybrid failures. We introduce a novel computing model combining Byzantine process failures with a message adversary. We then define the Message-Adversary-tolerant Byzantine Reliable Broadcast (MBRB) abstraction and prove its optimal resilience condition. We present three key algorithms implementing this abstraction: a simple signature-based MBRB algorithm, a new primitive called $k2\ell$-cast for cryptography-free MBRB implementations, and an erasure-coding-based MBRB algorithm optimizing communication complexity. These contributions advance the understanding of fault-tolerant distributed systems and provide a foundation for designing resilient and efficient distributed algorithms, with applications in critical infrastructures, financial systems, and blockchain technologies.

8.1.2 Privacy Preserving and fully Distributed Identity Management Systems

Participants: Mathieu Gestin.

This PhD thesis [46] focuses on privacy preserving and fully distributed identity management systems. These systems aim to allow a user to authenticate and be authorized by a service provider while only revealing strictly necessary information. In addition, these systems must be resilient to the presence of malicious processes. In this context, we are interested in two points. Firstly, anonymous credentials and their privacy properties. We identify a shortcoming that reduces this property in state of the art, and we correct it with a new type of signature: hidden issuer anonymous credentials. Next, we look at the distributed algorithms used for the auxiliary properties of distributed identity management systems, in particular for certificate revocation and public key management. We analyze these problems formally, particularly from the point of view of their *consensus number*. Finally, these analyses allow us to propose algorithms for implementing a fully distributed identity management system that requires reduced synchronization. In other words, a system where the use of consensus algorithms is reduced to a minimum.

8.1.3 Towards more scalable and privacy-preserving distributed asset transfer systems

Participants: Arthur Rauch.

Since 2018, there has been a notable increase in the popularity of peer-to-peer dis- tributed systems. In particular, blockchain technology has seen the emergence of numerous applications, spanning from cryptocurrency to digital identity, and healthcare systems. Several authors have dedicated efforts to analyzing the tensions between decentralized sys- tems such as the Blockchain, and privacy regulations. Most existing blockchains adopt a full replication model. From a legal perspective, the fully replicated nature of blockchains means that personal data is likely to be stored throughout the world, on blockchain nodes distributed across different countries. From a technical perspective, full replication pro- vides good fault tolerance at the cost of scalability. To achieve scalability, we must develop solutions that can ensure fault-tolerance with more reasonable levels of replication, while protecting privacy and avoiding clashes with national or cross-border regulatory laws. To address these issues, this PhD thesis [47] proposes two systems. The first is based on horizontal partitioning (sharding) of the blockchain to better distribute the costs of data storage and processing among subsets of peers. The second does not rely on consensus. It can therefore process independent transactions concurrently. Furthermore, it uses a set of cryptographic primitives to anonymize users' data exchanges and verify their legitimacy, without revealing or storing sensitive data.

8.1.4 AMECOS: A modular event-based framework for concurrent object specification

Participants: Timothé Albouy, Mathieu Gestin.

This work [21] introduces a modular framework for specifying distributed systems that is called AMECOS. Specifically, this framework departs from the traditional use of sequential specification, which presents limitations both on the specification expressiveness and implementation efficiency of inherently concurrent objects, as documented by Castañeda, Rajsbaum and Raynal in CACM 2023. The framework focuses on the interactions between the various system components, specified as concurrent objects. Interactions are described with sequences of object events. This provides a modular way of specifying distributed systems and separates legality (object semantics) from other issues, such as consistency. The work demonstrates the usability of the framework by (i) specifying various well-known concurrent objects, such as registers, shared memory, message-passing, reliable broadcast, and consensus, (ii) providing hierarchies of ordering semantics (namely, consistency hierarchy, memory hierarchy, and

reliable broadcast hierarchy), and (iii) presenting a novel axiomatic proof of the impossibility of the well-known Consensus problem.

This is a joint work with Antonio Fernández Anta (IMDEA Networks Institute), Chryssis Georgiou (University of Cyprus), Nicolas Nicolaou (Algosys Ltd), and Junlang Wang (IMDEA Networks Institute).

8.1.5 Near-optimal communication Byzantine reliable broadcast under a message adversary

Participants: Timothé Albouy, Davide Frey, Michel Raynal, François Taïani.

This work [23] addresses the problem of Reliable Broadcast in asynchronous message-passing systems with n nodes, of which up to t are malicious (faulty), in addition to a *message adversary* that can drop some of the messages sent by correct (non-faulty) nodes. This work presents a Message-Adversary-Tolerant Byzantine Reliable Broadcast (MBRB) algorithm that communicates $O(|m| + n\kappa)$ bits per node, where $|m|$ represents the length of the application message and $\kappa = \Omega(\log n)$ is a security parameter. This communication complexity is optimal up to the parameter κ. This significantly improves upon the state-of-the-art MBRB solution (Albouy, Frey, Raynal, and Taïani, TCS 2023), which incurs communication of $O(n|m| + n^2\kappa)$ bits per node. This solution sends at most $4n^2$ messages overall, which is asymptotically optimal. Reduced communication is achieved by employing coding techniques that replace the need for all nodes to (re-)broadcast the entire application message m. Instead, nodes forward authenticated fragments of the encoding of m using an erasure-correcting code. Under the cryptographic assumptions of threshold signatures and vector commitments, and assuming $n > 3t + 2d$, where the adversary drops at most d messages per broadcast, the algorithm allows at least $\ell = n - t - (1+\epsilon)d$ (for any arbitrarily low $\epsilon > 0$) correct nodes to reconstruct m, despite missing fragments caused by the malicious nodes and the message adversary.

This is a joint work with Ran Gelles (Bar-Ilan University), Carmit Hazay (Bar-Ilan University), Elad Michael Schiller (Chalmers University of Technology), and Vassilis Zikas (Georgia Institute of Technology). A brief announcement was also published in DISC 2024 [22].

8.1.6 Good-case early-stopping latency of synchronous Byzantine reliable broadcast: the deterministic case

Participants: Timothé Albouy, Davide Frey, Michel Raynal, François Taïani.

This work [15] considers the good-case latency of Byzantine Reliable Broadcast (BRB), i.e., the time taken by correct processes to deliver a message when the initial sender is correct. This time plays a crucial role in the performance of practical distributed systems. Although significant strides have been made in recent years on this question, progress has mainly focused on either asynchronous or randomized algorithms. By contrast, the good-case latency of deterministic synchronous BRB under a majority of Byzantine faults has been little studied. In particular, it was not known whether a good-case latency below the worst-case bound of $t+1$ rounds could be obtained. This work answers this open question positively and proposes a deterministic synchronous Byzantine reliable broadcast that achieves a good-case latency of $\max(2, t+3-c)$ rounds (or equivalently $\max(2, f+t+3-n)$), where t is the upper bound on the number of Byzantine processes, $f \leq t$ the number of effectively Byzantine processes, and $c = n - f$ the number of effectively correct processes. The proposed algorithm does not put any constraint on t, and assumes an authenticated setting, in which individual processes can sign the messages they send, and verify the authenticity of the signatures they receive.

8.1.7 Partition Detection in Byzantine Networks

Participants: Yerom David Bromberg, Manon Sourisseau, François Taïani.

Detecting and handling network partitions is a fundamental requirement of distributed systems. Although existing partition detection methods in arbitrary graphs tolerate unreliable networks, they either assume that all nodes are correct or that a limited number of nodes might crash. In particular, Byzantine behaviors are out of the scope of these algorithms despite Byzantine fault tolerance being an active research topic for important problems such as consensus. Moreover, Byzantine-tolerant protocols, such as broadcast or consensus, always rely on the assumption of connected networks. This work [27] addresses the problem of detecting partition in Byzantine networks (without connectivity assumption). We present a novel algorithm, which we call NECTAR, that safely detects partitioned and possibly partitionable networks and prove its correctness. NECTAR allows all correct nodes to detect whether a network could suffer from Byzantine nodes. We evaluate NECTAR's performance and compare it to two existing baselines using up to 100 nodes running real code, on various realistic topologies. Our results confirm that NECTAR maintains a 100% accuracy while the accuracy of the various existing baselines decreases by at least 40% as soon as one participant is Byzantine. Although NECTAR's network cost increases with the number of nodes and decreases with the network's diameter, it does not go above around 500KB in the worst cases.

This is a joint work with Jérémie Decouchant (TU Delft, The Netherlands)

8.1.8 Sharding in Permissionless Systems in Presence of an Adaptive Adversary

Participants: Davide Frey, Arthur Rauch.

