UMR IRISA

Activity Report 2023

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

Université de Rennes | rennes | INSA INSTITUT NATIONAL DES SCIENCES APPLIQUÉES RENNES | cnrs | Inria | Inserm | CentraleSupélec | IMT Atlantique Bretagne-Pays de la Loire École Mines-Télécom | Université Bretagne Sud ubs

Contents

Project-Team WIDE

Creation of the Project-Team: 2018 June 01

Keywords

Computer sciences and digital sciences

A1.2.5. – Internet of things

A1.2.9. – Social Networks

A1.3.2. – Mobile distributed systems

A1.3.3. – Blockchain

A1.3.4. – Peer to peer

A1.3.5. – Cloud

A1.3.6. – Fog, Edge

A2.1.7. – Distributed programming

A2.6.1. – Operating systems

A2.6.2. – Middleware

A2.6.3. – Virtual machines

A3.5.1. – Analysis of large graphs

A4. – Security and privacy

A4.8. – Privacy-enhancing technologies

A7.1.1. – Distributed algorithms

A7.1.2. – Parallel algorithms

A7.1.3. – Graph algorithms

A9. – Artificial intelligence

A9.2. – Machine learning

A9.9. – Distributed AI, Multi-agent

Other research topics and application domains

B6.1.1. – 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]
- George Giakkoupis [INRIA, Researcher]
- Erwan Le Merrer [INRIA, Senior Researcher, HDR]

Faculty Members

- François Taiani [Team leader, UNIV RENNES, Professor, 1/2 Delegation until 31/8/23, 1/4 Delegation from 1/9/23, HDR]
- Yerom David Bromberg [UNIV RENNES, Professor Delegation, until Aug 2023, HDR]
- Barbe Mvondo Djob [UNIV RENNES, Associate Professor]
- Michel Raynal [UNIV RENNES, Emeritus, HDR]

Post-Doctoral Fellow

- Luis Santiago Luevano Garcia [INRIA, Post-Doctoral Fellow, from May 2023]

PhD Students

- Timothé Albouy [UNIV RENNES]
- Alexandre Duvivier [UNIV RENNES, CIFRE, from Sep 2023]
- Jade Garcia Bourree [INRIA]
- Adrien Gegout [UNIV RENNES, CIFRE, from Jun 2023]
- Mathieu Gestin [INRIA]
- Augustin Godinot [UNIV RENNES]
- Amélie Gonzalez [UNIV RENNES, from Oct 2023]
- Dimitri Lereverend [INRIA, from Oct 2023]
- Honore Cesaire Mounah [INRIA]
- Ludovic Paillat [HIVE, from Oct 2023, Université de Lorraine]
- Rémy Raes [INRIA]
- Arthur Rauch [INRIA]
- Manon Sourisseau [UNIV RENNES, from Oct 2023]

Technical Staff

- Cyrille Kenfack [INRIA, Engineer, from Sep 2023]

Interns and Apprentices

- Hugo Bertin [UNIV RENNES, Intern, from Mar 2023 until Aug 2023]
- Amélie Gonzalez [UNIV RENNES, Intern, from Feb 2023 until Aug 2023]
- Dimitri Lereverend [ENS RENNES, Intern, until Jun 2023]
- Victoire Nganfang [INRIA, Intern, from Aug 2023]
- Manon Sourisseau [UNIV RENNES, Intern, until Jul 2023]

Administrative Assistant

- Virginie Desroches [INRIA]

Visiting Scientists

- Bernabe Batchakui [Ecole Nationale Supérieur polytechnique Yaoundé, until Jan 2023]
- Anne-Marie Chana [Ecole Nationale Supérieur polytechnique Yaoundé, until Jan 2023]

2 Overall objectives

2.1 Overview

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

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

2.2 Planetary-Scale Geo-Distributed Systems

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

individual developers can only hope to grasp parts of the resulting systems. Similarly, weak consistency models and loose coordination mechanisms tend to lead to counter-intuitive behaviors, while only providing weak overall guarantees. This lack of systematic guarantees and understandability make it harder for practitioners to design, deploy, and validate the distributed systems they produce, leading to rising costs and high entry barriers.

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

2.3 Highly Personalized On-Line Services

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

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

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

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

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

2.4 Social Collaboration Platforms

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

The challenges faced by such social applications blend in many ways issues already discussed in the two previous subsections and cast them in an application-driven context. These social collaboration

platforms require mechanisms that go beyond pure message propagation, with stricter consistency and robustness guarantees. Because they involve connected users, these applications must provide usable solutions, in particular in terms of latency and availability. At the same time, because they manipulate real-world transactions and objects (money, cars, accommodations) they must also provide a high level of consistency and guarantees. Many of these applications further operate at a planetary scale, and therefore also face stark scalability issues, that make them highly interesting case studies to investigate innovative architectures combining decentralized and centralized elements.

Formalizing and characterizing the needs and behaviors of these new applications seems particularly interesting in order to provide the fertile ground for new systems and novel theoretical work. The area of social applications also offers avenues for knowledge transfer and societal impact, along two dimensions. First, practical and usable approaches, back by a deep understanding of the foundation of distribution and coordination, are likely to find applications in future systems. Second, developers of complex social applications are often faced with a lack of robust scalable services[1] that can be easily exploited to harness the latest understanding of large-scale distributed coordination. We therefore think these applications offer an opportunity to design and deliver modular reusable bricks that can be easily appropriated by a large population of innovative developers without requiring the level of deep understanding usually necessary to implement these solutions from scratch. Providing such reusable bricks is however difficult, as many interesting formal properties are not composable, and a unified composable theory of distributed systems still need to be fully articulated.

3 Research program

3.1 Overview

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

More specifically, our program revolves around four key challenges.

- Challenge 1: Designing Hybrid Scalable Architectures,
- Challenge 2: Constructing Personalizable Privacy-Aware Distributed Systems,
- 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).

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

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 [46, 58] systems and hybrid infrastructures (i.e. architectures that combines centralized or hierarchical elements, often hosted in well-provisioned data-centers, and a decentralized part, often hosted in a peer-to-peer overlay) [50]. 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 [47] 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 [41].

A number of efforts in this directions focus on specific hardware, e.g. fog nodes that are responsible for connected IoT devices [48]. 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 [78], 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 [77], Kubernetes [45]) 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 [49, 51, 52] 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 [53, 60]. On the other hand, we will consider blockchains as a tool for tracking and rewarding data contributions. Ultimately, we aim at enabling users to configure the level of privacy and personalization they wish to experience.

Privacy preserving decentralized aggregation As a second setting we would like to consider target applications running on constrained devices like in the Internet-of-Things (IoT). This setting makes it particularly important to operate on decentralized data in a light-weight privacy-preserving manner, and further highlights the synergy between this objective and Objective 1. For example, we plan to provide data subjects with the possibility to store and manage their data locally on their own devices, without having to rely on third-party managers or aggregators, but possibly storing less private information or results in the cloud. Using this strategy, we intend to design protocols that enable users themselves, or third-party companies to query distributed data in aggregate form, or to run data analytics processes on a distributed set of data repositories, thereby gathering knowledge without violating the privacy of other users. For example, we have started working on the problem of computing an aggregate function over a subset of the data in a distributed setting. This involves two major steps: selection and aggregation. With respect to selection, we envision defining a decentralized data-selection operation that can apply a selection predicate without violating privacy constraints. With respect to aggregation, we will continue our investigation of lightweight protocols that can provide privacy with limited computational complexity [42].

3.4 Network Diffusion Processes

Social, biological, and technological networks can serve as conduits for the spread of ideas, trends, diseases, or viruses. In social networks, rumors, trends and behaviors, or the adoption of new products, spread from person to person. In biological networks, diseases spread through contact between individuals, and mutations spread from an individual to its offsprings. In technological networks, such as the Internet and the power grid, viruses and worms spread from computer to computer, and power failures often lead to cascading failures. The common theme in all the examples above is that the rumor, disease,

or failure starts out with a single or a few individual nodes, and propagates through the network, from node to node, to reach a potentially much larger number of nodes.

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

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

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

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

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

Some immunization problems are known to be hard to approximate in general graphs, even for the perfect-contagion model, e.g., the fixed-budget version of the fire-fighter problem cannot be approximated to any $n^{1-\epsilon}$ factor [44]. This strand of work will consider restricted graph families, such as trees or graphs of small treewidth, for such problems. In addition, for some immunization problems, there is a large gap between the best known approximation algorithm and the best known inaproximability result, and we would like to make progress in reducing these gaps.

