



Activity Report 2022

Team WIDE

the World Is Distributed: Exploring the tension between
scale and coordination

Joint team with Centre Inria de l'Université de Rennes

D1 – Large Scale Systems



Contents

Project-Team WIDE	1
1 Team members, visitors, external collaborators	2
2 Overall objectives	3
2.1 Overview	3
2.2 Planetary-Scale Geo-Distributed Systems	3
2.3 Highly Personalized On-Line Services	3
2.4 Social Collaboration Platforms	4
3 Research program	5
3.1 Overview	5
3.2 Hybrid Scalable Architectures	5
3.3 Personalizable Privacy-Aware Distributed Systems	7
3.4 Network Diffusion Processes	8
3.5 Systemizing Modular Distributed Computability and Efficiency	9
3.6 Evolution of our research program (2022-2026)	10
4 Application domains	12
5 Social and environmental responsibility	12
6 Highlights of the year	12
6.1 Awards	12
6.2 Dissemination	12
6.3 Major events	13
7 New software and platforms	13
7.1 New software	13
7.1.1 Basalt	13
7.1.2 Killerdroid-packer	13
7.1.3 python-andromak	13
7.1.4 Donar	14
8 New results	14
8.1 Distributed Algorithms	14
8.1.1 A Modular Approach to Construct Signature-Free BRB Algorithms under a Message Adversary	14
8.1.2 Reaching Consensus in the Presence of Contention-Related Crash Failures	14
8.1.3 Good-case Early-Stopping Latency of Synchronous Byzantine Reliable Broadcast: The Deterministic Case	15
8.1.4 A visit to mutual exclusion in seven dates	15
8.1.5 Distributed computability: Relating k-immediate snapshot and x-set agreement	15
8.1.6 Optimal Algorithms for Synchronous Byzantine k-Set Agreement	16
8.1.7 Self-stabilizing Byzantine Fault-Tolerant Repeated Reliable Broadcast	16
8.1.8 Election in Fully Anonymous Shared Memory Systems: Tight Space Bounds and Algorithms	17
8.2 Privacy Preserving and Attack-Resilient Systems	17
8.2.1 Hidden Issuer Anonymous Credential	17
8.2.2 Donar: Anonymous VoIP over Tor	17
8.2.3 RAPTEE: Leveraging trusted execution environments for Byzantine-tolerant peer sampling services	18
8.2.4 Odile: A scalable tracing tool for non-rooted and on-device Android phones	18
8.3 Network and Graph Algorithms	18
8.3.1 Expanders via Local Edge Flips in Quasilinear Time	18

8.3.2	Distributed Self-Stabilizing MIS with Few States and Weak Communication	19
8.4	Scaling and Understanding AI systems	19
8.4.1	Randomized Smoothing under Attack: How Good is it in Practice?	19
9	Partnerships and cooperations	20
9.1	International initiatives	20
9.1.1	Associate Teams in the framework of an Inria International Lab or in the framework of an Inria International Program	20
9.2	International research visitors	21
9.2.1	Visits of international scientists	21
9.2.2	Visits to international teams	21
9.2.3	H2020 projects	21
9.3	National initiatives	22
9.4	Regional initiatives	23
10	Dissemination	23
10.1	Promoting scientific activities	23
10.1.1	Scientific events: organisation	23
10.1.2	Scientific events: selection	24
10.1.3	Journal	24
10.1.4	Invited talks	25
10.1.5	Leadership within the scientific community	25
10.1.6	Scientific expertise	26
10.1.7	Research administration	26
10.2	Teaching - Supervision - Juries	26
10.2.1	Teaching	26
10.2.2	Supervision	27
10.2.3	Juries	28
10.3	Popularization	28
10.3.1	Internal or external Inria responsibilities	28
10.3.2	Articles and contents	28
10.3.3	Interventions	28
11	Scientific production	28
11.1	Major publications	28
11.2	Publications of the year	29
11.3	Other	32
11.4	Cited publications	32

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
- A2.6.2. – 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

- B6.1.1. – Software engineering
- 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, HDR]
- Georgios Giakkoupis [INRIA, Researcher]
- Erwan Le Merrer [INRIA, Senior Researcher, HDR]
- Michel Raynal [UNIV RENNES I, Emeritus, HDR]

Faculty Members

- François Taiani [Team leader, UNIV RENNES I, Professor, until Aug 2022, HDR]
- David Bromberg [UNIV RENNES I, Professor, HDR]
- Barbe Mvondo Djob [UNIV RENNES I, Associate Professor]

PhD Students

- Timothé Albouy [UNIV RENNES I]
- Jade Garcia Bourrée [INRIA, from Oct 2022]
- Mathieu Gestin [INRIA]
- Augustin Godinot [UNIV RENNES I, from Nov 2022]
- Honore Cesaire Mounah [UNIV RENNES I, from Dec 2022]
- Arthur Rauch [INRIA]

Technical Staff

- Muhammed Selcuk Kok [INRIA]
- Mathieu Simonin [INRIA]
- Ali Yesilkanat [INRIA]

Interns and Apprentices

- Honore Cesaire Mounah [UNIV RENNES I, from Jun 2022 until Oct 2022]
- Isabella Ziccardi [Università dell'Aquila, Italy, from Feb 2022 until May 2022, Supervised by George Giakkoupis]

Administrative Assistant

- Virginie Desroches [INRIA]

2 Overall objectives

2.1 Overview

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

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

2.2 Planetary-Scale Geo-Distributed Systems

Modern large-scale systems often encompass thousands of server nodes, hosted in tens of datacenters distributed over several continents. To address the challenges posed by such systems, alternative distributed architectures are today emerging that emphasize *decentralized* and *loosely coupled* interactions. This evolution can be observed at multiple levels of an application's distributed stack: the growing interest, both practical and theoretical, for weak consistency models is such an example. In spite of their potential counter-intuitive behaviors, weakly consistent data-structures allow developers to trade strict coordination guarantees for the ability to deliver a reactive and scalable service even when hit by arbitrary network delays or system partitions. At a higher, more architectural level, similar motivations explain the push for *micro-services* on the server side of on-line applications and the growth of rich *browser-based programming technologies* on their client side. Micro services help development teams decompose complex applications into a set of simpler and loosely-connected distributed services. In a parallel evolution, modern browsers embark increasingly powerful networking APIs such as WebRTC. These APIs are prompting a fresh rethink of the typical distribution of capabilities between servers and clients. This is likely to lead to more services and computations being offloaded to browsers, in particular within hybrid architectures. The above evolutions, away from tightly synchronized and monolithic deployments towards heterogeneous, composite and loosely coordinated distributed systems, raise a number of difficult challenges at the crossroad of theoretical distributed algorithms, system architecture, and programming frameworks. One of these challenges pertains to the growing complexity arising from these systems: as richer and more diverse services are being composed to construct whole applications, individual developers can only hope to grasp parts of the resulting systems. Similarly, weak consistency models and loose coordination mechanisms tend to lead to counter-intuitive behaviors, while only providing weak overall guarantees. This lack of systematic guarantees and understandability make it harder for practitioners to design, deploy, and validate the distributed systems they produce, leading to rising costs and high entry barriers.

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

2.3 Highly Personalized On-Line Services

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

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

data-mining and machine learning techniques to deliver enhanced, personalized, services to users and companies, and often exploit highly parallelized data analytics frameworks such as Spark, and Flink.

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

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

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

2.4 Social Collaboration Platforms

On-line social networks have had a fundamental and lasting impact on the Internet. In recent years, numerous applications have appeared that go beyond the services originally provided by “pure” on-line social networks, such as posting messages or maintaining on-line “friendship” links. These new applications seek to organize and coordinate users, often in the context of the sharing economy, for instance in order to facilitate car-sharing (e.g. BlaBla car, www.blablacar.com), short-term renting (e.g. AirBnB, www.airbnb.com), and peer-to-peer financial services (e.g. Lending Club, www.lendingclub.com). Some systems, such as Bitcoin or Ethereum, have given rise to new distributed protocols combining elements of cryptography and distribution that are now largely discussed in the research community, and have attracted the attention of policy makers and leading financial actors.

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

Formalizing and characterizing the needs and behaviors of these new applications seems particularly interesting in order to provide the fertile ground for new systems and novel theoretical work. The area of social applications also offers avenues for knowledge transfer and societal impact, along two dimensions. First, practical and usable approaches, back by a deep understanding of the foundation of distribution and coordination, are likely to find applications in future systems. Second, developers of complex social applications are often faced with a lack of robust scalable services¹ that can be easily exploited to harness the latest understanding of large-scale distributed coordination. We therefore think these applications offer an opportunity to design and deliver modular reusable bricks that can be easily appropriated by a large population of innovative developers without requiring the level of deep understanding usually necessary to implement these solutions from scratch. Providing such reusable bricks is however difficult,

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

as many interesting formal properties are not composable, and a unified composable theory of distributed systems still need to be fully articulated.

3 Research program

3.1 Overview

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

More specifically, our program revolves around four key challenges.

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

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

The above four challenges explore these tensions from *four complementary angles*: from an architectural perspective (Challenge 1), from the point of view of a fundamental system-wide guarantee (privacy protection, Challenge 2), looking at one universal scalable mechanism (network diffusion, Challenge 3), and considering the interplay between modularity and computability in large-scale systems (Challenge 4). These four challenges range from practical concerns (Challenges 1 and 2) to more theoretical questions (Challenges 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 [49, 61] 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) [53]. 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 [50] 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 [44].