This work introduced SplitChain, a protocol intended to support the creation of scalable proof-of-stake and account-based blockchains without undermining decentralization and security. This is achieved by using sharding, i.e. by splitting the blockchain into several lighter chains managed by their own disjoint sets of validators called shards. These shards balance the load by processing disjoint sets of transactions in parallel. SplitChain distinguishes itself from other sharded blockchains by reducing the synchronization constraints among shards while maintaining security guarantees in an asynchronous setting. A dedicated routing protocol enables transactions to be redirected between shards with a low number of hops and messages. Finally, the protocol is designed to dynamically adapt the number of shards to the system load to avoid over-dimensioning issues encountered in static sharding-based solutions. A preliminary version of this work appeared as a brief announcment at SIROCCO 2024 [24] and the full version was published at NETYS 2024 [25]. Olivier Deloubriere is working together with Patricio Inzaghi under the supervision of Davide Frey on the implementation of a prototype of SplitChain.

This is joint work with Emmanuelle Anceame from Inria Team PIRAT.

8.1.9 Process-commutative distributed objects: From cryptocurrencies to Byzantine-Fault-Tolerant CRDTs

Participants: Davide Frey, Michel Raynal, François Taïani.

With this contribution, published in Theoretical Computer Science [17], we explored the territory lying between best-effort Byzantine-Fault-Tolerant Conflict-free Replicated Data Types (BFT CRDTs) and totally ordered distributed ledgers, such as those implemented by Blockchains. We formally characterized a novel class of distributed objects that only requires a First In First Out (FIFO) order on the object operations from each process (taken individually). The formalization leverages Mazurkiewicz traces to define legal sequences of operations and ensure both Strong Eventual Consistency (SEC) and Pipleline Consistency (PC). We presented a generic algorithm that implements this novel class of distributed objects both in a crash- and Byzantine setting. We also illustrated the practical interest of the proposed approach using four instances of this class of objects, namely money transfer, Petri nets, multi-sets, and concurrent work stealing dequeues. This is joint work with Lucie Guillou, a former intern who is currently a PhD student at IRIF.

8.1.10 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 COAST Inria team and Mathieu Turuani from the PESTO Inria team.

8.1.11 Faster Randomized Repeated Choice and DCAS

Participants: George Giakkoupis.

At STOC 2021, Giakkoupis, Giv, and Woelfel presented an efficient randomized implementation of Double Compare-And-Swap (DCAS) from Compare-And-Swap (CAS) objects. DCAS is a useful and fundamental synchronization primitive for shared memory systems, which, contrary to CAS, is not available in hardware. The DCAS algorithm has $O(\log n)$ expected amortized step complexity against an oblivious adversary, where n is the number of processes in the system. The bottleneck of this algorithm is a building block, introduced in the same paper: A repeated choice (RC) object, which allows processes to propose values, and later agree on (and "lock in") one of the proposed values, which is roughly uniformly distributed among the "recently" proposed ones. The object can then be unlocked, and the process be repeated. The RC implementation introduced by Giakkoupis et al. has step complexity $O(\log n)$. In [26], we present a more efficient RC algorithm, with similar probabilistic guarantees, but expected step complexity $O(\log\log n)$. We then show how this improved RC object can be used to achieve an exponential improvement in the expected amortized step complexity of DCAS.

This work was done in collaboration with Dante Bencivenga (University of Calgary) and Philipp Woelfel (University of Calgary).

8.1.12 Self-stabilizing MIS computation in the beeping model

Participants: George Giakkoupis.

In [33], we consider self-stabilizing algorithms to compute a Maximal Independent Set (MIS) in the extremely weak beeping communication model. The model consists of an anonymous network with synchronous rounds. In each round, each vertex can optionally transmit a signal to all its neighbors

(beep). After the transmission of a signal, each vertex can only differentiate between no signal received, or at least one signal received. We also consider an extension of this model where vertices can transmit signals through two distinguishable beeping channels. We assume that vertices have some knowledge about the topology of the network. We revisit the not self-stabilizing algorithm proposed by Jeavons, Scott, and Xu (2013), which computes an MIS in the beeping model. We enhance this algorithm to be self-stabilizing, and explore three different variants, which differ in the knowledge about the topology available to the vertices and the number of beeping channels. In the first variant, every vertex knows an upper bound on the maximum degree Δ of the graph. For this case, we prove that the proposed self-stabilizing version maintains the same run-time as the original algorithm, i.e., it stabilizes after $O(\log n)$ rounds w.h.p. on any n-vertex graph. In the second variant, each vertex only knows an upper bound on its own degree. For this case, we prove that the algorithm stabilizes after $O(\log n \cdot \log\log n)$ rounds on any n-vertex graph, w.h.p. In the third variant, we consider the model with two beeping channels, where every vertex knows an upper bound of the maximum degree of the nodes in the 1-hop neighborhood. We prove that this variant stabilizes w.h.p. after $O(\log n)$ rounds.

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

8.2 Balanced Allocations and Sorting

8.2.1 An asymptotically optimal algorithm for generating bin cardinalities

Participants: Dimitrios Los.

Efficient random variate generation is a crucial task for a large number of applications in machine learning, randomized simulations and the natural sciences. In [16], we design an algorithm that samples from the uniform multinomial distribution (a.k.a. the balls-into-bins setting) improving the running time exponentially from naïve sampling.

More precisely, 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 $\Theta(\log n/\log\log n)$. We present an algorithm in the RAM model that generates the bin cardinalities of the final load vector in the optimal $\Theta(\log n/\log\log n)$ time in expectation and with high probability. Further, we demonstrate that this algorithm 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).

8.2.2 An Improved Drift Theorem for Balanced Allocations

Participants: Dimitrios Los.

In the balanced allocations framework, there are m jobs (balls) to be allocated to n servers (bins). The goal is to minimize the gap, the difference between the maximum and the average load. In 2015, Peres, Talwar and Wieder used the hyperbolic cosine potential function to analyze the challenging case where $m \geq n$, for a large family of load balancing processes, including the $(1+\beta)$-process and graphical balanced allocations. The key ingredient was to prove that the potential drops in every step, i.e., a drift inequality.

In [18], we improve the drift inequality so that (i) it is asymptotically tight (leading to tighter gap bounds), (ii) it assumes weaker preconditions (thereby resolving an open problem regarding weighted graphical allocations), (iii) it applies not only to processes allocating to more than one bin in a single step but also (iv) to processes allocating a varying number of balls depending on the sampled bin. Our applications include the aforementioned large family of processes, and also several new processes and

settings, including outdated information and memory. We hope that our techniques can be used to analyze further interesting settings and processes.

This work was done in collaboration with Thomas Sauerwald (University of Cambridge).

8.2.3 Naïvely sorting evolving data is optimal and robust

Participants: George Giakkoupis, Dimitrios Los.

Sorting is one of the fundamental tasks in computing and a crucial building block for many algorithms solving more complicated tasks. Furthermore, in the current computing landscape there is a growing number of applications which require algorithms to data that is constantly changing.

In [32], we study comparison sorting in the evolving data model, introduced by Anagnostopoulos, Kumar, Mahdian and Upfal (2011), where the true total order changes while the sorting algorithm is processing the input. More precisely, each comparison operation of the algorithm is followed by a sequence of evolution steps, where an evolution step perturbs the rank of a random item by a "small" random value. The goal is to maintain an ordering that remains close to the true order over time. Previous works have analyzed adaptations of classic sorting algorithms, assuming that an evolution step changes the rank of an item by just one, and that a fixed constant number b of evolution steps take place between two comparisons. In fact, the only previous result achieving optimal linear total deviation, by Besa Vial, Devanny, Eppstein, Goodrich and Johnson (2018a), applies just for $b = 1$.

We analyze a very simple sorting algorithm suggested by Mahdian (2014), which samples a random pair of adjacent items in each step and swaps them if they are out of order. We show that the algorithm achieves and maintains, with high probability, optimal total deviation, $O(n)$, and optimal maximum deviation, $O(\log n)$, under very general model settings. Namely, the perturbation introduced by each evolution step is sampled from a general distribution of bounded moment generating function, and we just require that the average number of evolution steps between two sorting steps be bounded by an (arbitrary) constant, where the average is over a linear number of steps.

The key ingredients of our proof are a novel potential function argument that inserts "gaps" in the list of items, and a general analysis framework which separates the analysis of sorting from that of the evolution steps, and is applicable to a variety of settings for which previous approaches do not apply. Our results settle conjectures and open problems in the three aforementioned works, and provide theoretical support that simple quadratic algorithms are optimal and robust for sorting evolving data, as empirically observed by Besa Vial, Devanny, Eppstein, Goodrich and Johnson (2018b).

This work was done in collaboration with Marcos Kiwi (Universidad de Chile).

8.3 Artificial Intelligence, Machine Learning and Auditing

8.3.1 Face recognition and Anti-Spoofing

Participants: Davide Frey, Luis Santiago Luevano Garcia.