3.5 Systemizing Modular Distributed Computability and Efficiency

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

We plan to investigate these foundations along two directions: first by studying novel approaches to some fundamental problems of *mutual exclusion and distributed coordination*, and second by exploring

how we can build a *comprehensive and modular framework* capturing the foundations of *distributed computation*.

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

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

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

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

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

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

To progress on this vision, we plan in the near future to investigate, from a distributed software point of view, the impact due to failures and asynchrony on the layered architecture of distributed computing systems. A first step in this direction will address the notions of *message adversaries* (introduced a long time ago in [76]) and *process adversaries* (investigated in several papers, e.g. [75, 56, 64, 65, 69]). The aim of these notions is to consider failures, not as "bad events", but as part of the normal behavior of a system. As an example, when considering round-based algorithms, a message adversary is a daemon which, at every round, is allowed to suppress some messages. The aim is then, given a problem P, to find the strongest adversary under which P can be solved ("strongest" means here that giving more power to the adversary makes the problem impossible to solve). This work will allow us to progress in terms of general *layered* theory of distributed computing, and allow us to better *map* distributed computing models and their relations, in the steps of noticeable early efforts in this direction [75, 40].

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 understanding of the lower levels of large-scale distributed infrastructures (for Djob). These novel challenges and opportunities lead us to propose the following four updated objectives.

Objective 1: Large-scale Trustless Sybil-Resistant Systems

We plan to contribute to the theoretical understanding of Blockchain-based and Byzantine-tolerant systems by exploring reusable abstractions that can allow programmers to develop Byzantine-tolerant applications more easily. We plan for example to extend existing work on weak consistency to a BFT setting, building for instance on recent proposals on Byzantine Fault-Tolerant CRDTs [63]. To address scale, we plan to explore novel scalable Byzantine fault-tolerant algorithms, both in the context of closed systems, and then in the more challenging case of open (aka permissionless) systems. Our line of attack is to focus on lightweight BFT primitives that can enable faster and more resource-efficient algorithms [54, 61]. In the case of open systems, we will leverage the expertise of our team in theoretical distributed algorithms and randomized algorithms to address Sybil attacks through novel countermeasures providing (hopefully) cheaper and more equitable alternatives to proof-of-work of proof-of-stake algorithms. One open, yet enticing, questions is whether anonymous computing models could provide a path to address this issue. We would also like to investigate how storage can be improved in Blockchains and BFT large-scale systems. Most of these systems are fully replicated, incurring formidable costs (up to 2.6PB of distributed storage in the case of Bitcoin). Coding techniques, that we have used in the past, and adaptable redundancy based on Byzantine quorums [71] are some avenues we would like to explore to address this challenge.

Objective 2: Robustness and Security at Scale

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

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

On the front of privacy, we plan to investigate the design of a distributed digital data vault able to securely store personal data, leveraging our experience on privacy-preserving decentralized systems [42], and on trusted-execution environments (e.g. SGX). We have started collaborating with the CIDRE team at Inria Rennes, with colleagues at KTH (Sweden), and with the company AriadNext (H2020 Soteria project) on these topics.

At an infrastructure level, and following the recruitment of Djob Mvondo, we plan to explore how progress in virtualization can help advance the team's agenda in terms of large-scale robustness, in particular in a cloud-computing setting [72, 73]. Specifically we would like to investigate how novel heterogeneous architectures that embed a range of ASICs and specialized units (GPU, FPGA, SMARTNIC, PIM-devices) can be leveraged to provide more robust and more efficient virtualized services.

Objective 3: 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) [55]. These algorithms, often coined *black-box algorithms* [74], 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 [43], or a k-nearest neighbor graph), by means of local decentralized operations (typically adjacent nodes interact by exchanging some of their incident edges). In all three cases, we will strive to derive time- and space- optimal algorithms, with strong robustness guarantees.

4 Application domains

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

5 Social and environmental responsibility

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

6 Highlights of the year

6.1 Awards

- Paper [26] co-authored by George Giakkoupis received the PODC 2023 Best Paper Award.

6.2 Other

- WIDE is currently involved in three Inria Challenges (Défis) (FedMalin, Alvearium, and OS Research in France).
- Two ANR proposals submitted by WIDE permanent members have been funded in 2023: The JCJC proposal "sGOV: Tailored governors for VMs to achieve real energy savings while ensuring performance" by Djob Mvondo, and the ANR proposal "Second Chance" by David Bromberg.
- WIDE together with its partner Broadpeak organized the Inria/IRISA Broadpeak Workshop on Streaming (WOS'23) on Tuesday, November 28, 2023, at Inria Rennes – IRISA, in Rennes. The event attracted close to 90 participants (on site and on-line).

7 New software, platforms, open data

7.1 New software

7.1.1 DecentralizedFlower

Name: DecentralizedFlower

Keyword: Decentralized Learning

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

Contact: Davide Frey

8 New results

8.1 Distributed Algorithms and Systems

8.1.1 Word-size RMR tradeoffs for recoverable mutual exclusion

Participants: George Giakkoupis.

In [26] we present tradeoffs between RMR complexity and memory word size for recoverable mutual exclusion (RME) algorithms using arbitrary synchronization primitives. Assuming that each memory location stores w bits, we show that n-process mutual exclusion has an RMR complexity of at least $\Omega\left(\min\left\{\log_w n, \log n/\log\log n\right\}\right)$ on the DSM and the CC model. For $w = (\log n)^{\Omega(1)}$, our lower bound asymptotically matches an upper bound by Katzan and Morrison (2020), whose RME mutual exclusion algorithm employs w-bit fetch-and-add operations. Our lower bound is the first one that does not restrict the type of atomic operations that can be executed on a memory location.

This work was done in collaboration with David Yu Cheng Chan (U. Calgary) and Philipp Woelfel (U. Calgary), and was the recipient of the PODC 2023 Best Paper Award.

8.1.2 Good-case Early-Stopping Latency of Synchronous Byzantine Reliable Broadcast: The Deterministic Case (Extended Version).

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

This work [36] considers the good-case latency of Byzantine Reliable Broadcast (BRB), i.e., the time taken by correct processes to deliver a message when the initial sender is correct. This time plays a crucial role in the performance of practical distributed systems. Although significant strides have been made in recent years on this question, progress has mainly focused on either asynchronous or randomized algorithms. By contrast, the good-case latency of deterministic synchronous BRB under a majority of Byzantine faults has been little studied. In particular, it was not known whether a goodcase latency below the worst-case bound of t + 1 rounds could be obtained. This work answers this open question positively and proposes a deterministic synchronous Byzantine reliable broadcast that achieves a good-case latency of max(2, t + 3 - c) rounds, where t is the upper bound on the number of Byzantine processes and c the number of effectively correct processes.

8.1.3 Optimal Algorithms for Synchronous Byzantine k-Set Agreement (Journal Version)

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 [15] 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).

8.1.4 Self-stabilizing Byzantine Fault-Tolerant Repeated Reliable Broadcast (Journal version)

Participants: Michel Raynal.

In this work [17], 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.

8.1.5 Basalt: A Rock-Solid Byzantine-Tolerant Peer Sampling for Very Large Decentralized Networks

Participants: Yérom-David Bromberg, Davide Frey, Djob Mvondo, François Taïani.

Recent large-scale Byzantine-Fault-Tolerant (BFT) algorithms provide scalability at a low cost by exploiting a secure Random Peer Sampling (RPS) service: a service that provides a stream of random network nodes where no attacking entity can become over-represented. Unfortunately, producing good peer samples untainted by Byzantine behavior in a large-scale network is particularly difficult, with existing solutions unable to withstand aggressive attacks. In this work [25], we propose a novel RPS algorithm, BASALT, that implements what we have termed a stubborn chaotic search over node IDs to counter attackers' attempts at becoming over-represented. Our evaluation based on a theoretical analysis, Monte Carlo simulations, and experiments on a live cryptocurrency network shows that BASALT delivers close-to-optimal protection against malicious behaviors and outperforms state-of-the-art solutions by a wide margin.

This is a joint work with Alex Auvolat from the Deuxfleurs association.