A number of efforts in this directions focus on specific hardware, e.g. fog nodes that are responsible for connected IoT devices [51]. 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 [82], 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 [81], Kubernetes [48]) 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 [52, 54, 55] 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 [56, 63]. 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 [45].

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 [71, 69, 70] that have been studied in this context so far. This includes the *randomized rumor spreading (RS)* model for information dissemination [60], *biased* versions of the *voter model* [65] modelling influence, and the (graph-based) *Moran processes* [73] 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 [71], 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 [47]. This strand of work will consider restricted graph families, such as trees or graphs of small treewidth, for such problems. In addition, for some immunization problems, there is a large gap between the best known approximation algorithm and the best known inapproximability result, and we would like to make progress in reducing these gaps.

3.5 Systemizing Modular Distributed Computability and Efficiency

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

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

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

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

In particular, looking at the problem of *mutual exclusion*, it is desirable that mutual exclusion algorithms be *abortable*. This means that a process that is trying to lock the resource can abort its attempt in case it has to wait too long. Abortability is difficult to achieve for mutual exclusion algorithms. We will try to extend our algorithms for the *cache-coherent* (CC) and the *distributed shared memory* (DSM) model in order to make them abortable, while maintaining expected constant *Remote Memory References* (RMRs) complexity, under optimistic system assumptions. In order to achieve this, the algorithm will use strong synchronization primitives, called compare-and-swap objects. As part of our collaboration with the University of Calgary, we will work on implementing those objects from registers in such a way that they also allow aborts. Our goal is to build on existing non-abortable implementations [62]. 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 [79]) and *process adversaries* (investigated in several papers, e.g. [78, 59, 67, 68, 72]). 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 [78, 43].

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 1) 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 under-

standing 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 [66]. 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 [57, 64]. In the case of open systems, we will leverage the expertise of our team in theoretical distributed algorithms and randomized algorithms to address Sybil attacks through novel countermeasures providing (hopefully) cheaper and more equitable alternatives to proof-of-work or proof-of-stake algorithms. One open, yet enticing, question is whether anonymous computing models could provide a path to address this issue. We would also like to investigate how storage can be improved in Blockchains and BFT large-scale systems. Most of these systems are fully replicated, incurring formidable costs (up to 2.6PB of distributed storage in the case of Bitcoin). Coding techniques, that we have used in the past, and adaptable redundancy based on Byzantine quorums [74] 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 [45], 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 [75, 76]. 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) [58]. These algorithms, often coined *black-box algorithms* [77], 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 [46], or a k-nearest neighbor graph), by means of local decentralized operations (typically adjacent nodes interact by exchanging some of their incident edges). In all three cases, we will strive to derive time- and space- optimal algorithms, with strong robustness guarantees.

4 Application domains

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

5 Social and environmental responsibility

- Davide Frey and Francois Taiani 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

6.1 Awards

Jade Garcia Bourrée and Augustin Goudinot, ranked second at a hackathon on black-box algorithms, that took place in December 2022. The hackathon was part of the conference on "Possibilités technologiques, limites juridiques : les algorithmes face à la régulation".

6.2 Dissemination

Michel Raynal issued a new book on "Concurrent Crash-Prone Shared Memory Systems", published in the collection "Theoretical Notions Synthesis Lectures on Distributed Computing Theory" published par Morgan & Claypool.

Erwan Le Merrer was interviewed by Le Media, as well as by Binaire (blog from Le Monde, reproduced by The Conversation and the Journal Du Dimanche) about the need to audit algorithms, and in particular the recommendation algorithm from YouTube.

6.3 Major events

David Bomberg acted as the co-chair for the Eurosys 2022 conference in Rennes, supported by the WIDE team for multiple tasks. The conference showcased an audience success with 340 attendees.

7 New software and platforms

7.1 New software

7.1.1 Basalt

Keywords: Peer-sampling, Blockchain

Functional Description: A number of novel blockchain and cryptocurrency implementations rely on random peer sampling. But existing protocols remain vulnerable to Sybil attacks. BASALT is a peer-sampling protocol that addresses this limitation by leveraging three main components. First it employs a novel sampling approach, termed stubborn chaotic search, that exploits ranking functions to define a dynamic target random graph (i.e. a set of target neighbors for each node) that cannot be controlled by Byzantine nodes. Second, it adopts a hit-counter mechanism that favors the exploration of new peers even in the presence of Byzantine nodes that flood the network with their identities. Finally, it incorporates hierarchical ranking functions that ensure that nodes sample their peers from a variety of address prefixes. The first two mechanisms ensure that the number of Byzantine nodes in a node's view cannot be increased arbitrarily by attackers. This offers protection from general Byzantine behaviors including those resulting from botnet attacks, as defined above. The third mechanism ensures that nodes sample their peers from a variety of address prefixes, thereby countering institutional attacks where the attacker controls a limited number of entire address prefixes.

Contact: Alex Auvolat-Bernstein

7.1.2 Killerdroid-packer

Name: Practical implementation of dissimulation techniques for Android application

Keywords: Malware, Android, Machine learning, Cybersecurity

Functional Description: Killerdroid-packer allows to intricate a Android guest application in a Android host application. A common use case is to hide a malicious code inside a benign one to evade typical detection systems. Killerdroid-packer has been proven effective and scalable. Applications generated with the tool can evade state-of-the art detection systems (Drebin, Mamadroid ...) and can be used to forge large scale datasets of malicious applications.

Publication: [hal-03146161](#)

Contact: David Bromberg

Participants: David Bromberg, Mathieu Simonin, Louison Gitzinger

Partner: DGA-MI

7.1.3 python-andromak

Name: Andromak: A library to assess Android detector quality

Keywords: Android, Anomaly detection, Cybersecurity, Distribution, Malware, Machine learning

Functional Description: State of the Art detectors are numerous, but actually comparing them in a thorough way is tedious. Andromak aims to make a step toward easier and sound comparisons between the top existing detectors.

Python-andromak is a library to: - (a) scale your training, - (b) evaluate the quality of your model against several input datasets and scenario, - (c) analyse the performance results (domain-specific metrics and statistical tests are shipped).

Contact: David Bromberg

7.1.4 Donar

Keywords: Audio, Anonymisation

Functional Description: Make anonymous calls via Tor

Publication: [hal-03923695](https://hal.archives-ouvertes.fr/hal-03923695)

Contact: Quentin Dufour

Participants: Quentin Dufour, David Bromberg, Davide Frey, Etienne Riviere

8 New results

8.1 Distributed Algorithms

8.1.1 A Modular Approach to Construct Signature-Free BRB Algorithms under a Message Adversary

Participants: Timothé Albouy, Davide Frey, Michel Raynal, Francois Taiani.

This work [21] explores how reliable broadcast can be implemented without signatures when facing a dual adversary that can both corrupt processes and remove messages. More precisely, we consider an asynchronous n -process message-passing system in which up to t processes are Byzantine and where, at the network level, for each message broadcast by a correct process, an adversary can prevent up to d processes from receiving it (the integer d defines the power of the message adversary). So, unlike previous works, this work considers that not only can computing entities be faulty (Byzantine processes), but, in addition, that the network can also lose messages. To this end, this work adopts a modular strategy and first introduces a new basic communication abstraction denoted k_2 -cast, which simplifies quorum engineering, and studies its properties in this new adversarial context. Then, this work deconstructs existing signature-free Byzantine-tolerant asynchronous broadcast algorithms and, with the help of the k_2 -cast communication abstraction, reconstructs versions of them that tolerate both Byzantine processes and message adversaries. Interestingly, these reconstructed algorithms are also more efficient than the Byzantine-tolerant-only algorithms from which they originate.

8.1.2 Reaching Consensus in the Presence of Contention-Related Crash Failures

Participants: Michel Raynal.

While consensus is at the heart of many coordination problems in asynchronous distributed systems prone to process crashes, it has been shown to be impossible to solve in such systems where processes communicate by message-passing or by reading and writing a shared memory. Hence, these systems must be enriched with additional computational power for consensus to be solved on top of them. This work [25][18] presents a new restriction of the classical basic computational model that combines process participation and a constraint on failure occurrences that can happen only while a predefined contention threshold has not yet been bypassed. This type of failure is called λ -constrained crashes, where λ defines the considered contention threshold. It appears that when assuming such contention-related crash failures and enriching the system with objects whose consensus number is $k \geq 1$, consensus

for n processes can be solved for any $n \geq k$ assuming up to k failures. The work proceeds incrementally. It first presents an algorithm that solves consensus on top of read/write registers if at most one crash occurs before the contention threshold $\lambda = n - 1$ has been bypassed. Then, it shows that if the system is enriched with objects whose consensus number is $k \geq 1$, then when $\lambda = n - k$, consensus can be solved despite up to k λ -constrained crashes, for any $n \geq k$, and when $\lambda = n - 2k + 1$, consensus can be solved despite up to $2k - 1$ λ -constrained crashes, assuming k divides n . Finally, impossibility results are presented for the number of λ -constrained failures that can be tolerated.

This is a joint work with Anaïs Durand (LIMOS, Clermont-Ferrand) and Gadi Taubenfeld (Reichman University, Israel)

8.1.3 Good-case Early-Stopping Latency of Synchronous Byzantine Reliable Broadcast: The Deterministic Case

Participants: Timothé Albouy, Davide Frey, Michel Raynal, Francois Taiani.

This work [22] 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, and an essential property for 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 under a Byzantine majority. In this work, we answer this open question positively and propose a deterministic synchronous Byzantine reliable broadcast that achieves a good-case latency of $\max(2, t + 3 - c)$ rounds, where t is the upper bound on the number of Byzantine processes, and c the number of effectively correct processes.