With the growing breakthrough of deep learning-based face recognition, the development of lightweight models that achieve high accuracy with computational and memory efficiency has become paramount, especially for deployment on embedded domains. While Vision Transformers have shown significant promising results in various computer vision tasks, their adaptability to resource-constrained devices remains a significant challenge. In this work, we first examined how pre-processing and training methods impact on the performance of Lightweight CNNs through evaluations on MobileNetV3 with a spoofing detection head, dubbed ”MobileNetV3-Spoof” [41]. Our results showed that pre-processing steps significantly boost the model’s ability to identify spoof samples, especially against complex attacks. Through detailed comparisons, we offer insights that could guide data curation and the creation of more effective and efficient antispoofing techniques suitable for real-world use in the era of digital face attacks. Then, we introduced SwiftFaceFormer [44], a new efficient, and lightweight family of face recognition

models inspired by the hybrid SwiftFormer architecture. Our proposal not only retains the representational capacity of its predecessor but also introduces efficiency improvements, enabling enhanced face recognition performance at a fraction of the computational cost. We also propose to enhance the verification performance of our original most lightweight variant by using a training paradigm based on Knowledge Distillation. Through extensive experiments on several face benchmarks, the presented SwiftFaceFormer demonstrates high levels of accuracy compared to the original SwiftFormer model, and very competitive results with respect to state-of-the-art deep face recognition models, providing a suitable solution for real-time, on-device face recognition applications. Our code is available at github.com/Inria-CENATAVTec/Assessing-Efficient-FAS-CVPR2024 and github.com/Inria-CENATAV-Tec/SwiftFaceFormer.

This work was done in coolaboration with Yoanna Martinez-Diaz and Heydi Mendez-Vazquez from Centro de Aplicaciones de Tecnologías de Avanzada Cuba, and with Miguel Gonzalez Mendoza from Monterrey Institute of Technology, Mexico.

8.3.2 LLMs hallucinate graphs too: a structural perspective

Participants: Erwan Le Merrer.

It is known that LLMs do hallucinate, that is, they return incorrect information as facts. In this work [35], we introduce the possibility to study these hallucinations under a structured form: graphs. Hallucinations in this context are incorrect outputs when prompted for well known graphs from the literature (e.g. Karate club, Les Misérables, graph atlas). These hallucinated graphs have the advantage of being much richer than the factual accuracy – or not – of a statement; this work thus argues that such rich hallucinations can be used to characterize the outputs of LLMs. Our first contribution observes the diversity of topological hallucinations from major modern LLMs. Our second contribution is the proposal of a metric for the amplitude of such hallucinations: the Graph Atlas Distance, that is the average graph edit distance from several graphs in the graph atlas set. We compare this metric to the Hallucination Leaderboard, a hallucination rank that leverages 10,000 times more prompts to obtain its ranking.

8.3.3 Challenges in archiving the personalized web

Participants: Erwan Le Merrer.

The decision-making algorithms embedded within online platforms are determining content shown to users. This personalization steers the dissemination of information, in contrast with the idea of a universal World Wide Web. Personalization thus generates a combinatorial explosion of different versions of the web, rendering each user's experience distinct. This raises critical questions: what elements of a personalized web should be archived? How can the collected user journeys capture a representative picture of our times? Navigating personalization is essential to capture the contemporary web experience, yet it presents methodological and technical challenges. In this work [42], we identify key challenges in performing a representative sampling of personalization within online platforms.

8.3.4 Under manipulations, are some AI models harder to audit?

Participants: Erwan Le Merrer, Augustin Godinot, François Taïani.

Auditors need robust methods to assess the compliance of web platforms with the law. However, since they hardly ever have access to the algorithm, implementation, or training data used by a platform, the problem is harder than a simple metric estimation. Within the recent framework of manipulation-proof auditing, we study in this work [34] the feasibility of robust audits in realistic settings, in which models exhibit large capacities. We first prove a constraining result: if a web platform uses models that may fit

any data, no audit strategy—whether active or not—can outperform random sampling when estimating properties such as demographic parity. To better understand the conditions under which state-of-the-art auditing techniques may remain competitive, we then relate the manipulability of audits to the capacity of the targeted models, using the Rademacher complexity. We empirically validate these results on popular models of increasing capacities, thus confirming experimentally that large-capacity models, which are commonly used in practice, are particularly hard to audit robustly. These results refine the limits of the auditing problem, and open up enticing questions on the connection between model capacity and the ability of platforms to manipulate audit attempts.

8.3.5 Fairness auditing with multi-agent collaboration

Participants: Erwan Le Merrer, Jade Garcia Bourrée, Benoit Rottembourg.

Existing work in fairness auditing assumes that each audit is performed independently. In this work [39], we consider multiple agents working together, each auditing the same platform for different tasks. Agents have two levers: their collaboration strategy, with or without coordination beforehand, and their strategy for sampling appropriate data points. We theoretically compare the interplay of these levers. Our main findings are that (i) collaboration is generally beneficial for accurate audits,(ii) basic sampling methods often prove to be effective, and (iii) counter-intuitively, extensive coordination on queries often deteriorates audits accuracy as the number of agents increases. Experiments on three large datasets confirm our theoretical results. Our findings motivate collaboration during fairness audits of platforms that use ML models for decision-making.

8.4 Large scale Cloud environments

8.4.1 Efficient hypervisors' update prediction on Cloud black-box workloads

Participants: Barbe Mvondo Djob.

We investigate the problem of getting hints on the effects of virtualization system (aka hypervisor) updates impact on virtual machines (VMs). System administrators can be reluctant to apply updates due to vague hints regarding the updates' impact on running applications. The problem is challenging since VMs are black boxes by design, reducing the scope of the data that can be retrieved and analyzed. Additionally, cloning VMs is only sometimes possible for obvious legal and privacy concerns. In this work [37], we present UTwinVM for Updated Twin VM, a mechanism to obtain valuable hints concerning the impact of updates on applications running in VMs. UTwinVM key idea is to generate a digital twin of running VMs that mimics the original VMs workloads behavior if they were running on with the updated virtualization stack. To achieve that, UTwinVM records several metrics regarding running VMs and the virtualization system on the initial server and on another server where the intended updated system runs. Then, it leverages a non-linear negative squared solver to determine how the initial system differs from the updated one. Based on that, it runs through specific scripts, workloads in VMs will match workloads running in the production VMs as if they were running on the updated system. Consequently, system administrators can observe these recreated workloads to obtain hints on the potential performance impact on production VMs. UTwinVM is non-intrusive for VMs and does not require modifications to the virtualization system.

This is a joint work with Tong Xing and Antonio Barbalace of the University of Edinburgh, Scotland, UK.

8.4.2 Efficient Buffer Overflow Detection In Virtualized Clouds

Participants: David Bromberg.

Write buffer overflow is a widespread and prevalent memory safety violation in C/C++, reported as the top vulnerability in 2022 and 2023. Secure memory allocators are generally used to protect systems against attacks that may exploit buffer overflows. Existing allocators mainly rely on two types of countermeasures to prevent or detect write overflows: canaries and guard pages, each with pros and cons in terms of detection latency and memory footprint. For virtualized cloud applications, this paper follows the Out of Hypervisor (OoH) trend and introduces GuaNary, a safety guard against write overflows, allowing synchronous detection at a low memory footprint cost. OoH is a new virtualization research axis introduced in 2022 advocating the exposure of hardware features for virtualization to the guest OS so that its processes can take advantage of them. Based on the OoH principle, GuaNary leverages Intel Sub-Page write Permission (SPP), a recent hardware virtualization feature that allows to write-protect guest memory at the granularity of 128B (namely, sub-page) instead of 4KB. We implement a software stack, LeanGuard, which promotes the utilization of SPP from inside virtual machines by new secure allocators that use GuaNary. Our evaluation shows that for the same number of protected buffers, LeanGuard consumes 8.3× less memory than SlimGuard, a recent state-of-art secure allocator. Further, for the same memory consumption, LeanGuard allows protecting 25× more buffers than SlimGuard.

This is a joint work with Stella Bitchebe, Yves Kone, Pierre Olivier, Jalil Boukhobza, Daniel Hagimont, Alain Tchana

8.4.3 Low latency workloads on FaaS platforms

Participants: Barbe Mvondo Djob, François Taiani, Yerom David Bromberg.

We investigate if FaaS platforms can handle ultra-low latency workloads that run as low as less than 1μs and show that even for a warm start, the initialization time takes up to 99,99% of the total execution time. This is due to the resume process of warm sandboxes that takes more time as the number of the sandbox's allocated virtual CPUs (vCPUs) increases. We uncover that two operations use up to 93,1% of the resume time. The first is the insertion of the paused sandbox's vCPUs to a CPU-sorted run queue. The second is the update of a lock-protected variable, which represents the vCPUs' load on each CPU. This variable is used for frequency scaling. In this work, we introduce Horse [37], for hot resume. Horse presents two simple approaches. The first is parallel precomputed sorted merge (P2SM), a parallel algorithm that leverages pre-computed data to have a parallel sorted merge of two sorted lists in O(1). The second is to coalesce the updates on the lock-protected variable used for frequency scaling. We implement Horse in Xen and Firecracker, two mainstream virtualization platforms. Our evaluation with real-world FaaS traces shows that Horse achieves up to 7, 16× resume time improvement and reduces sandbox initialization overhead by up to 142, 84× with no impact on functions.