8.1.6 Asynchronous Byzantine reliable broadcast with a message adversary

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

This work [14] considers the problem of reliable broadcast in asynchronous authenticated systems, in which n processes communicate using signed messages and up to t processes may behave arbitrarily (Byzantine processes). In addition, for each message m broadcast by a correct (i.e., non-Byzantine) process, a message adversary may prevent up to d correct processes from receiving m. (This message adversary captures network failures such as transient disconnections, silent churn, or message losses.) Considering such a "double" adversarial context and assuming $n > 3t + 2d$, a reliable broadcast algorithm is presented. Interestingly, when there is no message adversary (i.e., $d = 0$), the algorithm terminates in two communication steps (so, in this case, this algorithm is optimal in terms of both Byzantine tolerance and time efficiency). It is then shown that the condition $n > 3t + 2d$ is necessary for implementing reliable broadcast in the presence of both Byzantine processes and a message adversary (whether the underlying system is enriched with signatures or not).

8.1.7 The Synchronization Power (Consensus Number) of Access-Control Objects: the Case of AllowList and DenyList.

Participants: Davide Frey, Mathieu Gestin, Michel Raynal.

We studied the synchronization power of AllowList and DenyList objects under the lens provided by Herlihy's consensus hierarchy. It specifies AllowList and DenyList as distributed objects and shows that, while they can both be seen as specializations of a more general object type, they inherently have different synchronization power. While the AllowList object does not require synchronization between

participating processes, a DenyList object requires processes to reach consensus on a specific set of processes. These results are then applied to a more global analysis of anonymity-preserving systems that use AllowList and DenyList objects. First, a blind-signature-based e-voting is presented. Second, DenyList and AllowList objects are used to determine the consensus number of a specific decentralized key management system. Third, an anonymous money transfer algorithm using the association of AllowList and DenyList objects is presented. Finally, this analysis is used to study the properties of these application, and to highlight efficiency gains that they can achieve in message passing environment. This paper appeared at DISC 2023 [28].

8.1.8 Differentiated Consistency for Worldwide Gossips

Participants: Davide Frey, François Taïani.

Eventual consistency is a consistency model that favors liveness over safety. It is often used in large-scale distributed systems where models ensuring a stronger safety incur performance that are too low to be deemed practical. Eventual consistency tends to be uniformly applied within a system, but we argue a demand exists for differentiated eventual consistency, e.g. in blockchain systems. In this work [18], we propose update-query consistency with primaries and secondaries (UPS) to address this demand. UPS is a novel consistency mechanism that works in pair with our novel two-phase epidemic broadcast protocol gossip primary-secondary (GPS) to offer differentiated eventual consistency and delivery speed. We propose two complementary analyses of the broadcast protocol: a continuous analysis and a discrete analysis based on compartmental models used in epidemiology. Additionally, we propose the formal definition of a scalable consistency metric to measure the consistency trade-off at runtime. We evaluate UPS in two simulated worldwide settings: a one-million-node network and a network emulating that of the Ethereum blockchain. In both settings, UPS reduces inconsistencies experienced by a majority of the nodes and reduces the average message latency for the remaining nodes.

This is a joint work with Achour Mostéfaoui (U. Nantes), Matthieu Perrin (U. Nantes), and Pierre-Louis Roman (EPFL, Switzerland).

8.1.9 GoldFinger: Fast and Approximate Jaccard for Efficient KNN Graph Constructions.

Participants: François Taïani.

In this work [19], we propose *GoldFinger*, a new *compact* and *fast-to-compute* binary representation of datasets to approximate Jaccard's index. We illustrate the effectiveness of GoldFinger on the emblematic big data problem of K-Nearest-Neighbor (KNN) graph construction and show that GoldFinger can drastically accelerate a large range of existing KNN algorithms with little to no overhead. As a side effect, we also show that the compact representation of the data protects users' privacy *for free* by providing k-anonymity and l-diversity. Our extensive evaluation of the resulting approach on several realistic datasets shows that our approach reduces computation times by up to 78.9% compared to raw data while only incurring a negligible to moderate loss in terms of KNN quality. We also show that GoldFinger can be applied to KNN queries (a widely-used search technique) and delivers speedups of up to ×3.55 over one of the most efficient approaches to this problem.

This is a joint work with Rachid Guerraoui and Anne-Marie Kermarrec from EPFL, Guilhem Niot from ENS Lyon, and Olivier Ruas from Inria Lille (Spirals Team).

8.1.10 Design of an Efficient Distributed Delivery Service for Group Key Agreement Protocols

Participants: Davide Frey, Ludovic Paillat.

End-to-end encrypted messaging applications such as Signal became widely popular thanks to their capability to ensure the confidentiality and integrity of online communication. While the highest security guarantees were long reserved to two-party communication, solutions for n-party communication remained either inefficient or less secure until the standardization of the MLS Protocol (Messaging Layer Security). This new protocol offers an efficient way to provide end-to-end secure communication with the same guarantees originally offered by the Signal Protocol for two-party communication. However, both solutions still rely on a centralized component for message delivery, called the Delivery Service in the MLS Protocol. The centralization of the Delivery Service makes it an ideal target for attackers and threatens the availability of any protocol relying on MLS. In order to overcome this issue, we proposed the design of a fully distributed Delivery Service that allows clients to exchange protocol messages efficiently and without any intermediary. It uses a Probabilistic Reliable-Broadcast mechanism to efficiently deliver messages and the Cascade Consensus Protocol to handle messages requiring an agreement. Our solution strengthens the availability of the MLS Protocol without compromising its security. This work appeared at FPS 2023 [34].

This is joint work with Claudia Lavigna Ignat from the COAST team, and Mathieu Turiani from the PESTO team.

8.2 Network and Graph Algorithms

8.2.1 Distributed self-stabilizing MIS with few states and weak communication

Participants: George Giakkoupis.

In [29] 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, changes its state with probability 1/2. The process stabilizes with probability 1 on any graph, and the resulting set of black vertices is an MIS. We show that the expected stabilization time is $O(\log n)$ on certain graph families, such as cliques and graphs of bounded arboricity.

Our main result is that the process stabilizes in $poly(\log n)$ rounds w.h.p. on $G_{n,p}$ random graphs, for $0 \le p \le poly(\log n) \cdot n^{-1/2}$ or $p \ge 1/poly(\log n)$. Further, we propose an extension of this process, with larger but still constant vertex state space, which stabilizes in $poly(\log n)$ rounds on $G_{n,p}$ w.h.p., for all $1 \le p \le 1$. Both processes readily translate into distributed/parallel MIS algorithms, which are self-stabilizing, use constant space (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-stabilizing, uses constant space and constant randomness, and stabilizes in $poly(\log n)$ rounds on $G_{n,p}$ random graphs.

This work was done in collaboration with Isabella Ziccardi (Bocconi University, Italy).

8.3 Scaling and Understanding AI systems

8.3.1 FBI: Fingerprinting models with Benign Inputs

Participants: Erwan Le Merrer.

Recent advances in the fingerprinting of deep neural networks are able to detect specific instances of models, placed in a black-box interaction scheme. Inputs used by the fingerprinting protocols are specifically crafted for each precise model to be checked for. While efficient in such a scenario, this nevertheless results in a lack of guarantee after a mere modification of a model (e.g. finetuning, quantization of the parameters). These works generalize [21, 21] fingerprinting to the notion of model families and their variants and extends the task-encompassing scenarios where one wants to fingerprint not only a

precise model (previously referred to as a detection task) but also to identify which model or family is in the black-box (identification task). The main contribution is the proposal of fingerprinting schemes that are resilient to significant modifications of the models. We achieve these goals by demonstrating that benign inputs, that are unmodified images, are sufficient material for both tasks. We leverage an information-theoretic scheme for the identification task. We devise a greedy discrimination algorithm for the detection task. Both approaches are experimentally validated over an unprecedented set of more than 1,000 networks 1 .

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

8.3.2 Modeling Rabbit-Holes on YouTube

Participants: Erwan Le Merrer.

