Curriculum Vitae

1 Personal Details

Name:	Davide Frey
Date of Birth:	30 March 1976
Nationality:	Italian
Family Status:	Civil Union, 1 child
Current Position:	Research Scientist (CRCN - Chargé de Recherche de Classe Normale)
HDR:	obtained in June 2019
Professional Address:	IRISA/INRIA Rennes Campus Universitaire de Beaulieu 35042 Rennes CEDEX
Email:	davide.frey@inria.fr, davide.frey@irisa.fr
Phone:	+33-02-99847565
Web Site:	http://people.irisa.fr/Davide.Frey/

2 Short Bio

I have been a researcher at Inria Rennes Bretagne-Atlantique since 2010. I received my PhD from Politecnico di Milano in Italy in 2006; I then worked as a post-doctoral researcher both at Washington University in St. Louis (MO), and at Inria Rennes before being recruited as a permanent researcher in 2010. My research interest focus on distributed systems and algorithms with theoretical and practical contributions. On the practical side, my contribution span areas such as content dissemination, social networks, blockchain, and decentralized recommendation and machine-learning systems. On the theoretical side, I have worked on combinatorial optimization and distributed algorithms, particularly in the context of designing lightweight alternatives to blockchains for a variety of applications. I co-supervised 6 PhD students to completion and I am currently supervising and co-supervising 3 PhD students on diverse topics: blockchain, privacy, and decentralized identity management. I have been active in several national and International projects. At the European level, I am currently involved in SOTERIA, an H2020 project, led by AriadNext, for which I am Inria's scientific lead. At the national level, I am currently involved in the Byblos ANR project on Beyond-blockchain data structures, and I coordinate the PriCLeSS Cominlabs projects which sees a collaboration of computer scientists and law researchers to address the privacy and legal implication of the storage of personal data on blockchain-like infrastructures.

3 Current and Previous Positions

1. November 2010 - now.
   Research Scientist (CR1 and then CRCN) in the WIDE (formerly ASAP) team at INRIA Rennes-Bretagne Atlantique, IRISA lab.

2. November 2007 - 2010.
   Postdoc Researcher with the ASAP team at INRIA Rennes-Bretagne Atlantique.
   Position supported through (ANR-RIAM) Solipsis project and Gossple ERC project
   Director of Research Program: Anne-Marie Kermarrec

3. July 2006 - July 2007.
   Postdoc (Visiting Researcher) in the MobiLab.

Department of Computer Science and Engineering.
Washington University in St Louis, MO, USA. Position supported by Office of Naval Research under MURI research contract N00014-02-1-0715.
Director of Research Program: Prof. Dr. Gruia-Catalin Roman

4. March 2003 - May 2006.
Ph.D. Student in Computer Science and Engineering at Dipartimento di Elettronica e Informazione, Politecnico di Milano, Italy, under the supervision of Prof. Gian Pietro Picco.

4 Other activities within Inria

- I was scientific correspondent for H2020 projects at Inria from 2013 to 2020.
- I was a member of the COST-GTRI at Inria from 2016 to 2019.
- I represented Inria at several BDVA (http://www.bdva.eu/) meetings and summits and contributed to the BDVA SRIA in 2015.

5 Education

1. June 2019: Habilitation à diriger des rechercher (HDR): "Epidemic Protocols: From Large Scale to Big Data". University of Rennes 1, France [64].
2. April 2006: PhD "Publish Subscribe on Large-Scale Dynamic Topologies: Routing and Overlay Management". Politecnico di Milano Italy [65].
3. April 2002: Masters Degree. Summa cum Laude. Politecnico di Milano Italy.

6 Student Supervision and Juries

Ongoing PhD Supervisions

- *Timothe Albouy*: started October 2021, co-supervised (50%) with Francois Taiani. Timothe is working byzantine tolerant algorithms for lightweight blockchain-like systems. This work has led to several publications [17, 14, 13].
- *Mathieu Gestin*: started October 2021, supervised (100%). Mathieu is working on decentralized anonymous identity systems in the context of the SOTERIA H2020 project. We published a paper on hiding the issuer's identity in anonymous credentials at PETS 2022 [3].
- *Arthur Rauch*: started October 2021, co-supervised (50%) with Emmanuelle Anceaume. Arthur is working on privacy-conscious storage in blockchain systems, within the context of the PriCLeSS Cominlabs project.

Completed PhD Supervisions I co-supervised the following PhD students to completion.

- *Amaury-Bouchra Pilet*: "Contributions to distributed multi-task machine learning", November 2021, co-supervised (50%) with Francois Taiani.
Amaury worked on decentralized and privacy-preserving machine learning and developed a novel approach to multi-task learning [20] and a protocol to group similar tasks [18]. In addition we published earlier work on privacy preserving averaging[22] and Byzantine resilient peer sampling [21].

- *Quentin Dufour*: "High-throughput real-time onion networks to protect everyone's privacy", February 2021, co-supervised (50%) with David Bromberg.
  Quentin worked on privacy in large scale systems in the context of the O'Browser ANR project. We published a first paper on data dissemination at Infocom 2019 [23] (see Form 3 for details). Then Quentin worked on an anonymous VoIP service that leverages the Tor network [15]. He co-founded non-profit providing low-tech "cloud-like" services running on old hardware.