8.1.4 A visit to mutual exclusion in seven dates

Participants: Michel Raynal.

Mutual exclusion (mutex) is one of the most fundamental synchronization problems encountered in shared memory systems. It appears in all computer science first-degree curricula. This work [20] presents nine mutex algorithms, each with its noteworthy features, spread over seven dates covering 1965-2020. Most of these algorithms are very well known and paved the way for new research directions. This work aims to present fundamental issues and basic principles that underlie the design of shared memory mutex algorithms in different contexts. So, differently from exhaustive surveys on shared memory mutex algorithms, it strives to give the reader a flavor of the many design facets of this still challenging problem.

This is a joint work with Gadi Taubenfeld (Reichman University, Herzliya 46150, Israel).

8.1.5 Distributed computability: Relating k -immediate snapshot and x -set agreement

Participants: Michel Raynal.

This work [17] introduces a generalization of the immediate snapshot object denoted k -immediate snapshot, requiring that the snapshot returned contains at least $(n - k)$ pairs. The case $k = n - 1$ corresponds to the original immediate snapshot object, which requires that the snapshot returned contains at least one $\langle \text{process id}, \text{value} \rangle$ pair, that corresponds to the process id that invoked the operation).

The work first shows that k -immediate snapshot is impossible to implement in an asynchronous read/write system, even if $k = n - 2$ and $t = 1$. Then, this work considers x -set agreement, another object stronger than the classical read/write t -crash read/write model (when $x \leq t$), and studies the relation with

the k -immediate snapshot object, establishing strong relations linking these two fundamental distributed computing abstractions. This work shows conditions under which x -set agreement can be solved in read/write systems enriched with k -immediate snapshot objects. It also shows when k -immediate snapshot and consensus are equivalent.

This is a joint work with Carole Delporte (IRIF, Université Paris Diderot, Paris, France), Hugues Fauconnier (IRIF, Université Paris Diderot, Paris, France), and Sergio Rajsbaum (Instituto de Matemáticas, UNAM, Mexico City, Mexico).

8.1.6 Optimal Algorithms for Synchronous Byzantine k -Set Agreement

Participants: Michel Raynal.

Considering a system made up of n processes prone to Byzantine failures, k -set agreement allows each process to propose a value and decide a value such that at most k different values are decided by the correct (i.e., non-Byzantine) processes, in such a way that, if all the correct processes propose the same value v , they will decide v (when $k = 1$, k -set agreement boils down to consensus). This work [24] presents a two-round algorithm that solves Byzantine k -set agreement on top of a synchronous message-passing system. This algorithm is based on two new notions denoted by Square and Regions which allow processes to locally build a global knowledge on which processes proposed some values. Two instances of the algorithm are presented. Assuming $n = 3t$, where t is the maximum number of Byzantine, the first instance solves 2-set agreement. The second one solves the more general case $2t < n \leq 3t$, where $k = n - tn - 2t$ is an integer. These two algorithm instances are optimal with respect to the number of rounds executed by the processes (namely two rounds). Combined with previous results, this work “nearly closes” the solvability of Byzantine k -set agreement in synchronous message-passing systems (more precisely, the only remaining case for which it is not known whether k -set agreement can or cannot be solved is when $k = n - tn - 2t$ is not an integer).

This is a joint work with Carole Delporte (IRIF, Université Paris Diderot, Paris, France), Hugues Fauconnier (IRIF, Université Paris Diderot, Paris, France), and Mouna Safir (IRIF, Université Paris Cité, Paris, France, and School of Computer Sciences, Mohammed VI Polytechnic University, Ben Guerir, Morocco).

8.1.7 Self-stabilizing Byzantine Fault-Tolerant Repeated Reliable Broadcast

Participants: Michel Raynal.

In this work [32], we study a well-known communication abstraction called Byzantine Reliable Broadcast (BRB). This abstraction is central in the design and implementation of fault-tolerant distributed systems, as many fault-tolerant distributed applications require communication with provable guarantees on message deliveries. Our study focuses on fault-tolerant implementations for message-passing systems that are prone to process-failures, such as crashes and malicious behaviors.

At PODC 1983, Bracha and Toueg, in short, BT, solved the BRB problem. BT has optimal resilience since it can deal with up to $5 < n/3$

Byzantine processes, where n is the number of processes. The present work aims at the design of an even more robust solution than BT by expanding its fault-model with self-stabilization, a vigorous notion of fault-tolerance. In addition to tolerating Byzantine and communication failures, self-stabilizing systems can recover after the occurrence of arbitrary transient-faults. These faults represent any violation of the assumptions according to which the system was designed to operate (as long as the algorithm code remains intact).

We propose, to the best of our knowledge, the first self-stabilizing Byzantine fault-tolerant (SSBFT) solution for repeated BRB (that follows BT’s specifications) in signature-free message-passing systems. Our contribution includes a self-stabilizing variation on a BT that solves asynchronous single-instance

BRB. We also consider the problem of recycling instances of single-instance BRB. Our SSBFT recycling for time-free systems facilitates the concurrent handling of a predefined number of BRB invocations and, by this way, can serve as the basis for SSBFT consensus.

This is a joint work with Romaric Duvignau and Elad M. Schiller from Chalmers University of Technology, Gothenburg, Sweden.

8.1.8 Election in Fully Anonymous Shared Memory Systems: Tight Space Bounds and Algorithms

Participants: Michel Raynal.

This work [33] addresses election in fully anonymous systems made up of n asynchronous processes that communicate through atomic read-write registers or atomic read-modify-write registers. Given an integer $d \in \{1, \dots, n-1\}$, two elections problems are considered: d -election (at least one and at most d processes are elected) and exact d -election (exactly d processes are elected). Full anonymity means that both the processes and the shared registers are anonymous. Memory anonymity means that the processes may disagree on the names of the shared registers. That is, the same register name A can denote different registers for different processes, and the register name A used by a process and the register name B used by another process can address the same shared register. Let n be the number of processes, m the number of atomic read-modify-write registers, and let $M(n, d) = \{k : \forall \ell : 1 < \ell \leq n : \gcd(\ell, k) \leq d\}$. The following results are presented for solving election in such an adversarial full anonymity context.

- It is possible to solve d -election when participation is not required if and only if $m \in M(n, d)$.
- It is possible to solve exact d -election when participation is required if and only if $\gcd(m, n)$ divides d .
- It is possible to solve d -election when participation is required if and only if $\gcd(m, n) \leq d$.
- Neither d -election nor exact d -election (be participation required or not) can be solved when the processes communicate through read-write registers only.

This is a joint work with Damien Imbs (LIS, Aix-Marseille University & CNRS & Univ. Toulon, Marseille, France) and Gadi Taubenfeld (Reichman University, Herzliya, Israel).

8.2 Privacy Preserving and Attack-Resilient Systems

8.2.1 Hidden Issuer Anonymous Credential

Participants: Davide Frey, Mathieu Gestin.

This work [14], carried out in the context of the SOTERIA H2020 project, focuses on improving the privacy-preserving properties of Anonymous Credential schemes. Anonymous Credential schemes provide unlinkability between uses of a credential, a property that is particularly useful for providing privacy in self-sovereign identity (SSI) systems. However, existing anonymous credentials systematically disclose the identity of the Issuer of a given credential to service providers, leading to information leaks. In this work, we introduced a new Anonymous Credential scheme that allows a user to hide the Issuer of a credential, while being able to convince the service providers that they can trust the credential, in the absence of a trusted setup. We proved the new scheme secure and showed that it is efficient enough to be used with laptops, and to be integrated into SSI frameworks or any other IMS. We published this work in Issue 4 of PoPETS 2022 [14], and our implementation is [available](#). This work was done in collaboration with Daniel Bosk from KTH, and Guillaume Piolle from the CIDRE Team.

8.2.2 Donar: Anonymous VoIP over Tor

Participants: David Bromberg, Davide Frey.

In this work, we introduced DONAR, a system enabling anonymous VoIP with good quality-of-experience (QoE) over Tor. No individual Tor link can match VoIP networking requirements. DONAR bridges this gap by spreading VoIP traffic over several links. It combines active performance monitoring, dynamic link selection, adaptive traffic scheduling, and redundancy at no extra bandwidth cost. DONAR enables high QoE: latency remains under 360 ms for 99% of VoIP packets during most (86%) 5-minute and 90-minute calls. This work, published at NSDI 2022 [23], was done in collaboration with former WIDE member Quentin Dufour who defended his thesis in February 2021, and Etienne Riviere from UC Louvain.

8.2.3 RAPTEE: Leveraging trusted execution environments for Byzantine-tolerant peer sampling services

Participants: David Bromberg, Davide Frey.

In this work, published at ICDCS 2022 [29], we focused on building an attack-resilient peer-sampling framework. Peer sampling is a first-class abstraction used in distributed systems for overlay management and information dissemination. The goal of peer sampling is to continuously build and refresh a partial and local view of the full membership of a dynamic, large-scale distributed system. Malicious nodes under the control of an adversary may aim at being over-represented in the views of correct nodes, increasing their impact on the proper operation of protocols built over peer sampling. State-of-the-art Byzantine resilient peer sampling protocols reduce this bias as long as Byzantines are not overly present. So, in this work, we explored the benefits brought to the resilience of peer sampling services when considering that a small portion of trusted nodes can run code whose authenticity and integrity can be assessed within a trusted execution environment, and specifically Intel’s software guard extensions technology (SGX). We proposed Raptee, a protocol that builds and leverages trusted gossip-based communications to hamper an adversary’s ability to increase its system-wide representation in the views of all nodes. This work was done in collaboration with Matthieu Pigaglio, Joachim Bruneau-Queyreix and, Laurent Réveillère from LABRI Bordeaux, as well as with Etienne Riviere from UC Louvain.