Numerous discussions have advocated the presence of a so called rabbit-hole (RH) phenomenon on social media, interested in advanced personalization to their users. This phenomenon is loosely understood as a collapse of mainstream recommendations, in favor of ultra personalized ones that lock users into narrow and specialized feeds. Yet quantitative studies are often ignoring personalization, are of limited scale, and rely on manual tagging to track this collapse. This precludes a precise understanding of the phenomenon based on reproducible observations, and thus the continuous audits of platforms. In this work [20], we first tackle the scale issue by proposing a user-sided bot-centric approach that enables large scale data collection, through autoplay walks on recommendations. We then propose a simple theory that explains the appearance of these RHs. While this theory is a simplifying viewpoint on a complex and planet- wide phenomenon, it carries multiple advantages: it can be analytically modeled, and provides a general yet rigorous definition of RHs. We define them as an interplay between i) user interaction with personalization and ii) the attraction strength of certain video categories, which cause users to quickly step apart of mainstream recommendations made to fresh user profiles. We illustrate these concepts by highlighting some RHs found after collecting more than 16 million personalized recommendations on YouTube. A final validation step compares our automatically-identified RHs against manually-identified RHs from a previous research work. Together, those results pave the way for large scale and automated audits of the RH effect in recommendation systems.

Join work with Gilles Tredan (LAAS/CNRS) and ALI Yesilkanat (Inria).

8.3.3 Algorithmic audits of algorithms, and the law

Participants: Erwan Le Merrer.

Algorithmic decision making is now widespread, ranging from health care allocation to more common actions such as recommendation or information ranking. The aim to audit these algorithms has grown alongside. In this paper, we focus on external audits that are conducted by interacting with the user side of the target algorithm, hence considered as a black box. Yet, the legal framework in which these audits take place is mostly ambiguous to researchers developing them: on the one hand, the legal value of the audit outcome is uncertain; on the other hand the auditors' rights and obligations are unclear. The contribution of this work [24] is to articulate two canonical audit forms to law, to shed light on these aspects: (i) The first audit form (we coin the Bobby audit form) checks a predicate against the algorithm, while the second (Sherlock) is more loose and opens up to multiple investigations. We find that: Bobby audits are more amenable to prosecution, yet are delicate as operating on real user data. This can lead to reject by a court (notion of admissibility). Sherlock audits craft data for their operation, most notably to build surrogates of the audited algorithm. It is mostly used for acts for whistleblowing, as even if accepted as a proof, the evidential value will be low in practice. (ii) These two forms require the prior respect of a proper right to audit, granted by law or by the platform being audited; otherwise the auditor will be also prone to prosecutions regardless of the audit outcome. This article thus highlights the relation of current audits with law, in order to structure the growing field of algorithm auditing.

Join work with Gilles Tredan (LAAS/CNRS) and Ronan Pons (Université d'Ottawa).

8.3.4 Exploring the Effectiveness of Lightweight Architectures for Face Anti-Spoofing

Participants: Luis S. Luevano.

Detecting spoof faces is crucial in ensuring the robustness of face-based identity recognition and access control systems, as faces can be captured easily without the user's cooperation in uncontrolled environments. Several deep models have been proposed for this task, achieving high levels of accuracy but at a high computational cost. Considering the very good results obtained by lightweight deep networks on different computer vision tasks, in this work [33] we explore the effectiveness of this kind of architectures for face anti-spoofing. Specifically, we asses the performance of three lightweight face models on two challenging benchmark databases. The conducted experiments indicate that face anti-spoofing solutions based on lightweight face models are able to achieve comparable accuracy results to those obtained by state-of-the-art very deep models, with a significantly lower computational complexity.

This is joint work with Yoanna Martínez-Díaz and Heydi Méndez-Vázquez from the Biometrics group at CENATAV and Miguel González-Mendoza from Tecnologico de Monterrey.

8.3.5 Effectiveness of Blind Face Restoration to Boost Face Recognition Performance at Low-Resolution Images

Participants: Luis S. Luevano.

This paper [32] studies the effectiveness of Blind Face Restoration methods to boost the performance of face recognition systems on low-resolution images. We investigate the use of three blind face restoration techniques, which have demonstrated impressive results in generating realistic high-resolution face images. Three state-of-the-art face recognition methods were selected to assess the impact of using the generated high-resolution images on their performance. Our analysis includes both, synthesized and native low-resolution images. The conducted experimental evaluation show that this is still an open research problem.

This is joint work with Yoanna Martínez-Díaz and Heydi Méndez-Vázquez from the Biometrics group at CENATAV.

8.3.6 Performance and explainability of feature selection-boosted tree-based classifiers for COVID-19 detection

Participants: Davide Frey.

This work [23] evaluates the performance and analyzes the explainability of machine learning models boosted by feature selection in predicting COVID-19-positive cases from self-reported information. In essence, this work describes a methodology to identify COVID-19 infections that considers the large amount of information collected by the University of Maryland Global COVID-19 Trends and Impact Survey (UMD-CTIS). More precisely, this methodology performs a feature selection stage based on the recursive feature elimination (RFE) method to reduce the number of input variables without compromising detection accuracy. A tree-based supervised machine learning model is then optimized with the selected features to detect COVID-19active cases. In contrast to previous approaches that use a limited set of selected symptoms, the proposed approach builds the detection engine considering a broad range of features including self-reported symptoms, local community information, vaccination acceptance, and isolation measures, among others. We considered three different supervised classifiers were used: random forests (RF), light gradient boosting (LGB), and extreme gradient boosting (XGB). Based on data

collected from the UMD-CTIS, we evaluated the detection performance of the methodology for four countries (Brazil, Canada, Japan, and South Africa) and two periods (2020 and 2021). The proposed approach was assessed in terms of various quality metrics: F1score, sensitivity, specificity, precision, receiver operating characteristic (ROC), and area under the ROC curve (AUC). This work also shows the normalized daily incidence curves obtained by the proposed approach for the four countries.

Joint work with Jesús Rufino, Juan Marcos Ramírez, Jose Aguilar, Jaya Champati, Antonio Fernández-Anta from Imdea Networks, Madrid Spain, Carlos Baquero from University of Minho, Portugal, and Rosa Elvira Lillo from University Carlos III, Madrid, Spain.

8.3.7 Consistent Comparison of Symptom-based Methods for COVID-19 Infection Detection

Participants: Davide Frey.

In this work [22] we carried out a comprehensive comparison of various COVID-19 detection methods based on self-reported information using the University of Maryland Global COVID-19 Trends and Impact Survey (UMD-CTIS), a large health surveillance platform, which was launched in partnership with Facebook.

We implemented and evaluated fifteen classifiers from three different categories: rule-based approaches, logistic regression techniques, and tree-based machine-learning models. These methods were evaluated using different metrics including F1-score, sensitivity, specificity, and precision. An explainability analysis has also been conducted to compare methods.

Our explainability analysis reveals that the relevance of the reported symptoms in COVID-19 detection varies between countries and years. However, there are two variables consistently relevant across approaches: stuffy or runny nose, and aches or muscle pain.

Regarding the categories of detection methods, evaluating detection methods using homogeneous data across countries and years provides a solid and consistent comparison. An explainability analysis of a tree-based machine-learning model can assist in identifying infected individuals specifically based on their relevant symptoms. This study is nonetheless limited by the self-report nature of data, which cannot replace clinical diagnosis.

Joint work with Jesús Rufino, Juan Marcos Ramírez, Jose Aguilar, Jaya Champati, Antonio Fernández-Anta from Imdea Networks, Madrid Spain, Carlos Baquero from University of Minho, Portugal, and Rosa Elvira Lillo from University Carlos III, Madrid, Spain.

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

9.1.1 CIFRE with Broadpeak

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

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

9.1.2 CIFRE with Blacknut: Efficient Containerized Cloud-Gaming Platforms

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

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

10 Partnerships and cooperations

10.1 International initiatives

10.1.1 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 recommandation 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 [6]. This is a fundamental societal issue, as the decisions and their explanations are most of 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 [3]. 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." [7]. 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 -> 2023

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 businesses. 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.

10.1.2 Participation in other International Programs

Decentralized Learning with Byzantine Agents, Inria-EPFL laboratory

Participants: François Taiani.

François Taïani co-supervised Geovani Rizk's PostDoc at EPFL on Decentralized Learning with Byzantine Agents, in collaboration with Rachid Gerraoui within the Inria-EPFL laboratory.