- *Pierre-Louis Roman*: "Exploring heterogeneity in loosely consistent decentralized data replication", December 2018, co-advised (50%) with Francois Taiani.
  Pierre-Louis worked on the management of heterogeneity in contexts such as recommendation [31], content dissemination [28, 1], and blockchain [49, 63, 56]. Pierre-Louis was then a post-doc researcher at the Univeristà della Svizzera Italiana in Lugano, Switzerland and he is now a post-doc researcher at EPFL, Switzerland.

- *Stephane Delbruel*: "Towards an Architecture for Tag-based Predictive Placement in Distributed Storage Systems", January 2017, co-advised (50%) with Francois Taiani.
  Stephane worked on the use of tags in the context of video placement [26], and proposed a solution for predicting video consumption [50, 27]. He was then a post-doc researcher in the Distrinet group at KU Leuven, Belgium. He is now a post-doc researcher at the University of Oslo, Norway.

- *Antoine Rault*: "User privacy in collaborative filtering systems", June 2016, co-advised (50%) with Anne-Marie Kermarrec.
  Antoine worked on privacy in the context of recommender systems [25, 51, 30] (see Form 2 for details). He did a postdoc at EPFL after leaving Inria.

- *Arnaud Jegou*: "Harnessing the power of implicit and explicit social networks through decentralization", September 2014, co-advised (50%) with Anne-Marie Kermarrec.
  Arnaud worked on decentralized social networks and recommender systems [39, 12, 37, 36, 33, 11, 10]. He is now a R&D engineer at ScaleDynamics, Cesson-Sevigne.

I have also supervised about 20 undergraduate students for master and L3 internships.

PhD Juries I have been reviewer for the following PhD theses.

- Alejandro Ranchal-Pedrosa "The Blockchain of Oz: Specifying Blockchain Failures for Scalable Protocols Offering Unprecedented Safety and Decentralization", University of Sydney, 2022.

- Matthieu Nicolas "Ré-identification sans coordination dans les types de données répliquées sans conflits (CRDTs)", Université de Lorraine, 2022.

- Hayman Salih Mohammed Mohammed "Performance Evaluation Study on Tangle Network", Sapienza University of Rome, 2021.

- Daniele Ucci "Privacy-preserving data sharing in collaborative environments", Sapienza University of Rome, 2017.

In addition, I was also an examiner in the PhD jury of

- Konstantinos Kloudas, "Leveraging Content Properties to Optimize Distributed Storage Systems", University of Rennes, 2013.

7 Invited Talks and Seminars

- Invited talk at the IRISAtech event, 2008.

- Invited talk at the EDBT Summer School, Presqu'île de Giens, France, August/September 2009.

- Invited talk at the Summer School, Masses de donnés distribués, Les Houches (France), May 2010.
- Invited talk at DIS Universita' di Roma, La Sapienza, Italy, May 2009.
- Invited talk at at ENS-Cachan Antenne de Bretagne, France. November, 2009.
- Invited (Keynote) talk at the W-PSDS workshop co-located with SRDS 2016 in Budapest, Hungary, on September 26, 2016.
- Invited talk at the SG-2 meeting of the BDVA Summit in Valencia, Spain, on December 1, 2016.
- Invited talk at the "Session d'information sur les appels à projets 2017 Big Data" at Business France, Paris, on January 6, 2017.
- Invited talk at the WOS7 workshop at Technicolor, on November 30, 2017.
- 2nd AriadNext Workshop on Remote ID Verification. AriadNext, Cesson-Sevigne, on Oct 19, 2021. https://www.eventbrite.fr/e/workshop-ariadnext-remote-identity-verification-tickets-173363473817
- "Donar: Anonymous VoIP over Tor". DistriNet and Cosic teams, KU-Leuven on October 5, 2022 (seminar announcment available here).
- "Asynchronous Byzantine Reliable Broadcast With a Message Adversary". School of Computer, University of Sydney, on July 17, 2022.

8 Science Communication

I co-authored an article on the Telecom ParisTech Magazine. Issue 185: Davide Frey and Alena Siarheyeva "Bubble Computing une architecture Internet des objets décentralisée qui protège la vie privée dans les villes intelligentes".

8.1 Program Committees

8.1.1 Conferences

Distributed and Peer-to-Peer Systems Middleware (core A), ICDCS (core A), and IPDPS (core A) are among the top conferences in distributed systems. Peer-to-peer was a well respected conference in the context of peer-to-peer systems.

Conference	Year	Place	Link
Middleware (ACM/IFIP International Middleware conference)	2023	Bologna, Italy	webpage
	2018 (industrial track)	Rennes, France	webpage
	2016	Trento, Italy	webpage
	2015	Vancouver, Canada	webpage
	2014	Bordeaux, France	webpage
	2013	Beijing, China	proceedings
	2012	Montreal, Canada	proceedings
	2011	Lisbon, Portugal	proceedings

continues on next page

Conference	Year	Place	Link
ICDCS (International Conference on Distributed Computing Systems)	2014	Madrid, Spain	webpage
IPDPS (IEEE International Parallel & Distributed Processing Symposium)	2014 2013	Phoenix, Arizona, USA Boston, Massachusetts US	webpage webpage
EDCC (European Dependable Computing Conference)	2010	Valencia, Spain	webpage
P2P (IEEE International Conference on Peer-to-Peer Computing)	2015 2014 2013 2012	Cambridge, MA, USA London, England Trento, Italy Toronto, Canada	proceedings proceedings proceedings proceedings
DEBS (ACM INTERNATIONAL CONFERENCE ON DISTRIBUTED AND EVENT-BASED SYSTEMS)	2023 2020 2019 2017	Neuchatel, Switzerland Montreal Quebec Darmstadt, Germany Barcelona, Spain	webpage webpage webpage webpage
DAIS (IFIP International Conference on Distributed Applications and Interoperable Systems)	2023 2021 2020 2019 2018 2013	Lisbon, Portugal online online Copenhagen, Denmark Madrid, Spain Florence, Italy	webpage webpage webpage proceedings proceedings proceedings
SBAC PAD (International Symposium on Computer Architecture and High Performance Computing)	2017	Campinas, Brazil	webpage
HiPC (High Performance Computing)	2015	Bangalore, India	webpage