8.4.4 Model serving frameworks evaluation

Participants: Yerom David Bromberg, Barbe Mvondo Djob.

In machine learning (ML), the inference phase is the process of applying pre-trained models to new, unseen data with the objective of making predictions. During the inference phase, end-users interact with ML services to gain insights, recommendations, or actions based on the input data. For this reason, serving strategies are nowadays crucial for deploying and managing models in production environments effectively. These strategies ensure that models are available, scalable, reliable, and performant for real-world applications, such as time series forecasting, image classification, natural language processing, and so on. In this work [91], we evaluate the performances of five widely-used

model serving frameworks (TensorFlow Serving, TorchServe, MLServer, MLflow, and BentoML) under four different scenarios (malware detection, cryptocoin prices forecasting, image classification, and sentiment analysis). We demonstrate that TensorFlow Serving is able to outperform all the other frameworks in serving deep learning (DL) models. Moreover, we show that DL-specific frameworks (TensorFlow Serving and TorchServe) display significantly lower latencies than the three general-purpose ML frameworks (BentoML, MLFlow, and MLServer).

This is a joint work with Pasquale De Rosa, Pascal Felber, and Valerio Schiovani of the University of Neuchatel, Switzerland.

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, Djob Mvondo, 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, Djob Mvondo, 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.

10 Partnerships and cooperations

10.1 International research visitors

10.1.1 Visits of international scientists

Other international visits to the team

Dimitrios Los

Status: Visiting Scientist

Institution of origin: Univ Cambridge

Country: UK

Dates: Apr - Aug 2024

Context of the visit: Collaboration with George Giakkoupis on Evolving Data Algorithms

Mobility program/type of mobility: Funded by London Mathematical Society (LMS) Early Career Fellowship

Sven Dietrich

Status: Visiting Professor

Institution of origin: CUNY

Country: US

Dates: Apr 2024

Context of the visit: Collaboration with David Bromberg on security in distributed systems.

Paulin Melatagia Yonta

Status: Visiting Professor

Institution of origin: Univ Yaoundé 1

Country: Cameroun

Dates: May 2024

Context of the visit: Collaboration with David Bromberg on AI for OS.

Daniel Cordeiro

Status: Visiting Professor

Institution of origin: Univ Yaoundé 1

Country: Brasil

Dates: September 2024

Context of the visit: Collaboration with David Bromberg on optimizing carbon footprint in cloud infrastructure.

10.1.2 Visits to international teams

Research stays abroad

Jade Garcia Bourrée

Visited institution: UQAM

Country: Canada

Dates: Summer of 2024

Context of the visit: Collaboration with Sébastien Gambs's group

Mobility program/type of mobility: internship

David Bromberg

Visited institution: University of Yaoundé 1

Country: Cameroun

Dates: May and December of 2024

Context of the visit: Collaboration with Paulin Melatagia

Mobility program/type of mobility: Visiting Professor

David Bromberg

Visited institution: University of Sao Paulo

Country: Brazil

Dates: February and July of 2024

Context of the visit: Collaboration with Daniel Cordeiro

Mobility program/type of mobility: Visiting Professor

David Bromberg

Visited institution: University of Goiania

Country: Brazil

Dates: February and July of 2024

Context of the visit: Collaboration with Fabio Costa

Mobility program/type of mobility: Visiting Professor

David Bromberg

Visited institution: University of UFABC

Country: Brazil

Dates: July 2024

Context of the visit: Collaboration with Emilio de Camargo Francesquini

Mobility program/type of mobility: Visiting Professor

David Bromberg

Visited institution: KAUST

Country: Arabia Saudia

Dates: May of 2024

Context of the visit: Collaboration with Marc Dacier's group

Mobility program/type of mobility: Visiting Professor

David Bromberg

Visited institution: University of Rome Sapienza

Country: Italy

Dates: June 2024

Context of the visit: Collaboration with Leornado Querzoni's group

Mobility program/type of mobility: Visiting Professor

David Bromberg

Visited institution: University of Lac Lac Tanganyika

Country: Burundi

Dates: June 2024

Context of the visit: Descovering research opportunities

Mobility program/type of mobility: Visiting Professor

David Bromberg

Visited institution: University of Berkeley

Country: USA

Dates: November 2024

Context of the visit: Collaboration with Natacha Crooks's group

Mobility program/type of mobility: Visiting Professor

10.2 European initiatives

10.2.1 H2020 projects

SOTERIA

Participants: Davide Frey, Luis Santiago Luevano-Garcia, Elie Raspaud, Harvey Williams.

SOTERIA project on cordis.europa.eu

Title: uSer-friendly digiTal sEcured peRsonal data and prIvacy plAtform

Duration: From October 1, 2021 to September 30, 2024

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France
- IPCENTER AT GMBH (IPCENTER), Austria
- NORIA ONLUS, Italy
- AUDENCIA (AUDENCIA), France
- STELAR SECURITY TECHNOLOGY LAW RESEARCH UG (HAFTUNGSBESCHRANKT) GMBH (STELAR), Germany

- Servicio Vasco de Salud Osakidetza (Osakidetza), Spain
- SCYTL ELECTION TECHNOLOGIES SL, Spain
- ERDYN ATLANTIQUE, France
- EDUPRO GROUP GMBH, Austria
- FONDATION DE L'INSTITUT DE RECHERCHE IDIAP (IDIAP), Switzerland
- ASOCIACION INSTITUTO DE INVESTIGACION SANITARIA BIOBIZKAIA (BIOBIZKAIA), Spain
- IDnow SAS (IDnow), France
- ASOCIATIA INFOCONS (INFOCONS), Romania
- FUNDACION VASCA DE INNOVACION E INVESTIGACION SANITARIAS (BIOEF), Spain
- ERDYN CONSULTANTS SAS, France
- CENTRE DE VISIO PER COMPUTADOR (CVC-CERCA), Spain
- KATHOLIEKE UNIVERSITEIT LEUVEN (KU Leuven), Belgium
- CENTRALESUPELEC, France

Inria contact: Davide Frey

Coordinator: Montaser Awal (IDnow)

Summary: SOTERIA aims to drive a paradigm shift on data protection and enable active participation of citizens to their own security, privacy and personal data protection. SOTERIA will develop and test in 3 large-scale real-world use cases, a citizen-driven and citizen-centric, cost-effective, marketable service to enable citizens to control their private personal data easily and securely. Led by an SME, this project will develop, using a user-driven and user-centric design, a revolutionary tool, uniquely combining, in a user-friendly manner, a high-level identification tool with a decentralised secured data storage platform, to enable all citizens, whatever their gender, age or ICT skills, to fully protect and control their personal data while also gaining enhanced awareness on potential privacy risks. SOTERIA solution will be tested and validated through 3 real-world large-scale use-cases, involving 6,500 European citizens, targeting 3 applications which usefulness has been highlighted during COVID-19 pandemic: e-learning, e-voting and e-health. This 3-year transdisciplinary project from both SSH and technology angles, will develop an innovative solution based on: a secured access interface relying on high-level identification, a smart platform processing data to transmit only the minimum personal data required, a secured data storage platform (decentralized architecture) under the full control of the citizen, an educational tool to raise awareness of citizens developed using a citizen-driven and citizen-centric approach. The technologies developed will i) empower citizens to monitor and audit their personal data; ii) restore trust on privacy, security and personal data protection of citizens in digital services; iii) be fully compliant to GDPR regulation and apply strictly the data minimization principle; iv) ensure cybersecurity.

10.3 National initiatives

Collaboration with the (PEReN) Pôle d'expertise et de régulation du numérique

Participants: Erwan Le Merrer.

Collaborating with the PEReN on what types of platform audits are feasible or not. In a collaboration through the Ph.D. thesis of Augustin Godinot.

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)

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).

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 Taiani, Davide Frey, Cyrille Kenfack, Remy Raes.

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 form Inria Lille. Davide Frey also supervises the work of Cyril Kenfack (Engineer) in order to contribute to a benchmarking environment for the with experimentation federated and decentralized learning platforms and algorithms.

Inria Challenge Project Alvearium

Participants: François Taiani, 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, MYRIADS, 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 OS

Participants: David Bromberg, Djob Mvondo.

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

10.4 Regional initiatives

Cominlabs Project PriCLeSS (2021-2024)

Participants: Davide Frey, Arthur Rauch, Michel Raynal, François Taiani.

Promoters of blockchain-based systems such as cryptocurrencies have often advocated for the anonymity these provide as a pledge of privacy protection, and blockchains have consequently been envisioned as a way to safely and securely store data. Unfortunately, the decentralized, fully-replicated and unalterable nature of the blockchain clashes with both French and European legal requirements on the storage of personal data, on several aspects such as the right of rectification and the preservation of consent. PriCLeSS aims to establish a cross-disciplinary partnership between Computer Science and Law researchers to understand and address the legal and technical challenges associated with data storage in a blockchain context.