10.2 European initiatives

10.2.1 H2020 projects

SOTERIA SOTERIA project on cordis.europa.eu

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

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

Partners:

- INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET AUTOMATIQUE (INRIA), France

- IPCENTER AT GMBH (IPCENTER), Austria
- NORIA ONLUS, Italy
- AUDENCIA, France
- STELAR SECURITY TECHNOLOGY LAW RESEARCH UG (HAFTUNGSBESCHRANKT) GMBH (STELAR), Germany
- Servicio Vasco de Salud Osakidetza (Osakidetza), Spain
- SCYTL ELECTION TECHNOLOGIES SL, Spain
- ERDYN ATLANTIQUE, France
- EDUPRO GROUP GMBH, Austria
- FONDATION DE L'INSTITUT DE RECHERCHE IDIAP (IDIAP), Switzerland
- ASOCIACION INSTITUTO DE INVESTIGACION SANITARIA BIOBIZKAIA (BIOBIZKAIA), Spain
- IDnow SAS (IDnow), France
- ASOCIATIA INFOCONS (INFOCONS), Romania
- FUNDACION VASCA DE INNOVACION E INVESTIGACION SANITARIAS (BIOEF), Spain
- ERDYN CONSULTANTS SARL, France
- CENTRE DE VISIO PER COMPUTADOR (CVC-CERCA), Spain
- KATHOLIEKE UNIVERSITEIT LEUVEN (KU Leuven), Belgium
- CENTRALESUPELEC, France

Inria contact: Davide Frey

Coordinator: Montaser Awal, IDnow

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

10.3 National initiatives

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

Participants: Erwan Le Merrer.

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

ANR JCJC Project sGOV (2023-2027)

Participants: Djob Mvondo, Yerom-David Bromberg.

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

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

Partners: IRISA (coordinator, U. Rennes 1). Budget: 286 814.5€

ANR Project ByBloS (2021-2025)

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

Blockchain-based systems have over the last 10 years profoundly impacted society and research. They come however with many inefficiencies, that are inherent to the problem they attempt to solve, Byzantine Tolerant Agreement, one of the most difficult problems of distributed computing. Many Blockchain-based applications do not require the strong guarantees that an agreement provides. Building on this insight, Byblos seeks to explore the design, analysis, and implementation of lightweight Byzantine decentralized mechanisms for the systematic construction of large-scale Byzantine-tolerant Privacy-Preserving distributed systems.

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

Inria Challenge Project FedMalin

Participants: François Taiani, Davide Frey, Cyrille Kenfack, Remy Raes.

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

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

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

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

Inria Challenge Project Alvearium

Participants: François Taiani, Davide Frey.

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

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

10.4 Regional initiatives

Cominlabs Project PriCLeSS (2021-2024)

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

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

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

11 Dissemination

11.1 Promoting scientific activities

Participants: Erwan Le Merrer, George Giakkoupis, François Taiani, Djob Mvondo.

11.1.1 Scientific events: organisation

Member of the organizing committees

- Erwan Le Merrer co-organized the first workshop on algorithmic audits of algorithms (WAAA), May 23rd 2023, on Zoom.
- François Taïani and Davide Frey co-organized the 2023 Workshop on Streaming (WOS'23), Tuesday Nov. 28 2023, at Inria Rennes and online, in collaboration with the company Broadpeak.
- François Taïani served as Publicity Co-Chair for the 53rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN-2023), Porto, Portugal, June 27-30, 2023.

11.1.2 Scientific events: selection

Member of the conference program committees

- Erwan Le Merrer was in the PC of the SIAM International Conference on Data Mining (SDM24), Houston, TX, U.S.A., April 2023.
- George Giakkoupis served on the PC of the 42nd ACM Symposium on Principles of Distributed Computing (PODC), Orlando, Florida, US, Jun 19-23 2023.
- 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 the 11th International Conference on Networked Systems (NETYS), Ben Guerir, Morocco, May 22-24, 2023.
- Djob Mvondo served on the PC of the 24th ACM/IFIP International Middleware Conference (MIDDLEWARE), Bologna, Italy, December 11-15, 2023.
- Davide Frey served on the PC of the 24th ACM/IFIP International Middleware Conference (MIDDLEWARE), Bologna, Italy, December 11-15, 2023.
- Davide Frey served on the PC of the 17th ACM International Conference on Distributed and Event Based Systems (DEBS), Neuchatel, Switzerland, May 27-30, 2023.
- Davide Frey served on the PC of theInternational Conference on Distributed Applications and Interoperable Systems (DAIS), Lisbon, Portugal, June 19-23, 2023

11.1.3 Journal

Reviewer - reviewing activities

- George Giakkoupis reviewed papers for journal Distributed Computed (DIST).

11.1.4 Invited talks

- George Giakkoupis. Distributed self-stabilizing MIS with dew states and weak communication. 3rd Workshop Complexity and Algorithms (CoA 2023), LIP6, Sorbonne Université. Paris, France, Sep. 18 2023
- George Giakkoupis. Expanders via local edge flips in quasilinear time. U. Cambridge Randomized Algorithms Group Seminar (Online), Cambridge, UK, Feb. 24 2023
- François Taïani. Good-case Early-Stopping Latency of Synchronous Byzantine Reliable Broadcast: The Deterministic Case. Invited Talk, LIRIS, INSA Lyon, 22 November 2023

11.1.5 Leadership within the scientific community

- George Giakkoupis is a member of the steering committee (member-at-large) of the ACM Symposium on Principles of Distributed Computing (PODC) from 2023-2026.

11.1.6 Scientific expertise

- George Giakkoupis is a member of the Working Group GT CoA: Complexité et Algorithmes from 2023-2026.
- George Giakkoupis reviewed a grant proposal for the Israel Science Foundation (ISF).
- George Giakkoupis reviewed a grant proposal for the Vienna Science and Technology Fund (WWTF).
- François Taïani served on the evaluation Panel of the CHIST Era call "Security and Privacy in Decentralised and Distributed Systems".

11.1.7 Research administration

- Erwan Le Merrer is the head of the scientific council of the Société Informatique de France (SIF) since 2023.
- George Giakkoupis is a local correspondent of the Inria Centre Univ Rennes for the preparation of the Annual Activity Reports by the project teams.
- François Taïani is a Career Advisor for IRISA/Inria (Référent conseil-parcours professionel chercheurs) (since March 2019)

11.2 Teaching - Supervision - Juries

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

11.2.1 Teaching

- ENS L3: George Giakkoupis, Distributed Algorithms, 9h, 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
- 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.
- 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.

11.2.2 Supervision

- PhD in progress: Dimitri Lerévérend, Privacy-Preserving Decentralized Learning Through Model Fragmentation and Private Aggregation, started in 2023, supervised by Davide Frey, Romaric Gaudel (LACODAM team) and François Taïani.
- PhD in progress: Manon Sourisseau, Byzantine-Tolerant Netcodes For Tomorrow's Metaverse, started in 2023, supervised by François Taïani, David Bromber, and Jérémie Découchant (TU Delft).

- PhD in progress: Rémy Raes, Distributed Machine Learning in Ubiquitous Environments using Location-dependent Models, started in 2023, supervised by Davide Frey, François Taiani (WIDE team, Inria Rennes), Romain Rouvoy, Lionel Seinturier (Spirals team, Inria Lille).
- PhD in progress: Ludovic Paillat, Security for peer-to-peer cloud storage without central authority, started in 2023, supervised by Davide Frey, (WIDE team, Inria Rennes), Claudia Ignat (COAST team, Inria Nancy), Amine Ismail (HIVE), Mathieu Turiani (PESTO Team, Inria Nancy).
- PhD in progress: Jade Garcia Bourrée, Trust but verify: bot-driven audits of AI systems, started in October 2022, supervised by Erwan Le Merrer and Gilles Trédan (LAAS/CNRS).
- PhD in progress: Augustin Godinot, Auditing the mutations of AI-models, started on November 2022, supervised by Erwan Le Merrer, Gilles Trédan (LAAS/CNRS) François Taïani and Camilla Penzo (PEReN).
- PhD in progress: 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 François Taïani 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.
- PhD in progress: Amelie Gonzalez, Linux network stack optimization, Started on September 2023, supervised by Yerom-David Bromberg, Djob Mvondo, Julia Lawal (Inria Paris)
- PhD in progress: Alexandre Duvivier, CDN performance optimization, Started on October 2023, supervised by Yerom-David Bromberg, Djob Mvondo, Nicolas Le Scouarnec (Broadpeak)
- PhD in progress: Adrien Gegout, Efficient containerized Cloud Gaming, Started on October 2023, supervised by Davide Frey, Djob Mvondo, Pascal Manchon (Blacknut).
- PostDoc: Geovani Rizk (EPFL), Decentralized Learning with Byzantine Agents, in collaboration with Rachid Gerraoui within the Inria-EPFL laboratory.
- PostDoc: Luis Santiago Luevano Garcia, Private and Secure Computation on Personal Data, within the SOTERIA H2020 project.