Distributed Algorithms DISC (core A) is a flagship conference in the domain of distributed algorithms.

Conference	Year	Place	Link
DISC (International Symposium on Distributed Computing)	2022	Augusta Georgia, USA	webpage
SSS (International Symposium on Stabilization, Safety, and Security of Distributed Systems)	2012 2010 2009	Toronto, Canada New York, USA Lyon, France	webpage webpage webpage

8.1.2 Workshops

Workshop	Year	Place	Link
BlockDM (International Workshop on Blockchain and Data Management)	2021	online	webpage
SERIAL (Workshop on Scalable and Resilient Infrastructures for Distributed Ledgers)	2020	online	webpage
W-PSDS (Workshop on Planetary-Scale Distributed Systems)	2019	Lyon, France	webpage
MW4NG (Middleware for Next Generation Internet Computing)	2014	Bordeaux, France	webpage
WWW PhD Symposium (PhD Symposium of the International World Wide Web Conference)	2012	Lyon, France	webpage

8.2 Organization of Conferences and Workshops

I was involved in the organization of the following conferences and workshops.

8.2.1 Conferences

Middleware 2015: Workshop and tutorial co-chair for the ACM/IFIP/USENIX International Conference on Middleware (Middleware). http://2015.middleware-conference.org/committees/

ICDCN 2013: program co-chair of the 14th International Conference on Distributed Computing and Networking, Mumbai, January 2013; https://dblp.org/db/conf/icdcn/icdcn2013.html.

8.2.2 Workshops

WOS 2022: Workshop on Streaming, Inria Rennes, 24 November, 2022; https://team.inria.fr/wide/wos/.

WOS 2021: 10th Inria Workshop on Systems, Inria Rennes, 12 October, 2021; https://team.inria.fr/wide/previous-wos/.

SNS 2012: co-chair of the Workshop on Social Networks Systems colocated with Eurosys, Bern, April 2012; https://dblp.org/db/conf/sns/sns2012.html.

SNDS 2010: co-chair of the Workshop on Social Networks and Distributed Systems (SNDS 2010), colocated with PODC 2010, Zurich Switzerland, July 2010.

I served as a referee for the following journals.

- IEEE Transactions on Parallel and Distributed Systems.
- IEEE Transactions on Cloud Computing
- IEEE Transactions on Computers
- IEEE Internet Computing
- Journal of Systems and Software
- IEEE Transactions on Mobile Computing.
- PPNA Peer-to-Peer Networking and Applications.
- IEEE Transactions on Cloud Computing.
- ACM Transactions on Sensor Networks.
- Elsevier International Journal of Computer and Telecommunications Networking (COMNET).
- ARIMA Journal.

9 Teaching Activity

Although my position does not require teaching, I consider teaching as a great opportunity to communicate on major breakthroughs from my research results and those of other researchers. Moreover, teaching courses allows me to reflect more deeply on existing results and in several occasions I have come up with new ideas while teaching or preparing my courses. Below, I report a list of the courses I have taught since my recruitment as a CR in 2010. The specified number of hours is the number of hours I taught in each course (CM= teaching, TD = teaching assistance, TP = lab, ETD = equivalent to teaching assistance), for a total of 505.5 ETD hours.