8.2.4 Odile: A scalable tracing tool for non-rooted and on-device Android phones

Participants: David Bromberg.

In this work, published at RAID 2022 [80], We introduced a new dynamic binary instrumentation tool, named Odile, to help reverse engineers to perform on-device analysis for non-rooted Android devices. Odile provides a new scalable tracing approach that we call delegated instrumentation. It leverages Android’s instrumentation module and mainly relies on ART reverse engineering. We demonstrate the effectiveness of Odile in tracing various app types (including benign apps and malware). In particular, we show how much Odile outperforms Frida, the state-of-the-art tool in the domain. Odile enables to do runtime analysis of app behavior, which is becoming paramount for reverse engineers and app market maintainers (e.g., Google Play) to ensure that running apps do not include some malicious payload. This work was done in collaboration with Lavoisier Wapet, and Alain Tchana from Grenoble INP.

8.3 Network and Graph Algorithms

8.3.1 Expanders via Local Edge Flips in Quasilinear Time

Participants: George Giakkoupis.

Mahlmann and Schindelhaue (2005) proposed the following simple process, called *flip-chain*, for transforming any given connected d -regular graph into a d -regular expander: In each step, a random 3-path $abcd$ is selected, and edges ab and cd are replaced by two new edges ac and bd , provided that ac and bd do not exist already. A motivation for the study of the flip-chain arises in the design of overlay networks, where it is common practice that adjacent nodes periodically exchange random neighbors, to maintain good connectivity properties. It is known that the flip-chain converges to the uniform distribution over connected d -regular graphs, and it is conjectured that an expander graph is obtained after $O(nd \log n)$ steps, w.h.p., where n is the number of vertices. However, the best known upper bound on the number of steps has been $O(n^2 d^2 \sqrt{\log n})$, and the best bound on the mixing time of the chain is $O(n^{16} d^{36} \log n)$.

In [26] we provide a new analysis of a natural flip-chain instantiation, which shows that starting from any connected d -regular graph, for $d = \Omega(\log^2 n)$, an expander is obtained after $O(nd \log^2 n)$ steps, w.h.p. This result is tight within logarithmic factors, and almost matches the conjectured bound. Moreover, it justifies the use of edge flip operations in practice: for any d -regular graph with $d = \text{poly}(\log n)$, an expander is reached after each vertex participates in at most $\text{poly}(\log n)$ operations, w.h.p. Our analysis is arguably more elementary than previous approaches. It uses the novel notion of the *strain* of a cut, a value that depends both on the crossing edges and their adjacent edges. By keeping track of the cut strains, we form a recursive argument that bounds the time before all sets of a given size have large expansion, after all smaller sets have already attained large expansion.

8.3.2 Distributed Self-Stabilizing MIS with Few States and Weak Communication

Participants: George Giakkoupis, Isabella Ziccardi.

In this work in progress we study a simple random process that computes a maximal independent set (MIS) on a general n -vertex graph. Each vertex has a binary state, black or white, where black indicates inclusion into the MIS. The vertex states are arbitrary initially, and are updated in parallel: In each round, every vertex whose state is “inconsistent” with its neighbors’, i.e., it is black and has a black neighbor, or it is white and all neighbors are white, flips its state independently with probability $1/2$. The process stabilizes with probability 1 on any graph, and the resulting set of black vertices is an MIS. It is easy to see that the expected stabilization time is $O(\log n)$ on certain graph families, such as cliques and trees. However, analysing the process on graphs beyond these simple cases seems challenging, which may explain why the process has not appeared in the literature before.

Our main result is that the process stabilizes in $\text{poly}(\log n)$ rounds w.h.p. on $G_{n,p}$ random graphs, for all $0 \leq p = \tilde{O}(n^{-1/2})$. Further, an extension of this process, with larger but still constant vertex state space, stabilizes in $\text{poly}(\log n)$ rounds on $G_{n,p}$ graphs w.h.p., for all $1 \leq p \leq 1$. We conjecture that this improved bound holds for the original process as well. In fact, we believe that the original process stabilizes in $\text{poly}(\log n)$ rounds on any given n -vertex graph w.h.p. Both processes readily translate into distributed/parallel MIS algorithms, which are self-stabilizing, use constant memory (and constant random bits per round), and assume restricted communication as in the beeping or the synchronous stone age models. To the best of our knowledge, no previously known MIS algorithm is self-stabilising, uses constant memory and constant randomness, and stabilizes in $\text{poly}(\log n)$ rounds in general or random graphs.

8.4 Scaling and Understanding AI systems

8.4.1 Randomized Smoothing under Attack: How Good is it in Practice?

Participants: Erwan Le Merrer.

Randomized smoothing is a recent and celebrated solution to certify the robustness of any classifier. While it indeed provides a theoretical robustness against adversarial attacks, the dimensionality of

current classifiers necessarily imposes Monte Carlo approaches for its application in practice. This work [28] questions the effectiveness of randomized smoothing as a defense, against state of the art black-box attacks. This is a novel perspective, as previous research works considered the certification as an unquestionable guarantee. We first formally highlight the mismatch between a theoretical certification and the practice of attacks on classifiers. We then perform attacks on randomized smoothing as a defense. Our main observation is that there is a major mismatch in the settings of the RS for obtaining high certified robustness or when defeating black box attacks while preserving the classifier accuracy.

Join work with Teddy Furon (Inria) and Thibault Maho (Inria).

9 Partnerships and cooperations

9.1 International initiatives

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

Audita

Title: Data auditing systems for recommendation decision-making algorithms

Duration: 2021 -> 2023

Coordinator: Anne-Marie Kermarrec (anne-marie.kermarrec@epfl.ch)

Partners:

- Ecole Polytechnique Fédérale de Lausanne Lausanne (Suisse)

Inria contact: Erwan Le Merrer

Summary: Although they still remain largely unnoticed, we are today surrounded by algorithms taking decisions on our behalf. These decisions range from apparently mundane choices, such as picking a VoD movie, or selecting on-line ads, to more life-changing decisions, such the granting of a credit by a bank, the triage of patients at a hospital, or the setting of a prison term for a convicted person. In their vast majority, decision-making algorithms exploit user data to predict the likely outcome of a decision. For instance, a credit will be granted to a customer based on the likelihood that this customer will default, based on her past credit history. In spite of the pervasiveness of such decision-making algorithms, users and institutions remain largely uninformed of their precise internal workings, and in particular tend to ignore how these algorithms operate on their data. This is a fundamental societal issue, as the decisions and their explanations are most of the time not provided, which lead citizens to feel confused and powerless. A decision-making algorithm essentially functions as a black-box, that consumes data collected from users (inputs), and produces decisions (outputs), while all intermediary steps remain hidden. Yet nowadays, these algorithms are executed at the service providers premises. Filter bubbles are a salient example of a problematic effect of a decision-making algorithm on users: those of recommender systems. Filter bubbles are a phenomenon where a recommendation algorithm locks the users into some narrow information bubbles with low entropy on information sources. Recommenders are then deciding which recommendations to display, while users have no understanding about the lack of diversity or the under/over-representation of particular groups of recommended items. Facing those concerns, a 2019 white paper entitled “Understanding algorithmic decision-making: Opportunities and challenges” from the European parliament, states that “Frameworks, composed of metrics, methodologies and tools that assess the impact of an Algorithmic Decision Systems and test its desired properties should be developed.” The proposed Audita associated team aims at tackling this challenge, by the proposal of a taxonomy of feasible audit tasks, and of specific audit algorithm for recommendation systems.

MLNS2

Title: Machine Learning, Network, System and Security

Duration: 2021 ->

Coordinator: Bernabé Batchakui (bbatchakui@gmail.com)

Partners:

- Université de Yaoundé Yaoundé (Cameroun)

Inria contact: David Bromberg

Summary: Nowadays there are no satisfactory solutions to stop the proliferation of: (i) simboxes, and (ii) malware over Android devices. They constitute a severe threat to any business. In one hand, simboxes enable massive interconnect bypass frauds, and hence provide low cost international calls while leveraging cellular networks from telecom operators without their authorization. In another hand, malware may interrupt and disable applications, retrieve and spoof personal information and identity, access sensitive information, control all applications executing on users' device, and even overcharge users for functionality that is widely available. The aim of this collaboration is to tackle the two aforementioned challenges from a system perspective. In particular we aim to adequately design and investigate efficient techniques to fight against simbox frauds and malware proliferation. Addressing such challenges require multidisciplinary knowledge such as Machine Learning, Network, System, and Security (MLNS2). Having these four areas of expertise in the same research team is rare, and this is one of the strengths of this collaboration. Our scientific goal is to bridge the gap between each of these four areas of expertise while leveraging our ongoing joint works.