11.2.3 Juries

- Erwan Le Merrer was a jury member for Thibault Maho's PhD thesis: Security of deep neural-networks under realistic scenarios, Inria, 14 December 2024.
- George Giakkoupis was an examiner for Robin Vacus' PhD thesis: Algorithmic perspective to flocking and foraging, U. Paris Cité, 18 Dec 2023.
- George Giakkoupis was an examiner for Louis de Monterno' PhD thesis: Synchronization in dynamic networks, École polytechnique, 19 Oct 2023.
- François Taïani was a reviewer for Kadir Korkmaz's PhD thesis: Securing Blockchains Against IoT Cyberattacks, Université de Bordeaux, 24 March 2023
- François Taïani was a reviewer and committee chairman for Laurent Prosperi's PhD thesis: Varda: a language for programming distributed systems by composition, Sorbonne Université, 5 September 2023

- François Taïani was a reviewer for Fatima Elhattab's PhD thesis: Robust and Privacy-Preserving Federated Learning, INSA Lyon, 22 November 2023
- François Taïani was committee chairman for Anne Bumiller's PhD thesis: Beyond Risk Scores: Context-Aware Adaptive Authentication, Univ Rennes, 9 November 2023
- François Taïani was committee chairman for Clément Courageux-Sudan's PhD thesis: End-to-end simulation of the energy consumption of fog infra- structures and their applications, ENS-Rennes, 8 Decembre 2023
- François Taïani was a reviewer for Omar Hasan's HDR thesis: Privacy Preservation in Trust-Deficient Decentralized Systems, INSA Lyon, 16 June 2023
- François Taïani was a reviewer for Gil Utard's HDR thesis: Calcul sur les données volumineuses et stockage distribué à grande échelle, U Picarie, 12 July 2023
- Davide Frey was a reviewer for Abdulaye Diallo's PhD thesis: Ecriture de contrats intelligents – essai de methodologie en droit et en informatique, University of Grenoble, December 2023.

11.3 Popularization

Participants: Erwan Le Merrer, Rémy Raes.

11.3.1 Articles and contents

- Erwan Le Merrer was interviewed by Science & vie magazine, under the title "Face aux IA - La course au nouveau test de Turing", published in September 2023.
- Erwan Le Merrer published an article in the Blog Binaire (Le Monde), under the title "ChatGPT et test de Turing inversé", published in May 2023.

11.3.2 Education

- Rémy Raes participated to three events in relation with dissemination activites towards high school pupils in general and girls in particular.
  - 21 June, academic forum to tell high school teachers about what can be done with their students and computer science labs of the university.
  - From 26 to 27 October, "Rendez-vous des jeunes mathématiciennes et informaticiennes" (RJMI), an event to welcome high school girls in the lab during two days to make them work on scientific projects (filles-et-maths.fr/rjmi/).
  - 17 November, he participated to an event to welcome secondary school pupils to Inria Lille Interface showcase room to present his research career and computer science in general.

12 Scientific production

12.1 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. URL: https://hal.inria.fr/hal-03346710.

[2] D. Bosk, D. Frey, M. Gestin and G. Piolle. 'Hidden Issuer Anonymous Credential'. In: *Proceedings on Privacy Enhancing Technologies* 2022 (June 2022), pp. 571–607. DOI: 10.56553/popets-2022-0123. URL: https://hal.archives-ouvertes.fr/hal-03789485.

[3] Y.-D. Bromberg, Q. Dufour and D. Frey. 'Multisource Rumor Spreading with Network Coding'. In: *INFOCOM 2019 - IEEE International Conference on Computer Communications*. Paris, France: IEEE, Apr. 2019, pp. 1–10. URL: https://hal.inria.fr/hal-01946632.

[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] G. Damaskinos, R. Guerraoui, A.-M. Kermarrec, V. Nitu, R. Patra and F. Taïani. 'FLeet: Online Federated Learning via Staleness Awareness and Performance Prediction'. In: *Middleware '20: Proceedings of the 21st International Middleware Conference*. 21st International Middleware Conference. Delft (virtual), Netherlands, 7th Dec. 2020. DOI: 10.1145/3423211.3425685. URL: https://hal.archives-ouvertes.fr/hal-03390450.

[6] D. Frey, M. Gestin and M. Raynal. 'The Synchronization Power (Consensus Number) of Access-Control Objects: the Case of AllowList and DenyList'. In: DISC 2023 - 37th International Symposium on Distributed Computing. L'aquila, Italy: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, pp. 1–32. DOI: 10.4230/LIPIcs.DISC.2023.21. URL: https://inria.hal.science/hal-04399298.

[7] G. Giakkoupis. 'Expanders via local edge flips in quasilinear time'. In: STOC 2022 - 54th Annual ACM SIGACT Symposium on Theory of Computing. Rome, Italy: ACM, 25th May 2022, pp. 64–76. DOI: 10.1145/3519935.3520022. URL: https://hal.inria.fr/hal-03792482.

[8] G. Giakkoupis, M. Jafari Giv and P. Woelfel. 'Efficient Randomized DCAS'. In: STOC 2021 - 53rd Annual ACM SIGACT Symposium on Theory of Computing. Rome (Virtual), Italy: ACM, 21st June 2021, pp. 1–64. DOI: 10.1145/3406325.3451133. URL: https://hal.inria.fr/hal-03195692.

[9] R. Guerraoui, A.-M. Kermarrec, O. Ruas and F. Taïani. 'Smaller, Faster & Lighter KNN Graph Constructions'. In: WWW '20 - The Web Conference 2020. Taipei Taiwan, France: ACM, 20th Apr. 2020, pp. 1060–1070. DOI: 10.1145/3366423.3380184. URL: https://hal.inria.fr/hal-02888286.

[10] 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. URL: https://hal.inria.fr/hal-02376726.

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

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

[13] T. Maho, T. Furon and E. L. Merrer. 'SurFree: a fast surrogate-free black-box attack'. In: CVPR 2021 - Conference on Computer Vision and Pattern Recognition. Proc. of IEEE Conference on Computer Vision and Pattern Recognition, CVPR. Virtual, France, 19th June 2021, pp. 10430–10439. URL: https://hal.archives-ouvertes.fr/hal-03177639.

12.2 Publications of the year

International journals

[14] T. Albouy, D. Frey, F. Taïani and M. Raynal. 'Asynchronous Byzantine reliable broadcast with a message adversary'. In: *Theoretical Computer Science* 978 (Nov. 2023), p. 114110. DOI: 10.1016/j.tcs.2023.114110. URL: https://inria.hal.science/hal-04212154.

[15] C. Delporte-Gallet, H. Fauconnier, M. Raynal and M. Safir. 'Optimal algorithms for synchronous Byzantine k-set agreement'. In: *Theoretical Computer Science* 973 (Sept. 2023), p. 114098. DOI: 10.1016/J.TCS.2023.114098. URL: https://inria.hal.science/hal-04395355.

[16] A. Durand, M. Raynal and G. Taubenfeld. 'Reaching agreement in the presence of contention-related crash failures'. In: *Theoretical Computer Science* 966-967 (July 2023), p. 113982. DOI: 10.1016/j.tcs.2023.113982. URL: https://uca.hal.science/hal-04323433.

[17] R. Duvignau, M. Raynal and E. M. Schiller. 'Self-stabilizing Byzantine fault-tolerant repeated reliable broadcast'. In: *Theoretical Computer Science* 972 (Sept. 2023), p. 114070. DOI: 10.1016/J.TCS.2023.114070. URL: https://inria.hal.science/hal-04395645.

[18] D. Frey, A. Mostefaoui, M. Perrin, P.-L. Roman and F. Taiani. 'Differentiated consistency for world-wide gossips'. In: *IEEE Transactions on Parallel and Distributed Systems* 35.11 (1st Jan. 2023), pp. 11461–11475. DOI: 10.1109/TPDS.2022.3209150. URL: https://inria.hal.science/hal-03797554.