Year	Course name	Level	University	Hours
2022-23	Cloud Computing Cloud	Master	UR1	6h CM
2022-23	Distributed Algorithms	L3SIF	ENS	11h CM
2022-23	Distributed Systems	Master	ENSAI	21h CM
2012-22	Scalable Distributed Systems	Master	EIT/ICT-Labs Master school (UR1)	10h CM
2021-22	Cloud Computing Cloud	Master	UR1	6h CM
2021-22	Distributed Algorithms	L3SIF	ENS	11h CM
2021-22	Distributed Systems	Master	ENSAI	12h CM
2020-21	Scalable Distributed Systems	Master	EIT/ICT-Labs Master school (UR1)	10h CM
2020-21	Cloud Computing Cloud	Master	UR1	6h CM
2020-21	Big-Data Storage and Proc. Infrastructures	Master	UR1	10h CM
2020-21	Distributed Systems	Master	ENSAI	12h CM
2019-20	Distributed Computing & Blockchain	Master	UM6P	30h (CM+TD+TP)
2019-20	Scalable Distributed Systems	Master	EIT/ICT-Labs Master school (UR1)	10h CM
2019-20	Cloud Computing Cloud	Master	UR1	6h CM
2019-20	Big-Data Storage and Proc. Infrastructures	Master	UR1	10h CM
2019-20	Distributed Systems	Master	ENSAI	12h CM
2018-19	Distributed Computing & Blockchain	Master	UM6P	15h CM
2018-19	Scalable Distributed Systems	Master	EIT/ICT-Labs Master school (UR1)	10h CM
2018-19	Cloud Computing Cloud	Master	UR1	6h CM
2018-19	Big-Data Storage and Proc. Infrastructures	Master	UR1	10h CM
2018-19	Distributed Systems	Master	ENSAI	12h CM
2017-18	Distributed Systems	Master	ENSAI	12h CM
2017-18	Cloud Computing Cloud	Master	UR1	6h CM
2017-18	Scalable Distributed Systems	Master	EIT/ICT-Labs Master school (UR1)	10h CM
2017-18	Programming Technologies for the Cloud	Master	UR1	10h CM + 4h TD + 12h TP
2017-18	Big-Data Storage and Proc. Infrastructures	Master	UR1	10h CM
2016-17	Scalable Distributed Systems	Master	EIT/ICT-Labs Master school (UR1)	10h CM
2015-16	Scalable Distributed Systems	Master	EIT/ICT-Labs Master school (UR1)	10h CM
2014-15	Scalable Distributed Systems	Master	EIT/ICT-Labs Master school (UR1)	10h CM
2014-15	Scalable Distributed Systems	Master	UR1	5h CM
2013-14	Scalable Distributed Systems	Master	EIT/ICT-Labs Master school (UR1)	10h CM
2013-14	Scalable Distributed Systems	Master	UR1	10h CM
2012-13	Scalable Distributed Systems	Master	UR1	10h CM

In addition I have also tutored students of apprenticeship courses (alternance) since 2017.

I also gave invited lectures at Universita' di Roma La Sapienza (Italy), ENS Cachan, Paris, ENS Cachan Antenne de Bretagne, and at Lancaster University (UK).

10 Projects

10.1 SOTERIA H2020

. I am PI for Inria and Task leader in this Innovation Action led by AriadNext by IDnow.SOTERIA will develop and test a citizen-driven and citizen-centric, cost-effective, decentralized data vault allowing citizens to control their private personal data easily and securely. Our taks in the project involve Privacy and data protection, and Hardware-based Privacy Components.

10.2 Byblos ANR

. I am a Task leader for this ANR Project led by Francois Taiani, leader of the WIDE team. Byblos starts from the observation that many applications—including cryptocurrency—do not require full Byzantine agreement, and can be implemented with much lighter, and hence more scalable and more efficient, guarantees. It seeks to design novel algorithms for lightweight abstractions for permissioned (closed) and permissionless (open) systems.

10.3 PriCLeSS Cominlabs

. I coordinate this cross-disciplinary project that sees a partnership between Computer Science and Law researchers. PriCLeSS (Privacy-Conscious Legally-Sound Blockchain Storage) aims to understand and address the legal and technical challenges associated with data storage in a blockchain context. Priceless involves **four** interdisciplinary research teams at **three** different institutions, spanning the topics of **distributed systems**, **distributed algorithms**, **privacy**, and **law**.

10.4 PAMELA ANR

PAMELA is a collaborative ANR project involving ASAP, Inria Lille, UMPC, Mediego and Snips. The project aims at developing machine learning theories and algorithms in order to learn local and personalized models from data distributed over networked infrastructures. This project seeks to provide first answers to modern information systems built by interconnecting many personal devices holding private user data in the search of personalized suggestions and recommendations. In the context of this project, I am co-supervising the PhD thesis of Amaury Bouchra-Pilet.

10.5 OBrowser ANR

I am Principal Investigator for Inria Rennes, in O'Browser, a collaborative ANR project coordinated by Univ. Nantes. The project emerges from the vision of designing and deploying distributed application on millions of machines using web-enabled technologies without relying on a cloud or a central authority. OBrowser proposes to build collaborative applications through a decentralized execution environment composed of users' browsers that autonomously manages issues such as communication, naming, heterogeneity, and scalability. In the context of this project, I am co-supervising the PhD thesis of Quentin Dufour, and I am collaborating with OrangeLabs Lannion on a decentralized architecture for network diagnostics (see Section **??**). I

10.6 BooostEurope

In 2018, I received funding from the Region Bretagne to prepare an H2020 submission on privacy preserving data analytics. I used this funding to prepare the submission to ICT-13b that I mentioned earlier. But I will use the remaining funds to start preparing a new submission to ICT-12-2020.

10.7 DeSceNt CominLabs

The DeSceNt project aims to ease the writing of distributed programs on a federation of plug computers. Plug computers are a new generation of low-cost computers, such as Raspberry pi (25$), VIA- APC (49$), and ZERO Devices Z802 (75$), which offer a cheap and readily available infrastructure to deploy domestic on-line software. Plug computers open the opportunity for everyone to create cheap nano-clusters of domestic servers, host data and services and federate these resources with their friends, colleagues, and families based on social links. I actively worked in the context of Descent with several contributions [48, 26, 50, 9, 25, 30, 34].