9.2 International research visitors

9.2.1 Visits of international scientists

Inria International Chair

Other international visits to the team

9.2.2 Visits to international teams

Sabbatical programme

Research stays abroad

9.2.3 H2020 projects

SOTERIA [SOTERIA project on cordis.europa.eu](https://cordis.europa.eu/project/SOTERIA)

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, 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
- FONDATION DE L'INSTITUT DE RECHERCHE IDIAP (IDIAP), Switzerland
- ASOCIACION INSTITUTO DE INVESTIGACION SANITARIA BIOCRUCES BIZKAIA (BIOCRUCES BIZKAIA), Spain
- ARIADNEXT (ARIADNEXT), France
- ASOCIATIA INFOCONS (INFOCONS), Romania
- FUNDACION VASCA DE INNOVACION E INVESTIGACION SANITARIAS (BIOEF), Spain
- CENTRE DE VISIO PER COMPUTADOR (CVC), Spain
- KATHOLIEKE UNIVERSITEIT LEUVEN (KU Leuven), Belgium
- CENTRALESUPELEC (CentraleSupélec), France

Inria contact: Davide Frey

Coordinator:

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.

9.3 National initiatives

ANR Project PAMELA (2016-2022)

Participants: Davide Frey, George Giakkoupis, François Taïani.

PAMELA is a collaborative ANR project involving Inria/IRISA, Inria Lille (MAGNET team), UMPC, Mediego and Snips. The project aims at developing machine learning theories and algorithms in order to learn local and personalized models from data distributed over networked infrastructures. This project seeks to provide fundamental answers to modern information systems built by interconnecting many personal devices holding private user data in the search of personalized suggestions and recommendations. A significant asset of the project is the quality of its industrial partners, Snips and Mediego, who bring in their expertise in privacy protection and distributed computing as well as use cases and datasets.

Collaboration with the Conseil Supérieur de l'Audiovisuel (CSA)

Participants: Erwan Le Merrer.

Collaborating with the CSA (the French regulator of audio-visual content) on the YouTube recommender system, in order to assess the presence or not of the so called *rabbit-hole* phenomenon.

ANR Project ByBloS (2021-2025)

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

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 I) in Rennes, LIRIS (INSA Lyon) in Lyon, and LS2N (Université de Nantes) in Nantes. Budget: 252 220€

9.4 Regional initiatives

Cominlabs Project PriCLESS (2021-2023)

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

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. Budget:

10 Dissemination

10.1 Promoting scientific activities

Participants: Davide Frey, François Taïani, George Giakkoupis, Barbe Thys-tere Mvondo Djob, David Bromberg, Erwan Le Merrer, Mathieu Gestin, Jade Garcia Bourrée.

10.1.1 Scientific events: organisation

General chair, scientific chair

Member of the organizing committees

- David Bromberg served as General Chair of EuroSys 2022, the Seventeenth European Conference on Computer Systems, Rennes, France, April 5-8, 2022 (ACM).
- François Taïani served as Web-Chair of EuroSys 2022, the Seventeenth European Conference on Computer Systems, Rennes, France, April 5 - 8, 2022 (ACM).
- Djob Mvondo served as Poster Chair of EuroSys 2022, the Seventeenth European Conference on Computer Systems, Rennes, France, April 5-8, 2022 (ACM).
- François Taïani and Davide Frey co-organized the Workshop on Streaming (WOS'22), on November 24, in Rennes.

10.1.2 Scientific events: selection

Chair of conference program committees

Member of the conference program committees

- George Giakkoupis served on the PC of the 41st ACM Symposium on Principles of Distributed Computing (PODC), Salerno, Italy, Jul 25-29 2022.
- George Giakkoupis served on the PC of the 40th International Symposium on Theoretical Aspects of Computer Science (STACS), Hamburg, Germany, Mar 7-10 2023.
- François Taïani served on the PC of SRDS 2022, the 41th International Symposium on Reliable Distributed Systems (Vienna, Austria, September 19-22, 2022)
- François Taïani served on the PC of Middleware 2022 (23rd International Middleware Conference, Quebec, QC, Canada, November 7 - 11, 2022)
- François Taïani served on the PC of DSN 2022 (52nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2022, Baltimore, MD, USA, June 27-30, 2022)
- Djob Mvondo serves on the PC of EuroSys 2023, the Eighteenth European Conference on Computer Systems, Roma, Italy, May 8-12, 2023 (ACM).
- Djob Mvondo serves on the PC of MiddleWare 2023, the twenty-fourth ACM/IFIP Middleware conference, Bologna, Italy, December 11-15, 2023.
- Davide Frey served on the PC of DISC, the International Symposium on Distributed Computing, Augusta Georgia, USA, October 2022.
- Davide Frey serves on the PC of Middleware 2023, the twenty-fourth ACM/IFIP Middleware conference, Bologna, Italy, December 11-15, 2023.
- Davide Frey serves on the PC of DEBS, International ACM conference on Distributed and Event-Based Systems, Neuchatel, Switzerland, June 2023.
- Davide Frey serves on the PC of DAIS, the IFIP International Conference on Distributed Applications and Interoperable Systems, Lisbon, Portugal, June 2023.

Reviewer

10.1.3 Journal

Member of the editorial boards

- David Bromberg serves as an Associate Editor for ACM Transactions on Autonomous and Adaptive Systems (TAAS)

Reviewer - reviewing activities

- George Giakkoupis was a reviewer for Information Processing Letters (IPL).
- George Giakkoupis was a reviewer for Distributed Computing (DIST).
- Djob Mvondo was a reviewer for the PeerJ Computer Science journal.
- Davide Frey was a reviewer for Information Processing Letters (IPL)

10.1.4 Invited talks

- George Giakkoupis, Simple efficient distributed processes on graphs, Keynote talk, 29th International Colloquium on Structural Information and Communication Complexity (SIROCCO), Paderborn, Germany, Jun 28 2022.
- George Giakkoupis, Expanders via local edge flips in quasilinear time, 2nd Workshop Complexity and Algorithms (CoA), The Henri Poincare Institute (IHP), Paris, France, Sep. 27 2022.
- François Taïani, Asynchronous Byzantine Reliable Broadcast With a Message Adversary, Invited Talk, EPFL (Lausanne, Switzerland), 31 May 2022.
- Djob Mvondo, Towards enhanced performance and energy savings for new Cloud services, NEC Europe Laboratories, Heidelberg, Germany, 21 September 2022.
- David Bromberg, Major Android Safety Breach, Beware of your Android, Invited talk, University of Yaoundé 1, Cameroon, June 2022.
- David Bromberg, Major Android Safety Breach, Beware of your Android, Invited talk, University of Dshang, Cameroon, June 2022.
- David Bromberg, Major Android Safety Breach, Beware of your Android, Invited talk, University of Bordeaux, France, November 2022.
- David Bromberg, Odile : A scalable tracing tool for non-rooted and off-the-shelf Android phones, Invited talk, University of Yaoundé 1, Cameroon, November 2022.
- Davide Frey, Asynchronous Byzantine Reliable Broadcast With a Message Adversary, Invited Talk, University of Sydney, Australia, July 2022.
- Davide Frey, Donar: Anonymous VoIP over Tor, Invited Talk, KU Leuven, Belgium, October 2022.
- Mathieu Gestin, Hidden Issuer Anonymous Credential, Invited Talk, University of Sydney, Australia, July 2022.
- Jade Garcia Bourée, From image to graph watermarks, Invited Talk, EPFL, Switzerland, July 2022.
- Erwan Le Merrer, Auditing and modeling the rabbit-hole effect in YouTube, Invited Talk, EPFL, Switzerland, September 2022.

10.1.5 Leadership within the scientific community

- Erwan Le Merrer co-organized the "Journée infrastructures pour la souveraineté numérique", November 22nd, in CNAM Paris.
- Erwan Le Merrer co-organized the SUPSEC workshop on AI for supervision, September 19/20, in Inria Rennes.
- François Taïani was the chair of the hiring committee for the position of *Professeur des Universités at University of Rennes 1 - ESIR*.

- David Bromberg was the chair of the hiring committee for the position of *Professeur des Universités at University of Rennes 1 - ESIR*.
- David Bromberg was member of the hiring committee for the position of *Maître de conférences at University of Rennes 1 - ESIR*.
- François Taïani served on the promotion committee for computer science (*comité d'audition de repyramidage pour la section 27*) of Université Grenoble Alpes (UGA).
- Djob Mvondo co-organized a hackathon on Unikraft at Lyon 14-15 May 2022.
- David Bromberg organized a hackathon on AI for Malware detection at Yaoundé, Cameroon 7-9 November 2022.

10.1.6 Scientific expertise

- George Giakkoupis reviewed a grant proposal for the Israel Science Foundation (ISF).
- François Taïani reviewed an ERC starting grant proposal for the European Research Council.
- David Bromberg is an expert for the ministry of higher education and research concerning the French Research Tax Credit (CIR).
- Davide Frey was a project reviewer for the ANR and for the National Science Center of Poland

10.1.7 Research administration

- Erwan Le Merrer serves on the scientific board of Inria's REGALIA since 2021.
- Erwan Le Merrer serves on the administrative board of the Société Informatique de France since 2019.
- François Taïani served on the panel of the ARED programme (*Allocations de Recherche Doctorale*) of the region Brittany.
- François Taïani has served as Career Advice Person, (Réfèrent conseil-parcours professionnel chercheurs) for IRISA/Inria Rennes Bretagne Atlantique since 2019.
- David Bromberg leads the IRISA department (D1) entitled large scale distributed systems since 2021.
- David Bromberg serves on the administrative board of the Société Informatique de France since 2022.

10.2 Teaching - Supervision - Juries

Participants: Davide Frey, Francois Taiani, George Giakkoupis, Barbe Thystere Mvondo Djob, Erwan Le Merrer.