[19] R. Guerraoui, A.-M. Kermarrec, G. Niot, O. Ruas and F. Taïani. 'GoldFinger: Fast & Approximate Jaccard for Efficient KNN Graph Constructions'. In: *IEEE Transactions on Knowledge and Data Engineering* 35.11 (1st Nov. 2023), pp. 11461–11475. DOI: 10.1109/TKDE.2022.3232689. URL: https://inria.hal.science/hal-04394851.

[20] E. Le Merrer, G. Trédan and A. Yesilkanat. 'Modeling Rabbit-Holes on YouTube'. In: *Social Network Analysis and Mining* 13.1 (Dec. 2023), p. 100. DOI: 10.1007/s13278-023-01105-9. URL: https://hal.science/hal-03620039.

[21] T. Maho, T. Furon and E. L. Merrer. 'FBI: Fingerprinting models with Benign Inputs'. In: *IEEE Transactions on Information Forensics and Security* (2023), pp. 1–18. DOI: 10.1109/tifs.2023.3301268. URL: https://hal.science/hal-04176514.

[22] J. Rufino, J. M. Ramírez, J. Aguilar, C. Baquero, J. Champati, D. Frey, R. E. Lillo and A. Fernández-Anta. 'Consistent comparison of symptom-based methods for COVID-19 infection detection'. In: *International Journal of Medical Informatics* 177 (Sept. 2023), p. 105133. DOI: 10.1016/j.ijmedinf.2023.105133. URL: https://inria.hal.science/hal-04406757.

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

National journals

[24] E. Le Merrer, R. Pons and G. Trédan. 'Algorithmic audits of algorithms, and the law'. In: *AI and Ethics* (27th Sept. 2023), pp. 1–21. DOI: 10.1007/s43681-023-00343-z. URL: https://inria.hal.science/hal-03583919.

International peer-reviewed conferences

[25] A. Auvolat, Y.-D. Bromberg, D. Frey, D. Mvondo and F. Taïani. 'Basalt: A Rock-Solid Byzantine-Tolerant Peer Sampling for Very Large Decentralized Networks'. In: Middleware 2023 - 24th International Middleware Conference. Bologna, Italy: ACM, 27th Nov. 2023, pp. 111–123. DOI: 10.1145/3590140.3629109. URL: https://inria.hal.science/hal-04394966.

[26] D. Y. C. Chan, G. Giakkoupis and P. Woelfel. 'Word-Size RMR Tradeoffs for Recoverable Mutual Exclusion'. In: PODC 2023 - ACM Symposium on Principles of Distributed Computing. Orlando (FL), United States: ACM, 16th June 2023, pp. 79–89. DOI: 10.1145/3583668.3594597. URL: https://inria.hal.science/hal-04395095.

[27] A. Durand, M. Raynal and G. Taubenfeld. 'Comment se mettre d'accord quand les autres dorment ?' In: AlgoTel 2023 - 25èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications. Cargèse, France, 22nd May 2023, pp. 1–4. URL: https://uca.hal.science/hal-04076960.

[28] D. Frey, M. Gestin and M. Raynal. 'The Synchronization Power (Consensus Number) of Access-Control Objects: the Case of AllowList and DenyList'. In: DISC 2023 - 37th International Symposium on Distributed Computing. L'aquila, Italy: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, pp. 1–32. DOI: 10.4230/LIPIcs.DISC.2023.21. URL: https://inria.hal.science/hal-04399298.

[29] G. Giakkoupis and I. Ziccardi. 'Distributed Self-Stabilizing MIS with Few States and Weak Communication'. In: PODC 2023 - ACM Symposium on Principles of Distributed Computing. Orlando (FL), United States: ACM, 16th June 2023, pp. 310–320. DOI: 10.1145/3583668.3594581. URL: https://inria.hal.science/hal-04393730.

[30] A. Gonzalez, D. Mvondo and Y.-D. Bromberg. 'Takeaways of Implementing a Native Rust UDP Tunneling Network Driver in the Linux Kernel'. In: *Proceedings of the 12th Workshop on Programming Languages and Operating Systems*. PLOS 2023 - 12th Workshop on Programming Languages and Operating Systems. Koblenz, Germany: ACM, 2023. DOI: 10.1145/3623759.3624547. URL: https://hal.science/hal-04235526.

[31] T. Maho, T. Furon and E. Le Merrer. 'Model Fingerprinting with Benign Inputs'. In: ICASSP 2023 - IEEE International Conference on Acoustics, Speech and Signal Processing. ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). Ialyssos, Greece: IEEE, 2023, pp. 1–4. DOI: 10.1109/ICASSP49357.2023.10094751. URL: https://hal.science/hal-04112859.

[32] Y. Martínez-Díaz, L. S. Luevano and H. Méndez-Vázquez. 'Effectiveness of Blind Face Restoration to Boost Face Recognition Performance at Low-Resolution Images'. In: *Lecture Notes in Computer Science*. IWAIPR 2023 - International Workshop on Artificial Intelligence and Pattern Recognition. Vol. 14335. Varadero, Cuba: Springer Nature Switzerland, 20th Dec. 2023, pp. 455–467. DOI: 10.1007/978-3-031-49552-6_39. URL: https://hal.science/hal-04393649.

[33] Y. Martínez-Díaz, H. Méndez-Vázquez, L. S. Luevano and M. González-Mendoza. 'Exploring the Effectiveness of Lightweight Architectures for Face Anti-Spoofing'. In: *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. CVPRW 2023 - IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. Vancouver, Canada, 2023, pp. 6392–6402. DOI: 10.1109/CVPRW59228.2023.00680. URL: https://hal.science/hal-04393650.

[34] L. Paillat, C.-L. Ignat, D. Frey, M. Turuani and A. Ismail. 'Design of an Efficient Distributed Delivery Service for Group Key Agreement Protocols'. In: *Lecture Notes in Computer Science (LNCS)*. FPS 2023 - 16th International Symposium on Foundations & Practice of Security. Bordeaux, France, 11th Dec. 2023, pp. 1–16. URL: https://inria.hal.science/hal-04337821.

Conferences without proceedings

[35] A. Godinot, E. Le Merrer, G. Trédan, C. Penzo and F. Taïani. 'Change-Relaxed Active Fairness Auditing'. In: *Actes de la Conférence Nationale d'Intelligence Artificielle (CNIA) 2023* ⟨https://ut3-toulouseinp.hal.science/hal-04310171⟩. RJCIA 2023 - 21e Rencontres des Jeunes Chercheurs en Intelligence Artificiel. CNIA. Strasbourg, France, July 2023, pp. 91–96. URL: https://hal.science/hal-04395914.

Reports & preprints

[36] T. Albouy, D. Frey, M. Raynal and F. Taïani. *Good-case Early-Stopping Latency of Synchronous Byzantine Reliable Broadcast: The Deterministic Case (Extended Version)*. 7th Mar. 2023. URL: https://inria.hal.science/hal-04017887.

[37] D. Y. C. Chan, G. Giakkoupis and P. Woelfel. *Word-Size RMR Trade-offs for Recoverable Mutual Exclusion*. 16th May 2023. URL: https://inria.hal.science/hal-04098408.

[38] M. Déprés, A. Mostefaoui, M. Perrin and M. Raynal. *Send/Receive Patterns versus Read/Write Patterns: the MB-Broadcast Abstraction (Extended Version)*. 3rd May 2023. URL: https://hal.science/hal-04087447.

Other scientific publications

[39] L. S. Luevano, M. González-Mendoza, Y. Martínez-Díaz and H. Méndez-Vázquez. 'Exploring the Potential for Real-Time Vision Transformer-Level Precision on Face Recognition Scenarios through Binarization on Embedded Systems'. In: CVPRW 2023 - IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. Vancouver, Canada, 2023. URL: https://inria.hal.science/hal-04393662.

12.3 Other

12.4 Cited publications

[40] Y. Afek and E. Gafni. 'Asynchrony from synchrony'. In: *ICDCN*. 2013, pp. 225–239.

[41] A. Ahmed and E. Ahmed. 'A survey on mobile edge computing'. In: *2016 10th International Conference on Intelligent Systems and Control (ISCO)*. Jan. 2016, pp. 1–8. DOI: 10.1109/ISCO.2016.7727082. URL: http://dx.doi.org/10.1109/ISCO.2016.7727082.