10.8 Socioplug

Socioplug was a collaborative ANR project involving INRIA (ASAP team), the university of Nantes, and LIRIS (INSA Lyon and Universite Claude Bernard Lyon). The project emerged from the observation that the features offered by the Web 2.0 or by social media do not come for free. Rather they bring the implicit cost of privacy. Instead of concentrating information of cloud platforms owned by a few economic players, we worked on services made possible by cheap low-end plug computers available in every home or workplace. I was active in Socioplug, particularly with the supervision of the PhD thesis of Pierre-Louis Roman, as well as with several contributions[49, 48, 28, 31, 52, 34, 35]

10.9 Brow2Brow: Browser-to-browser serverless toolboxes

Brow2Brow was an "Action de Development Technologique", i.e. a collaborative development project that aims at providing a middleware and software library for browser-to-browser applications. I co-led Brow2Brow with Stephane Grumbach from the DICE Team from INRIA Grenoble (Antenne de Lyon). The project aimed at providing an alternative to the current model followed by Web2.0 applications by exploiting the recently introduced WebRTC standard. The main outcome of the project on our side consisted of the development of WebGC [57, 58], a library for gossip-based applications on browsers.

10.10 EIT/ICT-Labs AllYours

The AllYours EIT/ICT Labs project was a collaborative initiative that involved the ASAP team, TrentoRise (Italy), and the Eindhoven EIT/ICT nodes. Our work in this project concentrated on refining and testing two versions of the AllYours software: P2P AllYours, and Mobile AllYours. In the context of the project I supervised the work of Heverson Ribeiro who was a post-doc engineer in the team in the development of the p2p-AllYours software (see Section 14).

11 Summary of Research Activities

My research interests revolve around the grand theme of *large-scale systems*, and include several spin-off themes that hightlight my desire and ability to work on a variety of topics. After working on peer-to-peer systems I developed an interest protecting *privacy* in the context of *decentralized recommenders* and more generally in *large-scale systems*. More recently I have also been working on *decentralized machine learning*, which is the focus of the FedMalin Inria challenge in which I am participating. Motivated by the current need for scalable, energy efficient, and sovereign solutions for large-scale applications, I have recently been working on *lightweight Blockchain alternatives* both in practice and theory.

Large-Scale Systems An important topic in my work on large-scale systems consists of gossip (aka epidemic) protocols. I started working on epidemic protocols in the context of video streaming [42, 43, 41] during my postdoc at Inria. I contributed to designing and implementing a large-scale video-streaming application [43] capable of adapting to heterogeneous bandwidth characteristics [42]. Then I contributed to improving this solution with heuristics targeting latency and overall performance [41]. In this context, I started implementing the YALPS and GossipLib libraries which I employed in a lot of my subsequent work within the ASAP and WIDE teams, particularly in the context of decentralized recommenders.

Some of my work on large-scale systems actually draws inspiration on my work on decentralized recommenders. This is the case of Behave [35], an application of the Gossple [40] framework to the optimization of content-delivery networks.

In a similar context, I worked on optimizing data placement in distributed storage systems. In [26], we carried out an extensive analysis of a YouTube dataset and showed that tags can be used to predict the countries from which videos will be viewed. This allowed us to propose a data placement strategy that can optimize video storage. In a later paper [27], I leveraged my experience on Gossple [40] and Behave [35] to design and evaluate a distributed architecture that estimates the aggregated affinity of a new video with all the users in a country. This prediction mechanism makes it possible to place videos in the best locations to optimize download times and bandwidth usage [27].

I also did some work on cluster-based storage. In particular, I collaborated with Kostas Kloudas, a PhD student Anne-Marie Kermarrec, on the design of a cluster-based backup system [38]. We proposed an approach consisting of two techniques. The first consists of a probabilistic method for computing set intersections [106], itself based on probabilistic counting [123]. The second lies in a deduplication-friendly bucket-based load-balancing strategy.

Following the introduction of WebRTC, I worked on turning my GossipLib library into a web-oriented library, through the design of WebGC [57, 58], a WebRTC-based implementation of gossip-based overlay and dissemination protocols. In this case, I supervised the work of an engineer and of two master students who worked on the implementation. This interest on WebRTC then led me to work on Spray [9] a novel peer-sampling protocol designed for browser-based deployments that can be subject to flash crowds. Finally, my most recent work on large-scale decentralized systems consists in applying network coding to epidemic dissemination [23].

Decentralized Recommenders My first contribution on decentralized recommenders dates back to 2010. In Gossple [40], the main paper of Anne-Marie Kermarrec's ERC grant, we defined an architecture to automatically infer personalized connections in an Internet-scale decentralized system and applied it to the problem of query expansion. I worked in collaboration with Vincent Leroy, a PhD student of Anne-Marie Kermarrec and we proposed the use of two layers of gossip protocols, a lower random-peer-sampling layer, and an upper interest-based k-nearest-neighbor layer. We also addressed the problem of anonymity by proposing the use of proxy nodes for exchanging user profiles, a technique we called *gossip on behalf*. In addition to my contributions to the architecture, I implemented the Gossple system on top of the GossipLib and YALPS libraries and ran experiments on the PlanetLab platform.

I also worked on combining the implicit social network defined by Gossple with explicit ones like Facebook. I implemented this idea of mine within the PhD thesis of my first officially co-advised PhD student, Arnaud Jegou. We defined Social Market [39], a platform that allows users to identify and build connections to other

users that can provide interesting goods, or information, while backing up these connections with trust information, obtained by tracing paths on the explicit social network. We later augmented the protocol with privacy mechanisms that provide provable guarantees for the need to hide trust values from third parties [12]. Around this time, I also started my work on the design and the implementation of WhatsUp and HyRec, a decentralized and a hybrid/federated recommender.