10.2.1 Teaching

- Bachelor: George Giakkoupis, Distributed Algorithms, 10h, L3 parcours SI, ISTIC, ENS Rennes, France.
- Master: Barbe Thystere Mvondo Djob, Network and Security for IOT, 45h, ESIR M1, Rennes, France
- Master: Barbe Thystere Mvondo Djob, FabLab for IOT, 30h, ESIR, M1, Rennes, France
- Master: Barbe Thystere Mvondo Djob, Cloud Computing for IOT, 45h, ESIR M2, Rennes, France

- Engineering School: David Bromberg, FabLab for IOT, 30h, ESIR, M1, Rennes, France
- Engineering School: David Bromberg, Web for IOT, 45h, ESIR, M1, Rennes, France
- Master: Erwan Le Merrer, Network Science, 10h, ESIR Rennes, France.
- Engineering School: Francois Taiani, Synchronization and Parallel Programming, 30h, 2nd year of Engineering School (M1), ESIR / U. Rennes I, France.
- Engineering School: Francois Taiani, Distributed Systems, 12h, 3rd year of Engineering School (M2), ESIR / U. Rennes I, France.
- Engineering School: François Taiani, Introduction to Operating Systems, 24h, 1st year of Engineering School (L3), ESIR / U. Rennes I, France.
- Master: Davide Frey, Scalable Distributed Systems, 10h, M1, EIT/ICT Labs Master School, U. Rennes I, France.
- ENS L3 : Davide Frey, Distributed Algorithms, 11h, ENS Rennes, France.
- Master: Davide Frey, Cloud Computing, 12h, M2-MIAGE, U. Rennes I, France.
- Master: Davide Frey, Distributed Systems/Systèmes Répartis, 21h, ENSAI, France.
- Master: Davide Frey, Apprentice Tutoring, 16h ETD, M2 Alternance U. Rennes I, France.
- Engineering School: François Taïani, Synchronization and Parallel Programming, 30h, 2nd year of Engineering School (M1), ESIR / U. Rennes I, France.
- Engineering School: François Taïani, Distributed Systems, 12h, 3rd year of Engineering School (M2), ESIR / U. Rennes I, France.
- Engineering School: François Taïani, Introduction to Operating Systems, 24h, 1st year of Engineering School (L3), ESIR / U. Rennes I, France.

10.2.2 Supervision

- PhD in progress: Thibault Maho, On the security of neural networks, started in 2020, supervised by Erwan Le Merrer and Teddy Furon (Linkmedia team).
- 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: Cesaire Honoré, Scheduling in heterogeneous architectures, started on December 2022, supervised by Yerom-David Bromberg and Djob Mvondo
- PhD in progress: Timothé Albouy, Towards Lightweight Scalable and Open Byzantine-Fault-Tolerant Distributed Objects, U. Rennes I, supervised by Francois Taiani and Davide Frey, started on Oct 18 2021.
- PhD in progress: Arthur Rauch, Frugal and Legal for Future Blockchain, Inria Rennes, supervised by Emmanuelle Anceaume and Davide Frey, started on Oct 1 2021.
- PhD in progress: Mathieu Gestin, Private Authenticated Storage for Online Services, Inria Rennes, supervised by Davide Frey, started on Oct 1 2021.
- Supervision of engineer Ali Yesilkanat, for developing audits of rabbit-holes on YouTube.
- Supervision of engineer Muhammed Selcuk Kok, for developing predictions of the rate of unemployment, in collaboration with Pôle Emploi.

10.2.3 Juries

- François Taïani was chairman for the PhD thesis of Pierre Jeanjean: IDE as Code, U. Rennes 1, April 29 2022.
- François Taïani was external reviewer for the PhD thesis of Arsany Guirguis: System Support for Robust Distributed Learning, EPFL (Switzerland), June 1 2022.
- David Bromberg was external reviewer for the PhD thesis of Brice Ekane Apah: Input-Output optimization in multi-tiers architectures, Grenoble (France), December 8 2022.
- David Bromberg was external reviewer for the PhD thesis of Moubarak ZOURE: Verification of anomalies in NFV-based network services outsourced to the cloud, March 18 2022.
- Davide Frey was external reviewer for the PhD thesis of Alejandro Ranchal Pedrosa, The Blockchain of Oz: Specifying Blockchain Failures for Scalable Protocols Offering Unprecedented Safety and Decentralization, University of Sydney.
- Davide Frey was external reviewer for the PhD thesis of Matthieu Nicolas, Ré-identification sans coordination dans les types de données répliquées sans conflits (CRDTs), December 20, 2022.

10.3 Popularization

Participants: Erwan Le Merrer.

10.3.1 Internal or external Inria responsibilities

- Erwan Le Merrer is at the scientific board of the REGALIA pilot project.

10.3.2 Articles and contents

- Erwan Le Merrer was interviewed for an article in The conversation: "Peut-on faire des sondages politiques avec YouTube ?", July 5th 2022.
- Erwan Le Merrer wrote an article in Binaire (blog Le Monde): "Le recommandeur de Youtube et les sondages électoraux", July 5th 2022.
- Erwan Le Merrer was interviewed for an article in the Inria's intranet: "Quand la campagne présidentielle inspire les scientifiques", June 30th 2022.

10.3.3 Interventions

- Erwan Le Merrer was interviewed for a video in Le Media TV, "L'élection approche, la guerre de l'information fait rage | Contre-matinale 117", March 29th 2022.

11 Scientific production

11.1 Major publications

- [1] 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](https://doi.org/10.1016/j.tcs.2021.06.021). URL: <https://hal.inria.fr/hal-03346710>.
- [2] P. Berenbrink, G. Giakkoupis and P. Kling. 'Tight Bounds for Coalescing-Branching Random Walks on Regular Graphs'. In: *SODA 2018 - Proceedings of the 29th ACM-SIAM Symposium on Discrete Algorithms*. New Orleans, United States: ACM, Jan. 2018, pp. 1715–1733. URL: <https://hal.inria.fr/hal-01635757>.

- [3] D. Bosk, D. Frey, M. Gestin and G. Piolle. ‘Hidden Issuer Anonymous Credential’. In: *Proceedings on Privacy Enhancing Technologies 2022* (June 2022), pp. 571–607. DOI: [10.56553/popets-2022-0123](https://doi.org/10.56553/popets-2022-0123). URL: <https://hal.archives-ouvertes.fr/hal-03789485>.
- [4] Y.-D. Bromberg, Q. Dufour, 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>.
- [5] K. Censor-Hillel, M. Ghaffari, G. Giakkoupis, B. Haeupler and F. Kuhn. ‘Tight Bounds on Vertex Connectivity Under Sampling’. In: *ACM Transactions on Algorithms* 13.2 (May 2017), 19:1–19:26. DOI: [10.1145/3086465](https://doi.org/10.1145/3086465). URL: <https://hal.inria.fr/hal-01635743>.
- [6] F. Chierichetti, G. Giakkoupis, S. Lattanzi and A. Panconesi. ‘Rumor Spreading and Conductance’. In: *Journal of the ACM (JACM)* 65.4 (Aug. 2018), 17:1–17:21. DOI: [10.1145/3173043](https://doi.org/10.1145/3173043). URL: <https://hal.inria.fr/hal-01942162>.
- [7] 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](https://doi.org/10.1145/3423211.3425685). URL: <https://hal.archives-ouvertes.fr/hal-03390450>.
- [8] G. Giakkoupis. ‘Expanders via local edge flips in quasilinear time’. In: STOC 2022 - 54th Annual ACM SIGACT Symposium on Theory of Computing. Rome, Italy: ACM, 25th May 2022, pp. 64–76. DOI: [10.1145/3519935.3520022](https://doi.org/10.1145/3519935.3520022). URL: <https://hal.inria.fr/hal-03792482>.
- [9] G. Giakkoupis, M. Jafari Giv and P. Woelfel. ‘Efficient Randomized DCAS’. In: STOC 2021 - 53rd Annual ACM SIGACT Symposium on Theory of Computing. Rome (Virtual), Italy: ACM, 21st June 2021, pp. 1–64. DOI: [10.1145/3406325.3451133](https://doi.org/10.1145/3406325.3451133). URL: <https://hal.inria.fr/hal-03195692>.
- [10] R. Guerraoui, A.-M. Kermarrec, O. Ruas and F. Taïani. ‘Smaller, Faster & Lighter KNN Graph Constructions’. In: WWW '20 - The Web Conference 2020. Taipei Taiwan, France: ACM, 20th Apr. 2020, pp. 1060–1070. DOI: [10.1145/3366423.3380184](https://doi.org/10.1145/3366423.3380184). URL: <https://hal.inria.fr/hal-02888286>.
- [11] H. Lakhlef, M. Raynal and F. Taïani. ‘Vertex Coloring with Communication Constraints in Synchronous Broadcast Networks’. In: *IEEE Transactions on Parallel and Distributed Systems* 30.7 (July 2019), pp. 1672–1686. DOI: [10.1109/TPDS.2018.2889688](https://doi.org/10.1109/TPDS.2018.2889688). URL: <https://hal.inria.fr/hal-02376726>.
- [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](https://doi.org/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](https://doi.org/10.1038/s42256-020-0216-z). URL: <https://hal.laas.fr/hal-03048809>.