[42] T. Allard, D. Frey, G. Giakkoupis and J. Lepiller. 'Lightweight Privacy-Preserving Averaging for the Internet of Things'. In: *M4IOT 2016 - 3rd Workshop on Middleware for Context-Aware Applications in the IoT*. Trento, Italy: ACM, Dec. 2016, pp. 19–22. DOI: 10.1145/3008631.3008635. URL: https://hal.inria.fr/hal-01421986.

[43] 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.

[44] 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.

[45] D. Bernstein. 'Containers and Cloud: From LXC to Docker to Kubernetes'. In: *IEEE Cloud Computing* 1.3 (Sept. 2014), pp. 81–84. DOI: 10.1109/MCC.2014.51. URL: http://dx.doi.org/10.1109/MCC.2014.51.

[46] M. Bertier, D. Frey, R. Guerraoui, A.-M. Kermarrec and V. Leroy. 'The Gossple Anonymous Social Network'. In: *ACM/IFIP/USENIX 11th International Middleware Conference (MIDDLEWARE)*. Ed. by I. Gupta and C. Mascolo. Vol. LNCS-6452. Middleware 2010. Bangalore, India: Springer, Nov. 2010, pp. 191–211. DOI: 10.1007/978-3-642-16955-7_10. URL: https://hal.inria.fr/inria-00515693.

[47] F. Bonomi. *Connected vehicles, the internet of things, and fog computing. VANET 2011, 2011*. Keynote speech at VANET. 2011.

[48] F. Bonomi, R. Milito, J. Zhu and S. Addepalli. 'Fog Computing and Its Role in the Internet of Things'. In: *1st MCC Workshop on Mobile Cloud Computing*. 2012. DOI: 10.1145/2342509.2342513. URL: http://doi.acm.org/10.1145/2342509.2342513.

[49] 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.

[50] A. Boutet, D. Frey, R. Guerraoui, A.-M. Kermarrec and R. Patra. 'HyRec: Leveraging Browsers for Scalable Recommenders'. In: *Middleware 2014*. Bordeaux, France, Dec. 2014. DOI: 10.1145/2663165.2663315. URL: https://hal.inria.fr/hal-01080016.

[51] A. Boutet, D. Frey, R. Guerraoui, A.-M. Kermarrec, A. Rault, F. Taïani and J. Wang. 'Hide & Share: Landmark-based Similarity for Private KNN Computation'. In: *DSN*. Rio de Janeiro, Brazil, 2015. DOI: 10.1109/DSN.2015.60. URL: https://hal.archives-ouvertes.fr/hal-01171492.

[52] 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.

[53] B. Cohen. *Incentives Build Robustness in BitTorrent*. 2003. URL: http://citeseer.ist.psu.edu/cohen03incentives.html.

[54] D. Collins, R. Guerraoui, J. Komatovic, P. Kuznetsov, M. Monti, M. Pavlovic, Y. Pignolet, D. Seredinschi, A. Tonkikh and A. Xygkis. 'Online Payments by Merely Broadcasting Messages'. In: *IEEE DSN*. 2020. DOI: 10.1109/DSN48063.2020.00023. URL: https://doi.org/10.1109/DSN48063.2020.00023.

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

[56] C. Delporte-Gallet, H. Fauconnier, R. Guerraoui and A. Tielmann. 'The disagreement power of an adversary'. In: *Distributed Computing* 24.3-4 (2011), pp. 137–147.

[57] A. J. Demers, D. H. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. E. Sturgis, D. C. Swinehart and D. B. Terry. 'Epidemic Algorithms for Replicated Database Maintenance'. In: *PODC*. 1987, pp. 1–12.

[58] D. Frey, R. Guerraoui, A.-M. Kermarrec, M. Monod, K. Boris, M. Martin and V. Quéma. 'Heterogeneous Gossip'. In: *Middleware 2009*. Urbana-Champaign, IL, United States, Dec. 2009. URL: https://hal.inria.fr/inria-00436125.

[59] W. M. Golab, V. Hadzilacos, D. Hendler and P. Woelfel. 'RMR-efficient implementations of comparison primitives using read and write operations'. In: *Distributed Computing* 25.2 (2012), pp. 109–162.

[60] R. Guerraoui, K. Huguenin, A.-M. Kermarrec, M. Monod and S. Prusty. 'LiFTinG: Lightweight Freerider-Tracking Protocol in Gossip'. In: *11th ACM/IFIP/USENIX International Middleware Conference (MIDDLEWARE)*. Bangalore, India, Nov. 2010. DOI: 10.1007/978-3-642-16955-7_16. URL: https://hal.inria.fr/inria-00505268.

[61] R. Guerraoui, P. Kuznetsov, M. Monti, M. Pavlovic and D. Seredinschi. 'The Consensus Number of a Cryptocurrency'. In: *ACM PODC*. 2019. DOI: 10.1145/3293611.3331589. URL: https://doi.org/10.1145/3293611.3331589.

[62] R. A. Holley and T. M. Liggett. 'Ergodic Theorems for Weakly Interacting Infinite Systems and the Voter Model'. In: *The Annals of Probability* 3.4 (1975), pp. 643–663.

[63] K. Huang, H. Wei, Y. Huang, H. Li and A. Pan. 'Byz-GentleRain: An Efficient Byzantine-tolerant Causal Consistency Protocol'. In: *CoRR* abs/2109.14189 (2021). arXiv: 2109.14189. URL: https://arxiv.org/abs/2109.14189.

[64] D. Imbs and M. Raynal. 'A liveness condition for concurrent objects: x-wait-freedom'. In: *Concurrency and Computation: Practice and experience* 23.17 (2011), pp. 2154–2166.

[65] F. Junqueira and K. Marzullo. 'A framework for the design of dependent-failure algorithms'. In: *Concurrency and Computation: Practice and Experience* 19.17 (2007), pp. 2255–2269.

[66] D. Kempe, J. M. Kleinberg and É. Tardos. 'Influential Nodes in a Diffusion Model for Social Networks'. In: *ICALP*. 2005, pp. 1127–1138.

[67] D. Kempe, J. M. Kleinberg and É. Tardos. 'Maximizing the Spread of Influence through a Social Network'. In: *Theory of Computing* 11 (2015), pp. 105–147.

[68] D. Kempe, J. M. Kleinberg and É. Tardos. 'Maximizing the spread of influence through a social network'. In: *KDD*. 2003, pp. 137–146.

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

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

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

[72] D. Mvondo, A. Tchana, R. Lachaize, D. Hagimont and N. D. Palma. 'Fine-Grained Fault Tolerance for Resilient pVM-Based Virtual Machine Monitors'. In: *50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN*. IEEE, 2020, pp. 197–208. DOI: 10.1109/DSN48063.2020.00037. URL: https://doi.org/10.1109/DSN48063.2020.00037.

[73] D. Mvondo, B. Teabe, A. Tchana, D. Hagimont and N. D. Palma. 'Memory flipping: a threat to NUMA virtual machines in the Cloud'. In: *2019 IEEE Conference on Computer Communications, INFOCOM 2019*. IEEE, 2019, pp. 325–333. DOI: 10.1109/INFOCOM.2019.8737548. URL: https://doi.org/10.1109/INFOCOM.2019.8737548.

[74] F. Pasquale. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard U. Press, 2015.

[75] M. Raynal and J. Stainer. 'Synchrony weakened by message adversaries vs asynchrony restricted by failure detectors'. In: *PODC*. Proceedings of the 2013 ACM symposium on Principles of distributed computing. Montréal, Canada: ACM, July 2013, pp. 166–175. DOI: 10.1145/2484239.2484249. URL: https://hal.inria.fr/hal-00920734.

[76] N. Santoro and P. Widmayer. 'Time is not a healer'. In: *Annual Symposium on Theoretical Aspects of Computer Science*. Springer. 1989, pp. 304–313.

[77] 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.

[78] L. Zhang, F. Zhou, A. Mislove and R. Sundaram. 'Maygh: Building a CDN from Client Web Browsers'. In: *8th ACM European Conference on Computer Systems*. EuroSys '13. Prague, Czech Republic: ACM, 2013, pp. 281–294. DOI: 10.1145/2465351.2465379. URL: http://doi.acm.org/10.1145/2465351.2465379.