Partners: WIDE@Inria (coordinator), CIDRE@Inria, GDD@LS2N (Université de Nantes) in Nantes.

11 Dissemination

11.1 Promoting scientific activities

Participants: George Giakkoupis, Erwan Le Merrer, Davide Frey, David Bromberg, François Taïani, Barbe Mvondo Djob.

11.1.1 Scientific events: organisation

Member of the organizing committees

- Davide Frey co-organized the PriCLeSS International Workshop at Inria of the University of Rennes, on September 9, 2024, the workshop was funded by the PriCleSS and Splitchain projects.
- Davide Frey organized the WIDE Workshop on Reliable Distributed Systems and Blockchain, at Inria of the University of Rennes, December , 2024.

11.1.2 Scientific events: selection

Chair of conference program committees

- Davide Frey is program co-chair for the 12th Workshop on Principles and Practice of Consistency for Distributed Data, PaPoC 2025.

Member of the conference program committees

- George Giakkoupis served on the PC of the 38th International Symposium on Distributed Computing (DISC), Madrid, Spain, Oct 28- Nov 1 2024.
- Erwan Le Merrer served on the PC of the SIAM International Conference on Data Mining, 2024.
- Erwan Le Merrer served on the PC of the 27TH European conference on artificial intelligence, 2024.
- Erwan Le Merrer served on the PC of the 17TH ACM Workshop on Artificial Intelligence and Security, 2024.
- Erwan Le Merrer served on the PC of the European Conference on Machine Learning and Data Mining, 2024.
- Erwan Le Merrer served on the PC of the 30TH ACM SIGKDD conference on knowledge discovery and data mining, 2024.
- François Taïani served on the PC of the 54th Annual IEEE/IFIP International Conference on Dependable Systems and Networks 2024, (DSN 2024).
- François Taïani served on the PC of the 20th European Dependable Computing Conference 2025, (EDCC 2025).
- Djob Mvondo served on the PC of the 25th ACM/IFIP International Middleware Conference 2024, (Middleware 2024).
- Djob Mvondo served on the PC of the 20th ACM European Systems Conference 2025, (EUROSYS 2025).
- David Bromberg served on the PC of the 20th ACM European Systems Conference 2025, (EUROSYS 2025).
- David Bromberg served on the PC of the 54th Annual IEEE/IFIP International Conference on Dependable Systems and Networks 2024, (DSN 2024).
- David Bromberg served on the PC of the 21st USENIX Symposium on Networked Systems Design and Implementation, (NSDI 2024).
- Davide Frey served on the PC of the 25th ACM/IFIP International Middleware Conference 2024, (Middleware 2024).
- Davide Frey served on the PC of the 26th ACM/IFIP International Middleware Conference 2025, (Middleware 2025).
- Davide Frey served on the PC of the 43rd International Symposium on Reliable Distributed Systems (SRDS 2024).

11.1.3 Journal

Reviewer - reviewing activities

- George Giakkoupis reviewed papers for journals Distributed Computed (DIST), IEEE Transactions on Networking (ToN), and Autonomous Agents and Multi-Agent Systems (AAMSFJ)
- Davide Frey reviewed papers for Blockchain Research and applications and for the Parallel Computing journal.

11.1.4 Invited talks

- George Giakkoupis. Expanders via local edge flips. Dagstuhl Seminar Graph Algorithms: Distributed Meets Dynamic Dagstuhl, Germany, Nov 24 2024
- George Giakkoupis. Sorting evolving data. Mixing Times Workshop with Applications to Theoretical Computer Science, University of Bath, UK, Sep 16 2024
- George Giakkoupis. Simple self-stabilizing distributed algorithms for Maximal Independent Set selection. 10th Workshop on Biological Distributed Algorithms (BDA), La Cité des Congrés, Nantes, France, Jun 17 2024
- Erwan Le Merrer. Algorithmic audits of AIs: What do we know (we don't know)?, 9th GDR RSD / ASF Winter School on Distributed Systems & Networks, 2024.
- Davide Frey. From network-wide to localized: some consensus but not too much! EPFL Lausanne, Switzerland. June 10, 2024.
- Davide Frey. From network-wide to localized: some consensus but not too much! Politecnico di Milano, Italy. October 11, 2024.
- David Bromberg. Major Android safety breach: Be aware of your Android. USP, Sao Paulo, Brazil July, 2024
- David Bromberg. Attack on the Metaverse. UFG, Goiania, Brazil, Janv. 2024
- David Bromberg. Major Android safety breach: Be aware of your Android. Sapienza, Rome, Italy, 204

11.1.5 Leadership within the scientific community

- George Giakkoupis is a member of the steering committee (member-at-large) of the ACM Symposium on Principles of Distributed Computing (PODC) from 2023-2026
- François Taïani serves as co-chair of the scientific committee of GDR RSD (Groupement de Recherche Réseaux & Systèmes Distribués) since 2024.
- David Bromberg is a member of the Steering Committee of the EuroSys ACM SIGOPS

11.1.6 Scientific expertise

- George Giakkoupis is a member of the Working Group GT CoA: Complexité et Algorithmes from 2023-2026
- George Giakkoupis reviewed grant proposals for the Israel Science Foundation (ISF), and the Agence Nationale de la Recherche (ANR)
- Erwan Le Merrer reviewed a rescrit JEI for the "crédit d'impôt recherche" program.
- François Taïani reviewed a grant protposal for the Swiss National Science Foundation (SNSF).
- Djob Mvondo reviewed a grant proposal for the Agence Nationale de la Recherche (ANR) and for Etoiles Montantes en Pays de la Loire
- David Bromberg is reviewer for the "crédit d'impôt recherche" program.

11.1.7 Research administration

- George Giakkoupis is a local correspondent of the Inria Centre Univ Rennes for the preparation of the Annual Activity Reports by the project teams.
- François Taïani is a member of the thesis committee of the Doctoral School Matisse (ED N. 601).
- François Taïani is a member of the Bureau du Comité des Projets (BCP) of the Centre Inria de l'Université de Rennes.
- François Taïani is a member of the local Inria secondment committee (Comission des délégations Inria) in Rennes.
- François Taïani is a Career Advice Person, (Référent conseil-parcours professionel chercheurs) for IRISA/Centre Inria Université de Rennes since 2019.
- David Bromberg is leading the D1 department at IRISA entitled Secured and large scale systems for IRISA/Centre Inria Université de Rennes since 2020.
- David Bromberg is leading the Laboratoire International de Recherche en Informatique et mathématiques appliquées (LIRIMA) since 2023.
- David Bromberg served on the recruitment committee for an MCF "Network" of University of Vannes.
- David Bromberg served on the recruitment committee for a PR "System" of University of Grenoble.
- Davide Frey was part of the recruitment committee for CRCN-ISFP positions at Inria of the University of Rennes, in 2024.
- François Taïani served as vice-chairman of the recruitment committee for an MCF "Réseaux couches hautes, systèmes, cloud" of ISTIC (University of Rennes).
- François Taïani served as vice-chairman of the recruitment committee for an MCF "Langage, programmation, sciences du logiciel" of ISTIC (University of Rennes).
- François Taïani served on the recruitment committee for an MCF "Système" of ENSEEIHT (INPT, Toulouse).
- François Taïani served on the recruitment committee for an MCF "Informatique - Systèmes d'exploitation, virtualisation et sécurité" of INSA Lyon.

11.2 Teaching - Supervision - Juries

Participants: George Giakkoupis, Erwan Le Merrer, Davide Frey, Yerom David Bromberg, François Taïani, Barbe Mvondo Djob, Brice Ekane Apah.

11.2.1 Teaching

- ENS L3: George Giakkoupis, Distributed Algorithms, 9h, L3 parcours SI, ISTIC, ENS Rennes, France.
- Engineering School: Barbe Thystere Mvondo Djob, Network and Security for IOT, 45h, ESIR M1, Rennes, France
- Engineering School: Barbe Thystere Mvondo Djob, Cloud Computing for IOT, 45h, ESIR M2, Rennes, France
- Engineering School: Djob Mvondo, Cybersecurity and Hacking, 30h, Polytechnic 2nd year, Paris, France

- 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, 24h, 1st year of Engineering School (L3), ESIR / 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, ENS Rennes, France.
- Master: Davide Frey, Distributed Systems/Systèmes Répartis, 21h, ENSAI, France.
- Master: Davide Frey, Recommender Systems, 21h, ENSAI, France.
- Master: Brice Ekane, Architecture Réseau, 38h, M2 Parcours Cloud & Réseau, ISTIC, U. Rennes, France.
- Master: Brice Ekane, Architecture Réseau Nouvelle Génération, 10h, M2 Parcours Cloud & Réseau & IL, ISTIC M2, U. Rennes, France.
- Licence: Brice Ekane, Réseau, 60h, L2, ISTIC, U. Rennes, France.