11.2 Publications of the year

International journals

- [14] D. Bosk, D. Frey, M. Gestin and G. Piolle. ‘Hidden Issuer Anonymous Credential’. In: *Proceedings on Privacy Enhancing Technologies 2022* (June 2022), pp. 571–607. DOI: [10.56553/popets-2022-0123](https://doi.org/10.56553/popets-2022-0123). URL: <https://hal.archives-ouvertes.fr/hal-03789485>.
- [15] R.-A. Cherrueau, M. Delavergne, A. van Kempen, A. Lebre, D. Pertin, J. Rojas Balderrama, A. Simonet and M. Simonin. ‘EnosLib: A Library for Experiment-Driven Research in Distributed Computing’. In: *IEEE Transactions on Parallel and Distributed Systems* 33.6 (1st June 2022), pp. 1464–1477. DOI: [10.1109/TPDS.2021.3111159](https://doi.org/10.1109/TPDS.2021.3111159). URL: <https://hal.inria.fr/hal-03324177>.

- [16] 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: *ACM Transactions on Intelligent Systems and Technology* 13.5 (31st Oct. 2022), pp. 1–30. DOI: [10.1145/3527621](https://doi.org/10.1145/3527621). URL: <https://hal.inria.fr/hal-03906055>.
- [17] C. Delporte, H. Fauconnier, S. Rajsbaum and M. Raynal. ‘Distributed computability: Relating k-immediate snapshot and x-set agreement’. In: *Information and Computation* 285 (May 2022), p. 104815. DOI: [10.1016/j.ic.2021.104815](https://doi.org/10.1016/j.ic.2021.104815). URL: <https://hal.inria.fr/hal-03920684>.
- [18] A. Durand, M. Raynal and G. Taubenfeld. ‘Contention-related crash failures: Definitions, agreement algorithms, and impossibility results’. In: *Theoretical Computer Science* 909 (Mar. 2022), pp. 76–86. DOI: [10.1016/j.tcs.2022.01.029](https://doi.org/10.1016/j.tcs.2022.01.029). URL: <https://hal.uca.fr/hal-03746704>.
- [19] D. Frey, A. Mostefaoui, M. Perrin, P.-L. Roman and F. Taïani. ‘Differentiated consistency for world-wide gossips’. In: *IEEE Transactions on Parallel and Distributed Systems* 34.1 (23rd Sept. 2022), pp. 1–15. DOI: [10.1109/TPDS.2022.3209150](https://doi.org/10.1109/TPDS.2022.3209150). URL: <https://hal.inria.fr/hal-03797554>.
- [20] M. Raynal and G. Taubenfeld. ‘A visit to mutual exclusion in seven dates’. In: *Theoretical Computer Science* 919 (June 2022), pp. 47–65. DOI: [10.1016/j.tcs.2022.03.030](https://doi.org/10.1016/j.tcs.2022.03.030). URL: <https://hal.inria.fr/hal-03920720>.

International peer-reviewed conferences

- [21] T. Albouy, D. Frey, M. Raynal and F. Taïani. ‘A Modular Approach to Construct Signature-Free BRB Algorithms under a Message Adversary’. In: OPODIS 2022 - 26th Conference on Principles of Distributed Systems. Brussels, Belgium, 13th Dec. 2022, pp. 1–44. URL: <https://hal.inria.fr/hal-03906141>.
- [22] 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.2022.32](https://doi.org/10.4230/LIPIcs.DISC.2022.32). URL: <https://hal.inria.fr/hal-03791921>.
- [23] 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, pp. 1–17. URL: <https://hal.inria.fr/hal-03923695>.
- [24] C. Delporte-Gallet, H. Fauconnier, M. Raynal and M. Safir. ‘Optimal Algorithms for Synchronous Byzantine k-Set Agreement’. In: SSS 2022 - 24th International Symposium on Stabilizing, Safety, and Security of Distributed Systems. Vol. LNCS-13751. Lecture Notes in Computer Science. Clermont-Ferrand, France: Springer International Publishing, 9th Nov. 2022, pp. 178–192. DOI: [10.1007/978-3-031-21017-4_12](https://doi.org/10.1007/978-3-031-21017-4_12). URL: <https://hal.inria.fr/hal-03920733>.
- [25] A. Durand, M. Raynal and G. Taubenfeld. ‘Reaching Consensus in the Presence of Contention-Related Crash Failures’. In: *Lecture Notes in Computer Science*. SSS 2022 - 24th International Symposium on Stabilization, Safety, and Security of Distributed Systems. Vol. LNCS-13751. Stabilization, Safety, and Security of Distributed Systems. Clermont-Ferrand, France: Springer International Publishing, 9th Nov. 2022, pp. 193–205. DOI: [10.1007/978-3-031-21017-4_13](https://doi.org/10.1007/978-3-031-21017-4_13). URL: <https://hal.archives-ouvertes.fr/hal-03853639>.
- [26] G. Giakkoupis. ‘Expanders via local edge flips in quasilinear time’. In: STOC 2022 - 54th Annual ACM SIGACT Symposium on Theory of Computing. Rome, Italy: ACM, 25th May 2022, pp. 64–76. DOI: [10.1145/3519935.3520022](https://doi.org/10.1145/3519935.3520022). URL: <https://hal.inria.fr/hal-03792482>.
- [27] T. Maho, T. Furon and E. Le Merrer. ‘FBI: Fingerprinting models with Benign Inputs’. In: CAID 2022 - Conference on Artificial Intelligence for Defense. Actes de la 4ème Conference on Artificial Intelligence for Defense (CAID 2022). Rennes, France, 16th Nov. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03879849>.
- [28] T. Maho, T. Furon and E. Le Merrer. ‘Randomized Smoothing under Attack: How Good is it in Practice?’ In: ICASSP 2022 - IEEE International Conference on Acoustics, Speech and Signal Processing. Singapore, Singapore: IEEE, 22nd May 2022, pp. 1–5. URL: <https://hal.inria.fr/hal-03591421>.

- [29] M. Pigaglio, J. Bruneau-Queyreix, Y.-D. Bromberg, D. Frey, E. Rivière and L. Réveillère. ‘RAPTEE: Leveraging trusted execution environments for Byzantine-tolerant peer sampling services’. In: ICDCS 2022 - 42nd IEEE International Conference on Distributed Computing Systems. Bologna, Italy: IEEE, 10th July 2022, pp. 1–11. URL: <https://hal.inria.fr/hal-03923712>.
- [30] A. Rauch, Q. Bramas, S. Devismes, P. Lafourcade and A. Lamani. ‘Exploration perpétuelle : ça s’en va et ça revient’. In: AlgoTel 2022 - 24èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications. Saint-Rémy-Lès-Chevreuse, France, 30th May 2022, pp. 1–4. DOI: [10.5281/zenodo.4640462](https://doi.org/10.5281/zenodo.4640462). URL: <https://hal.archives-ouvertes.fr/hal-03657044>.

Conferences without proceedings

- [31] D. Mvondo, A. Barbalace, J.-P. Lozi and G. Muller. ‘Towards User-Programmable Schedulers in the Operating System Kernel’. In: SPMA 22 - 11th workshop on Systems for Post-Moore Architectures. Rennes, France, 5th Apr. 2022, pp. 1–4. URL: <https://hal.inria.fr/hal-03750209>.

Scientific book chapters

- [32] R. Duvignau, M. Raynal and E. Schiller. ‘Self-stabilizing Byzantine Fault-Tolerant Repeated Reliable Broadcast’. In: *Stabilization, Safety, and Security of Distributed Systems*. Vol. 13751. Lecture Notes in Computer Science. Springer International Publishing, 9th Nov. 2022, pp. 206–221. DOI: [10.1007/978-3-031-21017-4_14](https://doi.org/10.1007/978-3-031-21017-4_14). URL: <https://hal.inria.fr/hal-03920748>.
- [33] D. Imbs, M. Raynal and G. Taubenfeld. ‘Election in Fully Anonymous Shared Memory Systems: Tight Space Bounds and Algorithms’. In: *Structural Information and Communication Complexity*. Vol. 13298. Lecture Notes in Computer Science. Springer International Publishing, 25th June 2022, pp. 174–190. DOI: [10.1007/978-3-031-09993-9_10](https://doi.org/10.1007/978-3-031-09993-9_10). URL: <https://hal.inria.fr/hal-03920724>.

Reports & preprints

- [34] T. Albouy, D. Frey, M. Raynal and F. Taïani. *Asynchronous Byzantine Reliable Broadcast With a Message Adversary*. 18th May 2022. URL: <https://hal.inria.fr/hal-03671451>.
- [35] D. Frey, M. Raynal, F. Taïani and T. Albouy. *A Modular Approach to Construct Signature-Free BRB Algorithms under a Message Adversary*. 8th Nov. 2022. URL: <https://hal.archives-ouvertes.fr/hal-03653878>.
- [36] M. S. Kok. *Is Okun's law still valid in France?* 25th May 2022. URL: <https://hal.inria.fr/hal-03678747>.
- [37] E. Le Merrer, R. Pons and G. Trédan. *Algorithmic audits of algorithms, and the law*. 22nd Feb. 2022. URL: <https://hal.inria.fr/hal-03583919>.
- [38] E. Le Merrer and G. Trédan. *Qu'est ce qu'un algorithme en boîte noire ? Tractatus des décisions algorithmiques*. 2022. URL: <https://hal.inria.fr/hal-03851597>.
- [39] E. Le Merrer and G. Trédan. *Surfing Personalization for Quantifying the Rabbit Hole Phenomenon on YouTube*. 25th Mar. 2022. URL: <https://hal.science/hal-03620039>.
- [40] E. Le Merrer, G. Trédan and A. Yesilkanat. *YouTube Recommendations Do Predict Polls: A note on the 2022 French presidential election*. Rapport LAAS n° 22136. Inria, 29th Apr. 2022. URL: <https://hal.inria.fr/hal-03655608>.