I later worked on a framework to develop dynamically adaptive decentralized recommendation systems based on the Gossple model [52, 31]. The most interesting aspect of this work lies in its use of gossip for decentralized coordination. Individual nodes can independently select, and update their own recommendation algorithm, while still collectively contributing to the overall system's mission.

Privacy in the context of large-scale systems. I started being interesed in privacy while working on decentralized recommender systems[11, 30, 10, 25]. More recently, my work on privacy has extended to other areas in large-scale systems. First, I addressed the problem of aonymous voice communication. Several colleagues and I started from the observation that despite there being a large number of deployed solutions for anonymous textual communication, and despite their importance in democratic efforts against totalitarian governments, there existed no deployed solution for anonymous voice calling. We therefore started to investigate the feasibility of anonymous VoIP over the Tor network. We observed that Tor makes it impossible to sustain high-quality reliable voice calls for as little as 30 seconds. We then proposed a middleware solution, built a a wrapper around the Tor client, that leverages the use of multiple Tor circuits to circumvent latency spikes and deliver consistently low and stable latency. This allowed our solution to sustain high-quality calls for as long as 90 minutes [15].

Second, in the context of the SOTERIA H2020 project, I started considering the domain of self-sovereign identities. In this context, we observed that while the SSI models improves privacy through the concept of anonymous credentials [121], existing scheme still enable partial user identification since the identity of the issuer of a credential must be known to a verifier. For example, a hospital restricting access only to vaccinated individuals must know the identity of the issuer of the vaccine certificate (the vaccination center in this case) to determine if it can be trusted. In our work [3], we proposed a novel cryptographic scheme that allows a verifier to determine if it can trust an issuer without knowing or determining its identity.

Decentralized Machine Learning My work on privacy in recommender systems motivated me to explore more general forms of privacy-preserving computation. In particular, I worked on two privacy-preserving averaging protocols [48, 22]. The first [48] applies Shamir's secret sharing scheme to gossip averaging, by splitting each value to be averaged into a set of random addends, and couples it with encryption in the first rounds to resist network adversaries. The second operates by disseminating random noise for a few communication rounds and then subtracting it from the original value to be averaged.

A protocol for decentralized averaging constitutes a central brick for training machine-learning models in a decentralized manner. Here the key idea consists in averaging the model parameters computed by different peers into a single aggregate model. Starting from this observation, in the PhD thesis of Amaury Bouchra-Pilet, we proposed the idea of averaging only subsets of a model across a subset of the peers as a means to address multi-task learning. The algorithm we proposed in [20] was one of the first to combine decentralization with the ability to address multiple tasks in neural networks. Moreover, the protocol can equally well be applied in decentralized and federated settings. In the same line of work, we also proposed a clustering algorithm designed to identify similar tasks in the context of multi-task learning [18]. I am now planning to continue working on this topic in the context of the FedMalin Inria Challenge.

Lightweight Blockchain Alternatives In recent years, I have worked on improving and going beyond Blockchain systems, first from a very practical, and later from a more theoretical perspective. With Pierre-Louis Roman, whom I co-supervised with Francois Taiani, and together with other colleagues, I worked on Dietcoin, an extension to the Bitcoin protocol that makes it possible for lightweight devices such as smartphones to verify the legitimacy of the transactions they are interested in. We published a preliminary version of Dietcoin at ARM 2016 [49], and demoed the full-fledged solution at VLDB 2019 [56]. The latest version is available as a technical report [63].

Unfortunately, even with optimizations, blockchain systems costly in terms of computation, storage, and network requirements (in particular proof-of-* protocols require each message to be disseminated to the entire network within a known bounded time). This led me and other researchers to realize that many applications, including cryptocurrency, do not require the global replication and synchronization that are still implemented in the vast majority of blockchain solutions. This observation marked the start of my work on lightweight blockchain-like systems. First, I worked on the definition of Byzantine tolerant causal broadcast [4], a primitive that was thought necessary to implement a cryptocurrency. But we soon realized that a weaker primitive, namely FIFO broadcast, was sufficient [7]. To this end, we provided the first concurrent specification of money transfer together with a broadcast-based algorithm that we proved correct. In subsequent work, we also formalized a larger set of problems that has the same requirements from a distributed computability perspective [19]. We defined the PC-Ledger (process-commutative ledger), a distributed data structure in which operations issued by different processes commute, we gave an algorithm to build it, and provided a formalization based on Mazurkiewicz's traces. More recently, we started extending our work on money-transfer objects to dynamic systems in which processes can join and leave. This les us to consider a new fault model that had never been studied before in the context of asynchronous systems. In particular, we combined a classical Byzantine adversary [124], with a message adversary, and abstraction that had only been studied in the sychronous message-passing model by Santoro and Widmayer in 1989 [118] and 2007 [101]. A message adversary makes it possible to model message loss resulting from link failures or disconnections (node churn). Our work was the first to consider this model in an asynchronous setting and to combine it with byzantine failures.

12 Ongoing Work and Research Perspectives

My current research represents the natural continuation of the work I carried out in the last ten years. In particular, I identified four main research domains. The first three correspond to the three collaborative projects I am currently involved in: Byblos ANR (led by Francois Taiani), PriCLeSS and SOTERIA, and thus to the PhD theses of Timothe Albouy, Arthur Rauch, and Mathieu Gestin. The fourth is a direction that I am currently pursuing with a master student, and that I plan to continue in the context of a CIFRE PhD grant with Blacknut.