11.2.2 Supervision

- PhD: Timothé Albouy, Towards Lightweight Scalable and Open Byzantine-Fault-Tolerant Distributed Objects, U. Rennes, supervised by François Taïani and Davide Frey, started on Oct 18 2021, defended on Dec 16 2024.
- PhD: Arthur Rauch, Towards more scalable and privacy-preserving distributed asset transfer systems, Inria Rennes, supervised by Emmanuelle Anceaume and Davide Frey, started on Oct 1 2021, defended on December 18, 2024.
- PhD: Mathieu Gestin, Privacy Preserving and fully Distributed Identity Management Systems, Inria Rennes, supervised by Davide Frey, started on Oct 1 2021, defended on December 18, 2024.
- PhD in progress: Dimitri Lerévérend, Privacy-Preserving Decentralized Learning Through Model Fragmentation and Private Aggregation, started in September 2023, supervised by Davide Frey, Romaric Gaudel (LACODAM 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 Taiani (WIDE team, Inria Rennes), Romain Rouvoy, Lionel Seinturier (Spira,ls team, Inria Lille).
- PhD in progress: Ludovic Paillat, Security for peer-to-peer cloud storage without central authority, started in 2023, supervised by Davide Frey, (WIDE team, Inria Rennes), Claudia Ignat (COAST team, Inria Nancy), Alexandru Dobrila (HIVE), Mathieu Turiani (PESTO Team, Inria Nancy).
- PhD in progress: Jade Garcia Bourrée, Trust but verify: bot-driven audits of AI systems, started in October 2022, supervised by Erwan Le Merrer and Gilles Trédan (LAAS/CNRS).
- PhD in progress: Augustin Godinot, Auditing the mutations of AI-models, started on November 2022, supervised by Erwan Le Merrer, Gilles Trédan (LAAS/CNRS), François Taïani and Camilla Penzo (PEReN).

- PhD in progress: Gurvan Richardeau, started on December 2024, supervised by Erwan Le Merrer, Gilles Trédan (LAAS/CNRS) and Camilla Penzo (PEReN).
- PhD in progress: Cesaire Honoré, Scheduling in heterogeneous architectures, started on December 2022, supervised by Yerom David Bromberg and Djob Mvondo
- PhD in progress: Amelie Gonzalez, Linux network stack optimization, Started on September 2023, supervised by Yerom-David Bromberg, Djob Mvondo, Julia Lawal (Inria Paris)
- PhD in progress: Alexandre Duvivier, CDN performance optimization, Started on October 2023, supervised by Yerom-David Bromberg, Djob Mvondo, Nicolas Le Scouarnec (Broadpeak)
- PhD in progress: Adrien Gegout, Efficient containerized Cloud Gaming, Started on October 2023, supervised by Davide Frey, Djob Mvondo, Pascal Manchon (Blacknut).
- 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 on November 2024, supervised by François Taiani (WIDE team, Inria Rennes), Gérald Oster (COAST team, Inria Nancy), and Alexandru Dobrila (Hive).
- PhD in progress: Victoire Nganfang, Android Malware Poliferation Mitigation, Started on November 2024, supervised by Yerom David Bromberg, Djob Mvondo, Valerio Schiovani (University of Neuchatel, Switzerland)
- PhD in progress: Caleb Fonyuy-Asheri, Heterogeneous VM migration , Started on November 2023, supervised by Yerom David Bromberg, Alain Tchana (Grenoble INP), Djob Mvondo, Renaud Lachaise (UGA)
- PostDoc: Geovani Rizk (EPFL), Decentralized Learning with Byzantine Agents, in collaboration with Rachid Gerraoui within the Inria-EPFL laboratory.
- PostDoc: Dimitrios Los, Distributed Algorithms on Evolving Data, started Sep 2024, supervised by George Giakkoupis, funded by Action Exploratoire DisEvo: Distributed Algorithms on Evolving Data

11.2.3 Juries

- Davide Frey was a reviewer (rapporteur) for the PhD thesis of Boubacar Kane "Les objets ajustés: Une approche bien fondée et efficace de la programmation concurrente", Institut Polytechnique de Paris.
- Davide Frey was a reviewer (rapporteur) for the PhD thesis of Hassan Nazeer Chaudhry "Efficient Processing of Graph-Based Data Streams", Politecnico di Milano, Italy.
- Davide Frey was a reviewer (rapporteur) for the PhD thesis of Badr Bellaj "Securing P2P Resource Sharing via Blockchain and GNN-Based Trust", Institut Polytechnique de Paris, defended on June 5, 2024.
- Davide Frey was a reviewer (rapporteur) for the PhD thesis of Anastasiia Kucherenko "Robustness Of Gossip Protocols", EPFL Lausanne, Switzerland, defended on June 10, 2024.
- François Taïani was an external examiner for Quentin Stokkink's PhD thesis: Systems for Digital Self-Sovereignty, TU Delft (NL), 8 April 2024.
- François Taïani was an external examiner for Mohamed Lechiakh's PhD thesis: Beyond personalization: towards beneficial and depolarized recommender systems, Mohammed VI Polytechnic University (UM6P, Morocco), 11 December 2024.
- François Taïani was examiner for Pierre-Antoine Rault's PhD thesis: Mécanismes de contrôle d'accès pour applications collaboratives sans autorité centrale, Université de Lorraine, 12 December 2024.

- François Taïani was examiner for Antoine Boutet's HDR thesis: Privacy issues in AI and geolocation : from data protection to user awareness, INSA Lyon, 10 December 2024.
- David Bromberg was a reviewer (rapporteur) for the PhD thesis of NGUETCHOUANG NGONGANG "Efficient Storage Virtualization in Cloud Environments", ENS Lyon, defended on September 12, 2024.
- David Bromberg was a reviewer (rapporteur) for the HDR thesis of Baptiste Lepers "Concurrence relachée mais concurrence prouvée",Université Grenoble Alpes, defended on December 12, 2024.
- Djob Mvondo was examiner for Armel Jeatsa Toulepi's PhD thesis: Optimisation de l'allocation de la mémoire cache CPU pour les fonctions cloud et les applications haute performance, University of Toulouse, 11 September 2024
- Brice Ekane was examiner for Théophile DUBUC PhD thesis: Détection d'anomalies de latence dans les systèmes distribués avec eBPF, ECOLE NORMALE SUPERIEURE DE LYON, 18 December 2024

11.3 Popularization

Participants: Erwan Le Merrer, Yerom David Bromberg.

11.3.1 Specific official responsibilities in science outreach structures

- Erwan Le Merrer lead the scientific board of the Société Informatique de France in 2024.
- David Bromberg is on the board of directors of the Société Informatique de France since 2021.

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

- Erwan Le Merrer was interviewed in journal Les Echos: "Face à des IA trop humaines, le défi de l'identification", 12 novembre 2024.

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. URL: https://inria.hal.science/hal-03791921.

[2] A. Auvolat, 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. 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. 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: Online 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. 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. 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. 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. URL: https://hal.inria.fr/hal-03195692.

[10] 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. URL: https://inria.hal.science/hal-04394851.

[11] 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. URL: https://hal.inria.fr/hal-02888286.

[12] E. Le Merrer, B. Morgan and G. Trédan. 'Setting the Record Straighter on Shadow Banning'. In: INFOCOM 2021 - IEEE International Conference on Computer Communications. Virtual, Canada: IEEE, May 2021, pp. 1–10. DOI: 10.1109/INFOCOM42981.2021.9488792. URL: https://hal.inria.fr/hal-03234771.

[13] E. Le Merrer and G. Trédan. 'Remote explainability faces the bouncer problem'. In: *Nature Machine Intelligence* 2.9 (2020), pp. 529–539. DOI: 10.1038/s42256-020-0216-z. URL: https://hal.laas.fr/hal-03048809.

[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] T. Albouy, D. Frey, M. Raynal and F. Taïani. 'Good-case early-stopping latency of synchronous byzantine reliable broadcast: the deterministic case'. In: *Distributed Computing* (22nd Mar. 2024), pp. 1–34. DOI: 10.1007/s00446-024-00464-6. URL: https://inria.hal.science/hal-04521960 (cit. on p. 17).

[16] 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. URL: https://hal.science/hal-04889571 (cit. on p. 20).

[17] D. Frey, L. Guillou, M. Raynal and F. Taïani. 'Process-commutative distributed objects: From cryptocurrencies to Byzantine-Fault-Tolerant CRDTs'. In: *Theoretical Computer Science* 1017 (Nov. 2024), p. 114794. DOI: 10.1016/j.tcs.2024.114794. URL: https://inria.hal.science/hal-04889162 (cit. on p. 18).

[18] D. Los and T. Sauerwald. 'An Improved Drift Theorem for Balanced Allocations'. In: *ACM Transactions on Algorithms* 20.4 (11th Oct. 2024), pp. 1–39. DOI: 10.1145/3673900. URL: https://hal.science/hal-04889279 (cit. on p. 20).