Other scientific publications

- [41] A. Clementi, G. Giakkoupis, E. Natale and F. d'Amore. ‘Search via Parallel Lévy Walks on Z^2 ’. In: HALG 2022 - Highlights of Algorithms. London, United Kingdom, 1st June 2022, pp. 1–13. URL: <https://hal.archives-ouvertes.fr/hal-03694177>.

11.3 Other

Scientific popularization

- [42] E. Le Merrer and G. Trédan. *What is a black box algorithm?: Tractatus of algorithmic decision-making*. 2023. URL: <https://hal.inria.fr/hal-03940259>.

11.4 Cited publications

- [43] Y. Afek and E. Gafni. ‘Asynchrony from synchrony’. In: *ICDCN*. 2013, pp. 225–239.
- [44] A. Ahmed and E. Ahmed. ‘A survey on mobile edge computing’. In: *2016 10th International Conference on Intelligent Systems and Control (ISCO)*. Jan. 2016, pp. 1–8. DOI: [10.1109/ISCO.2016.7727082](https://doi.org/10.1109/ISCO.2016.7727082). URL: <http://dx.doi.org/10.1109/ISCO.2016.7727082>.
- [45] 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](https://doi.org/10.1145/3008631.3008635). URL: <https://hal.inria.fr/hal-01421986>.
- [46] 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.
- [47] 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.
- [48] D. Bernstein. ‘Containers and Cloud: From LXC to Docker to Kubernetes’. In: *IEEE Cloud Computing* 1.3 (Sept. 2014), pp. 81–84. DOI: [10.1109/MCC.2014.51](https://doi.org/10.1109/MCC.2014.51). URL: <http://dx.doi.org/10.1109/MCC.2014.51>.
- [49] M. Bertier, D. Frey, R. Guerraoui, A.-M. Kermarrec and V. Leroy. ‘The Gossple Anonymous Social Network’. In: *ACM/IFIP/USENIX 11th International Middleware Conference (MIDDLEWARE)*. Ed. by I. Gupta and C. Mascolo. Vol. LNCS-6452. Middleware 2010. Bangalore, India: Springer, Nov. 2010, pp. 191–211. DOI: [10.1007/978-3-642-16955-7_10](https://doi.org/10.1007/978-3-642-16955-7_10). URL: <https://hal.inria.fr/inria-00515693>.
- [50] F. Bonomi. *Connected vehicles, the internet of things, and fog computing*. *VANET 2011, 2011*. Keynote speech at VANET. 2011.
- [51] F. Bonomi, R. Milito, J. Zhu and S. Addepalli. ‘Fog Computing and Its Role in the Internet of Things’. In: *1st MCC Workshop on Mobile Cloud Computing*. 2012. DOI: [10.1145/2342509.2342513](https://doi.org/10.1145/2342509.2342513). URL: <http://doi.acm.org/10.1145/2342509.2342513>.
- [52] 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>.
- [53] A. Boutet, D. Frey, R. Guerraoui, A.-M. Kermarrec and R. Patra. ‘HyRec: Leveraging Browsers for Scalable Recommenders’. In: *Middleware 2014*. Bordeaux, France, Dec. 2014. DOI: [10.1145/2663165.2663315](https://doi.org/10.1145/2663165.2663315). URL: <https://hal.inria.fr/hal-01080016>.
- [54] A. Boutet, D. Frey, R. Guerraoui, A.-M. Kermarrec, A. Rault, F. Taïani and J. Wang. ‘Hide & Share: Landmark-based Similarity for Private KNN Computation’. In: *DSN*. Rio de Janeiro, Brazil, 2015. DOI: [10.1109/DSN.2015.60](https://doi.org/10.1109/DSN.2015.60). URL: <https://hal.archives-ouvertes.fr/hal-01171492>.
- [55] 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>.
- [56] B. Cohen. *Incentives Build Robustness in BitTorrent*. 2003. URL: <http://citeseer.ist.psu.edu/cohen03incentives.html>.

- [57] D. Collins, R. Guerraoui, J. Komatovic, P. Kuznetsov, M. Monti, M. Pavlovic, Y.-A. Pignolet, D.-A. Seredinschi, A. Tonkikh and A. Xygkis. ‘Online Payments by Merely Broadcasting Messages’. In: *IEEE DSN*. 2020. DOI: [10.1109/DSN48063.2020.00023](https://doi.org/10.1109/DSN48063.2020.00023). URL: <https://doi.org/10.1109/DSN48063.2020.00023>.
- [58] A. Dash, A. Mukherjee and S. Ghosh. ‘A Network-centric Framework for Auditing Recommendation Systems’. In: *IEEE Conference on Computer Communications, INFOCOM*. 2019.
- [59] 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.
- [60] 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.
- [61] 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>.
- [62] 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.
- [63] R. Guerraoui, K. Huguenin, A.-M. Kermarrec, M. Monod and S. Prusty. ‘LiFTinG: Lightweight Freerider-Tracking Protocol in Gossip’. In: *11th ACM/IFIP/USENIX International Middleware Conference (MIDDLEWARE)*. Bangalore, India, Nov. 2010. DOI: [10.1007/978-3-642-16955-7_16](https://doi.org/10.1007/978-3-642-16955-7_16). URL: <https://hal.inria.fr/inria-00505268>.
- [64] R. Guerraoui, P. Kuznetsov, M. Monti, M. Pavlovic and D.-A. Seredinschi. ‘The Consensus Number of a Cryptocurrency’. In: *ACM PODC*. 2019. DOI: [10.1145/3293611.3331589](https://doi.org/10.1145/3293611.3331589). URL: <https://doi.org/10.1145/3293611.3331589>.
- [65] 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.
- [66] K. Huang, H. Wei, Y. Huang, H. Li and A. Pan. ‘Byz-GentleRain: An Efficient Byzantine-tolerant Causal Consistency Protocol’. In: *CoRR* abs/2109.14189 (2021). arXiv: [2109.14189](https://arxiv.org/abs/2109.14189). URL: <https://arxiv.org/abs/2109.14189>.
- [67] 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.
- [68] 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.
- [69] D. Kempe, J. M. Kleinberg and É. Tardos. ‘Influential Nodes in a Diffusion Model for Social Networks’. In: *ICALP*. 2005, pp. 1127–1138.
- [70] 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.
- [71] D. Kempe, J. M. Kleinberg and É. Tardos. ‘Maximizing the spread of influence through a social network’. In: *KDD*. 2003, pp. 137–146.
- [72] P. Kuznetsov et al. ‘Understanding non-uniform failure models’. In: *Bulletin of the EATCS* 106 (2012), pp. 53–77.
- [73] E. Lieberman, C. Hauert and M. Nowak. ‘Evolutionary dynamics on graphs’. In: *Nature* 433.7023 (2005), pp. 312–316.
- [74] D. Malkhi and M. Reiter. ‘Byzantine quorum systems’. In: *Distributed computing* 11.4 (1998), pp. 203–213.
- [75] D. Mvondo, A. Tchana, R. Lachaize, D. Hagimont and N. D. Palma. ‘Fine-Grained Fault Tolerance for Resilient pVM-Based Virtual Machine Monitors’. In: *50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN*. IEEE, 2020, pp. 197–208. DOI: [10.1109/DSN48063.2020.00037](https://doi.org/10.1109/DSN48063.2020.00037). URL: <https://doi.org/10.1109/DSN48063.2020.00037>.

- [76] D. Mvondo, B. Teabe, A. Tchana, D. Hagimont and N. D. Palma. ‘Memory flipping: a threat to NUMA virtual machines in the Cloud’. In: *2019 IEEE Conference on Computer Communications, INFOCOM 2019*. IEEE, 2019, pp. 325–333. DOI: [10.1109/INFOCOM.2019.8737548](https://doi.org/10.1109/INFOCOM.2019.8737548). URL: <https://doi.org/10.1109/INFOCOM.2019.8737548>.
- [77] F. Pasquale. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard U. Press, 2015.
- [78] M. Raynal and J. Stainer. ‘Synchrony weakened by message adversaries vs asynchrony restricted by failure detectors’. In: *PODC. Proceedings of the 2013 ACM symposium on Principles of distributed computing*. Montréal, Canada: ACM, July 2013, pp. 166–175. DOI: [10.1145/2484239.2484249](https://doi.org/10.1145/2484239.2484249). URL: <https://hal.inria.fr/hal-00920734>.
- [79] N. Santoro and P. Widmayer. ‘Time is not a healer’. In: *Annual Symposium on Theoretical Aspects of Computer Science*. Springer, 1989, pp. 304–313.
- [80] A. Tchana, L. L. Wapet and Y.-D. Bromberg. ‘Odile: A scalable tracing tool for non-rooted and on-device Android phones’. In: *25th International Symposium on Research in Attacks, Intrusions and Defenses, RAID 2022, Limassol, Cyprus, October 26-28, 2022*. ACM, 2022, pp. 252–262. DOI: [10.1145/3545948.3545951](https://doi.org/10.1145/3545948.3545951). URL: <https://doi.org/10.1145/3545948.3545951>.
- [81] 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.
- [82] L. Zhang, F. Zhou, A. Mislove and R. Sundaram. ‘Maygh: Building a CDN from Client Web Browsers’. In: *8th ACM European Conference on Computer Systems. EuroSys ’13*. Prague, Czech Republic: ACM, 2013, pp. 281–294. DOI: [10.1145/2465351.2465379](https://doi.org/10.1145/2465351.2465379). URL: <http://doi.acm.org/10.1145/2465351.2465379>.