Distributed Algorithms and Blockchain My recent work on the money transfer problem [7] and on byzantine-tolerant broadcast, has awaken my long-standing interest in the theory of distributed algorithms. I plan to continue working in this directions with several objectives.

First, **I plan to continue exploring the notion of message adversary** we started addressing in [17], and study its impact on different distributed objects. In addition, we want to explore variants of the message adversary that can cover the set of behaviors that occur in the presence of Byzantine failures. This more general type of adversary risks to be too strong for practical purposes as already observed by other researchers in a synchronous setting[118]. We thus aim to consider weaker variants that correspond to practically occurring cases.

Second, **I plan to study which other objects and data structures can be implemented without relying on strong assumption like the ability to solve consensus**. Existing work [68] showed that cryptocurrencies can be implemented without consensus and thus without a blockchain. In our work [7], we precisely characterized the problem of money transfer, and showed it does not even require causal ordering. My plan is to extend this type of analysis to other problems that are normally addressed by blockchain systems. Some applications, like self-sovreign identities, may for example benefit from more lightweight implementations than existing blockchains.

Finally, **I plan to explore if and how classical distributed algorithms, and particularly Byzantine fault tolerant ones, can systematically be transformed in protocols for large-scale dynamic systems**. This involves three main challenges: removing the assumption of a known network size, allowing nodes to join and leave arbitrarily and silently, and reducing the complexity of protocols to reach very large network sizes. This will necessarily involve relaxing at least some of the deterministic guarantees of classical algorithms. Thus we plan to take a probabilistic approach, as already proposed by some recent work [67, 75].

But in addition to defining protocols that can operate in such large-scale dynamic environments, we aim to define a set of transformations that can automatically, or semi-automatically transform a classical deterministic algorithm for static networks in to a non-deterministic one for large-scale dynamic networks. I plan to pursue this line of work in the PhD thesis of Timothe Albouy.

Large-Scale Systems In addition to theoretical work on distributed algorithms, I also plan to continue doing more system-level work by leveraging theoretical results. But even without new theoretical contributions, I can already identify some important directions of investigation at the system level.

First, one of the main issues associated with large-scale dynamic systems lies in the vulnerability to Sybil attacks. I plan to leverage my expertise and that of my colleagues in peer-to-peer systems to devise novel countermeasures that can provide cheaper and more equitable alternatives than proof-of-work of proof-of-stake algorithms. To this end we will expand the direction of Byzantine and Sybil resilient peer-sampling protocols that we started exploring in recent years [21]. In particular, we plan to explore how the very structure of the Internet and possible amendments to its protocols can provide ways to implement Sybil resilience.

Second, I plan to work on **distributed storage in the context of blockchain applications**. To the best of our knowledge, most existing BFT large-scale data structures for large-scale, dynamic systems, including distributed-ledger/blockchain implementations, implement a fully replicated model, which can can rapidly lead to formidable costs: the Bitcoin blockchain uses 266GB on each of its 10k full nodes, for a total 2.6PB of storage. The use of sharding, as a means for partial replication, provides an effective way to address this challenge. But even in sharded blockchains, the data within each shard is still fully replicated to provide adequate fault tolerance. This leads to significant use of storage resources and unnecessarily exposes potentially sensitive data to a large number of involved parties.

Finally, in the context of the PriCLeSS project, and thus in the PhD thesis of Arthur Rauch, **I am currently focusing on improving the storage of data within each shard by introducing more frugality and privacy.** To this end, we will devise parsimonious techniques that make it possible to store the content of a shard while striking an effective trade-off between privacy, redundancy, and fault tolerance. In particular, we plan to combine erasure-coding and encryption techniques. The former will make it possible to store the content of shards more effectively, thereby achieving a better trade-off between redundancy and fault tolerance. The latter will make it possible to safeguard the privacy of stored personal data and will enable nodes participating in a partially replicated system to benefit from the same level of trust as nodes that participate in fully replicated ones. Finally, **I plan to design distributed storage protocols that can take into account legal storage requirements** like those expressed by the General Data Protection Regulation (GDPR).

Recommender Systems and Decentralized Learning Research on recommendation algorithms invariably faces the problem of evaluating the performance of a solution with respect to another. Ideally, one would want to evaluate algorithms *online* in their target deployment environment, and this is often what companies that deploy recommender systems do. They expose a set of users to their website with a recommender algorithm and another set of users to the same website with a different recommender algorithm. Metrics such as click rate, sales, or time spent on the website, provide a measure of the effectiveness of the algorithm in engaging users. This technique known as A/B testing represents the web-oriented version of randomized control trials, but it is costly and often unavailable to algorithm designers due to the lack of a deployed system and of a large enough user base.

So proposed algorithms are often evaluated through offline (aka system centric) metrics. The process consists in computing recommendations on an offline dataset and using metrics such as recall and precisions to evaluate their effectiveness. Most recommendation algorithms receive an offline evaluation, but unfortunately, recent research has shown that these offline metrics often offer contradicting results when comparing different algorithms with one another [93, 76]. Our conjecture is that part of this confusion results from the bias induced by the use of datasets that were obtained from a recommender system that already uses a specific algorithm.