[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* (11th July 2024), pp. 1–23. DOI: 10.1007/s12243-024-01053-1. URL: https://inria.hal.science/hal-04829916 (cit. on p. 19).

[20] J. Rufino, J. M. Ramírez, J. Aguilar, C. Baquero, J. Champati, D. Frey, R. E. Lillo and A. Fernández-Anta. 'Performance and explainability of feature selection-boosted tree-based classifiers for COVID-19 detection'. In: *Heliyon* 10.1 (Jan. 2024), e23219. DOI: 10.1016/j.heliyon.2023.e23219. URL: https://inria.hal.science/hal-04406767.

International peer-reviewed conferences

[21] 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. URL: https://inria.hal.science/hal-04577664 (cit. on p. 16).

[22] T. Albouy, D. Frey, R. Gelles, C. Hazay, M. Raynal, E. M. Schiller, F. Taïani and V. Zikas. 'Brief Announcement: Towards Optimal Communication Byzantine Reliable Broadcast Under a Message Adversary'. In: DISC 2024 - 38th International Symposium on Distributed Computing. Vol. Proceedings of the 38th International Symposium on Distributed Computing (DISC 2024). Madrid, Spain: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024, pp. 1–7. DOI: 10.4230/LIPIcs.DISC.2024.41. URL: https://inria.hal.science/hal-04889285 (cit. on p. 17).

[23] T. Albouy, D. Frey, R. Gelles, C. Hazay, M. Raynal, E. M. Schiller, F. Taïani and V. Zikas. 'Near-Optimal Communication Byzantine Reliable Broadcast Under a Message Adversary'. In: *Proceedings of the 28th International Conference on Principles of Distributed Systems (OPODIS 2024)*. OPODIS 2024 - 28th International Conference on Principles of Distributed Systems. Lucques, Italy: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. DOI: 10.4230/LIPICS.OPODIS.2024.14. URL: https://inria.hal.science/hal-04880388 (cit. on p. 17).

[24] E. Anceaume, D. Frey and A. Rauch. 'Brief: Sharding in permissionless systems in presence of an adaptive adversary'. In: *31st International Colloquium on Structural Information and Communication Complexity (SIROCCO)*. 31st International Colloquium on Structural Information and Communication Complexity (SIROCCO). Vol. 14662. Lecture Notes in Computer Science. Vietri sul Mare, Italy, 23rd May 2024, pp. 481–487. DOI: 10.1007/978-3-031-60603-8_26. URL: https://hal.science/hal-04477243 (cit. on p. 18).

[25] E. Anceaume, D. Frey and A. Rauch. 'Sharding in permissionless systems in presence of an adaptive adversary'. In: *Proceedings of the International Conference on Networked Systems (NETYS)*. NETYS 2024 - 12th International Conference on Networked Systems. Rabat, Morocco: Springer, 2024, pp. 1–30. URL: https://cnrs.hal.science/hal-04794826 (cit. on p. 18).

[26] D. Bencivenga, G. Giakkoupis and P. Woelfel. 'Faster Randomized Repeated Choice and DCAS'. In: *Proceedings of the 43rd ACM Symposium on Principles of Distributed Computing*. PODC '24: 43rd ACM Symposium on Principles of Distributed Computing. Nantes, France: ACM, 17th June 2024, pp. 454–464. DOI: 10.1145/3662158.3662828. URL: https://inria.hal.science/hal-04896495 (cit. on p. 19).

[27] Y.-D. Bromberg, J. Decouchant, M. Sourisseau and F. Taïani. 'Partition Detection in Byzantine Networks'. In: ICDCS 2024 - IEEE 44th International Conference on Distributed Computing Systems. Jersey City, NJ, United States: IEEE, 2024, pp. 139–150. DOI: 10.1109/ICDCS60910.2024.00022. URL: https://inria.hal.science/hal-04677829 (cit. on p. 18).

[28] M. Déprés, A. Mostefaoui, M. Perrin and M. Raynal. 'What are the Relationships between Read/Write and Send/Receive in Crash-Prone Asynchronous Systems'. In: *AlgoTel 2024 – 26èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications*. AlgoTel 2024 – 26èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications. Saint-Briac-sur-Mer, France, 2024, pp. 1–4. URL: https://hal.science/hal-04555445.

[29] A. Durand, M. Raynal and G. Taubenfeld. 'Better Sooner Rather Than Later'. In: SIROCCO 2024 - 31st International Colloquium Structural Information and Communication Complexity. Vol. 14662. Lecture Notes in Computer Science. Vietri sul Mare, Italy: Springer Nature Switzerland, 23rd May 2024, pp. 226–237. DOI: 10.1007/978-3-031-60603-8_13. URL: https://uca.hal.science/hal-04685214.

[30] A. Durand, M. Raynal and G. Taubenfeld. 'Mieux vaut tôt que jamais'. In: *AlgoTel 2024 – 26èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications*. AlgoTel 2024 – 26èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications. Saint-Briac-sur-Mer, France, 2024, pp. 1–4. URL: https://uca.hal.science/hal-04554838.

[31] D. Frey, M. Gestin and M. Raynal. 'Consensus number du contrôle d'accès: le cas des AllowLists et DenyLists'. In: *AlgoTel 2024 – 26èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications*. AlgoTel 2024 – 26èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications. Saint-Briac-sur-Mer, France, 2024, pp. 1–4. URL: https://hal.science/hal-04551351.

[32] 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. URL: https://hal.science/hal-04888959 (cit. on p. 21).

[33] G. Giakkoupis, V. Turau and I. Ziccardi. 'Self-Stabilizing MIS Computation in the Beeping Model'. In: *Leibniz International Proceedings in Informatics (LIPIcs)*. 38th International Symposium on Distributed Computing (DISC 2024). Vol. 319. Madrid, Spain, 2024, 28:1–28:21. DOI: 10.4230/LIPIcs.DISC.2024.28. URL: https://inria.hal.science/hal-04896494 (cit. on p. 19).

[34] A. Godinot, E. Le Merrer, G. Trédan, C. Penzo and F. Taïani. 'Under manipulations, are some AI models harder to audit?' In: 2nd IEEE Conference on Secure and Trustworthy Machine Learning (SaTML 2024). Toronto (CA), Canada: IEEE, 14th Feb. 2024, pp. 1–21. URL: https://laas.hal.science/hal-04800332 (cit. on p. 22).

[35] E. Le Merrer and G. Trédan. 'LLMs hallucinate graphs too: a structural perspective'. In: *complex networks*. complex networks 2024. Istanbul, Turkey: Springer, 10th Dec. 2024. URL: https://hal.science/hal-04684742 (cit. on p. 22).

[36] D. Mvondo, T. Xing and A. Barbalace. 'UTwinVM: Reliable hints on the effects of hypervisor updates on VMs in the Cloud'. In: *25th ACM/IFIP International Middleware Conference*. Middleware '24: 25th International Middleware Conference. Hong Kong, Hong Kong SAR China: ACM, 2nd Dec. 2024, pp. 103–116. DOI: 10.1145/nnnnnnn.nnnnnnn. URL: https://hal.science/hal-04894609.

[37] D. B. T. Mvondo Djob, F. Taiani and Y.-D. Bromberg. 'HORSE: Ultra-low latency workloads on FaaS platforms'. In: *25th ACM/IFIP International Middleware Conference*. Middleware '24: 25th International Middleware Conference. Hong Kong, Hong Kong SAR China: ACM, 2nd Dec. 2024, pp. 445–453. DOI: 10.1145/3652892.3700784. URL: https://hal.science/hal-04894549 (cit. on pp. 23, 24).

[38] R. Raes, A. Luxey-Bitri, R. Rouvoy, D. Frey and F. Taïani. 'Venice: eschewing the cloud by leveraging local communication channels'. In: ICT4S 2024 - International Conference on Information and Communications Technology for Sustainability. Stockholm, Sweden, 2024, pp. 1–4. URL: https://hal.science/hal-04576743.

[39] M. de Vos, A. Dhasade, J. Garcia Bourrée, A.-M. Kermarrec, E. Le Merrer, B. Rottembourg and G. Trédan. 'Fairness Auditing with Multi-Agent Collaboration'. In: 27th European Conference on Artificial Intelligence (ECAI 2024). Frontiers in Artificial Intelligence and Applications. Santiago de Compostela, Spain: IOS Press, 16th Oct. 2024, pp. 1–14. DOI: 10.3233/FAIA240604. URL: https://laas.hal.science/hal-04800328 (cit. on p. 23).

National peer-reviewed Conferences

[40] M. de Vos, A. Dhasade, J. G. Bourrée, A. M. Kermarrec, E. Le Merrer, B. Rottembourg and G. Trédan. 'Auditer l'équité : l'union fait-elle la force ?' In: *AlgoTel 2024 – 26èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications*. AlgoTel 2024 – 26èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications. Saint-Briac-sur-Mer, France, May 2024, pp. 1–4. URL: https://hal.science/hal-04565809.

Conferences without proceedings

[41] L. S. Luevano, Y. Martínez-Díaz, H. Méndez-Vázquez, M. Gonzalez-Mendoza and D. Frey. 'Assessing the Performance of Efficient Face Anti-Spoofing Detection Against Physical and Digital Presentation Attacks'. In: FAS 2024 - 5th Face Anti-spoofing Workshop and Challenge workshop @ CVPR. Seattle, United States, June 2024, pp. 1–8. URL: https://hal.science/hal-04610076 (cit. on p. 21).

Scientific book chapters

[42] E. Le Merrer, C. Penzo, G. Tredan and L. Verney. 'Challenges in archiving the personalized web'. In: *Exploring the Archived Web during Transformative Age*. 3rd Sept. 2024, pp. 1–16. DOI: 10.36253/fup_best_practice. URL: https://hal.science/hal-04881990 (cit. on p. 22).

[43] E. Le Merrer, C. Penzo, G. Tredan and L. Verney. 'Challenges in archiving the personalized web'. In: *Exploring the Archived Web during a Highly Transformative Age*. 3rd Sept. 2024, pp. 1–16. DOI: 10.36253/979-12-215-0413-2.10. URL: https://hal.science/hal-04685057.

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

[44] *SwiftFaceFormer: An Efficient and Lightweight Hybrid Architecture for Accurate Face Recognition Applications*. 27th International Conference on Pattern Recognition (ICPR) 2024. Vol. 15314. Lecture Notes in Computer Science. Springer Nature Switzerland, 2nd Dec. 2024, pp. 244–258. DOI: 10.1007/978-3-031-78341-8_16. URL: https://inria.hal.science/hal-04393675 (cit. on p. 21).

Doctoral dissertations and habilitation theses

[45] T. Albouy. 'Foundations of Reliable Cooperation under Asynchrony, Byzantine Faults, and Message Adversaries'. Université de Rennes, 16th Dec. 2024. URL: https://inria.hal.science/tel-04764046 (cit. on p. 15).

[46] M. Gestin. 'Privacy Preserving and fully-Distributed Identity Management Systems'. Universite de rennes 1, 18th Dec. 2024. URL: https://hal.science/tel-04808780 (cit. on p. 16).

[47] A. Rauch. 'Towards more scalable and privacy-preserving distributed assettransfer systems'. Université de Rennes, 18th Dec. 2024, pp. 481–487. URL: https://inria.hal.science/tel-04857796 (cit. on p. 16).

Reports & preprints

[48] 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*. 15th May 2024. URL: https://inria.hal.science/hal-04578985.

[49] E. Bannier, S. Castellan, S. Derrien, F. Galassi, L. Garnier, L. Hoyet, A. l'Azou, N. Lahaye, M. J.-M. Macé, O. Martineau, A. Masson, T. Maugey, B. Ninassi, E. Rohou, M. Simonin and F. Taïani. *Reducing GHG emissions from business travel: A collaborative approach at IRISA/Inria*. Groupe de travail « missions » IRISA / Centre Inria de l'Université de Rennes, Mar. 2024, pp. 1–16. URL: https://univ-rennes.hal.science/hal-04506138.

[50] L. S. Luevano, L. Chang, M. González-Mendoza, Y. Martínez-Díaz, H. Méndez-Vázquez and G. Ochoa-Ruiz. *BinaryFaceNet: A Binarized Approach for Real-Time Very Low Resolution Face Recognition in Video Surveillance Scenarios*. 14th Jan. 2024. URL: https://inria.hal.science/hal-04393666.

[51] L. S. Luevano and D. Frey. *Analyzing Trusted Execution Environments: Comparing Commercial Implementations and Diverse Applications*. 14th Jan. 2024. URL: https://inria.hal.science/hal-04393667.

[52] L. S. Luevano, D. Frey, M. Sel and D. Singelee. *SOTERIA D5.4 HARDWARE-BASED PRIVACY*. 14th Jan. 2024. URL: https://inria.hal.science/hal-04393670.

[53] G. Richardeau, E. Le Merrer, C. Penzo and G. Trédan. *THE 20 QUESTIONS GAME TO DISTINGUISH LARGE LANGUAGE MODELS*. 2024. URL: https://hal.science/hal-04699271.

Other scientific publications

[54] R. Raes, R. Rouvoy, A. Luxey-Bitri, D. Romero, D. Frey and F. Taïani. 'Eschewing the Cloud by leveraging Local Communication Channels: Angus & Bob want to share cool things'. In: ICT4S 2024 - International Conference on ICT for Sustainability. Stockholm, Sweden, 2024, pp. 1–1. URL: https://hal.science/hal-04633670.

12.3 Cited publications

[55] Y. Afek and E. Gafni. 'Asynchrony from synchrony'. In: *ICDCN*. 2013, pp. 225–239 (cit. on p. 11).

[56] 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. URL: http://dx.doi.org/10.1109/ISCO.2016.7727082 (cit. on p. 6).

[57] T. Allard, D. Frey, G. Giakkoupis and J. Lepiller. 'Lightweight Privacy-Preserving Averaging for the Internet of Things'. In: *M4IOT 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. URL: https://hal.inria.fr/hal-01421986 (cit. on pp. 9, 12).

[58] 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. 12).

[59] 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. 10).

[60] 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. URL: http://dx.doi.org/10.1109/MCC.2014.51 (cit. on p. 7).

[61] 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. URL: https://hal.inria.fr/inria-00515693 (cit. on p. 6).

[62] F. Bonomi. *Connected vehicles, the internet of things, and fog computing. VANET 2011, 2011*. Keynote speech at VANET. 2011 (cit. on p. 6).

[63] 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. URL: http://doi.acm.org/10.1145/2342509.2342513 (cit. on p. 6).

[64] 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. 8).

[65] 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. URL: https://hal.inria.fr/hal-01080016 (cit. on p. 6).

[66] 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. URL: https://hal.archives-ouvertes.fr/hal-01171492 (cit. on p. 8).

[67] 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. 8).

[68] B. Cohen. *Incentives Build Robustness in BitTorrent*. 2003. URL: http://citeseer.ist.psu.edu/cohen03incentives.html (cit. on p. 8).

[69] 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. URL: https://doi.org/10.1109/DSN48063.2020.00023 (cit. on p. 11).

[70] A. Dash, A. Mukherjee and S. Ghosh. 'A Network-centric Framework for Auditing Recommendation Systems'. In: *IEEE Conference on Computer Communications, INFOCOM*. 2019 (cit. on p. 12).

[71] 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. 11).

[72] 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. 9).

[73] 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. 6).

[74] 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. 10).

[75] 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. URL: https://hal.inria.fr/inria-00505268 (cit. on p. 8).

[76] 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. URL: https://doi.org/10.1145/3293611.3331589 (cit. on p. 11).

[77] 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. 9).

[78] 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. URL: https://arxiv.org/abs/2109.14189 (cit. on p. 11).

[79] 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. 11).

[80] 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. 11).

[81] 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. 9).

[82] 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. 9).

[83] 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. 9).

[84] P. Kuznetsov et al. 'Understanding non-uniform failure models'. In: *Bulletin of the EATCS* 106 (2012), pp. 53–77 (cit. on p. 11).

[85] E. Lieberman, C. Hauert and M. Nowak. 'Evolutionary dynamics on graphs'. In: *Nature* 433.7023 (2005), pp. 312–316 (cit. on p. 9).

[86] D. Malkhi and M. Reiter. 'Byzantine quorum systems'. In: *Distributed computing* 11.4 (1998), pp. 203–213 (cit. on p. 12).

[87] 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. URL: https://doi.org/10.1109/DSN48063.2020.00037 (cit. on p. 12).

[88] 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. URL: https://doi.org/10.1109/INFOCOM.2019.8737548 (cit. on p. 12).

[89] F. Pasquale. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard U. Press, 2015 (cit. on p. 12).

[90] 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. URL: https://hal.inria.fr/hal-00920734 (cit. on p. 11).

[91] P. de Rosa, Y.-D. Bromberg, P. Felber, D. B. T. Mvondo Djob and V. Schiavoni. 'On the Cost of Model-Serving Frameworks: An Experimental Evaluation'. In: *2024 IEEE International Conference on Cloud Engineering (IC2E)*. 2024 IEEE International Conference on Cloud Engineering (IC2E). IEEE Computer Society. Paphos, Cyprus: IEEE, Sept. 2024, pp. 221–232. DOI: 10.1109/IC2E61754.2024.00032. URL: https://hal.science/hal-04894045 (cit. on p. 24).

[92] 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. 11).

[93] 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. 7).

[94] 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. URL: http://doi.acm.org/10.1145/2465351.2465379 (cit. on p. 6).