In the context of machine learning, and in the context of the FedMalin Inria Challenge, I also plan to carry on the work we started in the PhD thesis of Amaury-Bouchra Pilet on decentralized multi-task learning. Unfortunately, in the context of his thesis we only carried out experiments on small distributed networks and with a limited range of relatively small neural network models. I therefore plan to test the algorithm we

developed on larger neural networks and with a wider range of datasets, and clearly identify the cases in which it can benefit applications.

Privacy and Identity In the context of privacy, I am planning to continue working on the domain of privacy-preserving decentralized computation. This is one of the topics we are going to address in the SOTERIA H2020 project, which funds the PhD thesis of Mathieu Gestin, and which will fund a post-doc researcher I am currently recruiting.

Withing this project we plan to design protocols that can enable **privacy preserving computation on the personal data stored on smartphones**. The idea is that external actors should be able to issue queries and train predictive models using the knowledge available from large sets of users, without infringing the privacy of users. In this context, I plan to collaborate with partners in the SOTERIA consortium such as KU Leuven, for the design, and AriadNext for the deployment of these protocols.

One key aspect in privacy preserving computation lies in unlinkabilty, i.e. in preventing colluding attackers from linking different data items in order to identify a user or to discover confidential information. In this context we plan to employ techniques such as zero-knowledge proofs or blind digital signatures combined with threshold encryption to provide revocable unlinkability. We already started an effort in this direction during the master internship of Mathieu Gestin, in which we started designing an **anonymous credential scheme, which allows the user to hide the issuer of a credential, while being able to convince the service providers they can trust the information produced**. I am also planning to explore the use of hardware-based data protection technologies such as trusted-execution environments to support privacy protection both in the context of the SOTERIA project and in the context of recommender systems. Finally, in the context of the Phd thesis of Mathieu Gestin, we are planning to explore efficient solutions for the Self-Sovereign Identity (SSI) systems, leveraging the results of our work on consensus-less objects mentioned above.

Optimizing Cloud Gaming Services I am starting working with Blacknut, a local company providing a cloud-based video-game streaming platform. In particular, I am planning to work on optimization of their cloud-gaming streaming service. In particular, we plan to work on the optimization of container platforms with the objective to run multiple video-games under the same OS, in the same physical or virtual machine. This is particularly challenging, due to the real-time nature of video-games and their demand for GPU-based computation.

In addition, we are planning to employ prediction and recommendation techniques to identify which games are likely to be played in the near future thereby optimizing their allocation. These topics will likely be explored in the context of a CIFRE thesis.

13 Publications

As customary in the distributed systems community all papers list authors in alphabetical order except for the papers regarding my work in operations research. For those, I explained the extent of my contribution in Form 1-Section 7 and Form 2. Also, like in many other CS communities conferences are generally more competitive than journal and they constitute my major target for publication. I highlighted rank A conferences below. Besides the rating, top conferences in the domains of distributed algorithms and systems and privacy include DISC, Infocom, NSDI, Middleware, IPDPS, ICDCS, DSN, and PETS/PoPETS. Each entry highlights my contribution using the following code.
SC: scientific guidance, CT: core technical contribution, EX: experiments, WT: writing

13.1 International journals

[1] Davide Frey, Achour Mostefaoui, Matthieu Perrin, Pierre-Louis Roman, and François Ta'iani. "Differentiated Consistency for Worldwide Gossips". In: *IEEE Transactions on Parallel and Distributed Systems* 34.1 (2023), pp. 1–15. DOI: 10.1109/TPDS.2022.3209150.
Contribution: SC,CT,WT

[2] Jesús Rufino et al. "Using survey data to estimate the impact of the omicron variant on vaccine efficacy against COVID-19 infection". In: *Scientific Reports* 13.1 (Jan. 2023), p. 900. ISSN: 2045-2322. DOI: 10.1038/s41598-023-27951-3. URL: https://doi.org/10.1038/s41598-023-27951-3.
Contribution: SC,WT

[3] Daniel Bosk, Davide Frey, Mathieu Gestin, and Guillaume 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.
Contribution: SC,WT

[4] Alex Auvolat, Davide Frey, Michel Raynal, and François Taïani. "Byzantine-tolerant causal broadcast". In: *Theoretical Computer Science* 885 (2021), pp. 55–68. DOI: 10.1016/j.tcs.2021.06.021. URL: https://doi.org/10.1016/j.tcs.2021.06.021.
Rating: A
Contribution: SC,WT

[5] Carlos Baquero et al. "The CoronaSurveys System for COVID-19 Incidence Data Collection and Processing". In: *Frontiers Comput. Sci.* 3 (2021), p. 641237. DOI: 10.3389/fcomp.2021.641237. URL: https://doi.org/10.3389/fcomp.2021.641237.
Rating: IF=1.039
Contribution: EX,WT

[6] Augusto Garcia-Agundez et al. "Estimating the COVID-19 Prevalence in Spain With Indirect Reporting via Open Surveys". In: *Frontiers in Public Health* 9 (2021), p. 306. ISSN: 2296-2565. DOI: 10.3389/fpubh.2021.658544. URL: https://www.frontiersin.org/article/10.3389/fpubh.2021.658544.
Rating: IF=3.018
Contribution: EX,WT

[7] A. Auvolat, D. Frey, and F. Taiani M. Raynal. "Money Transfer Made Simple: a Specification, a Generic Algorithm, and its Proof". In: *Bull. EATCS* 132 (2020). http://bulletin.eatcs.org/index.php/beatcs/issue/view/34, pp. 21–44.
Contribution: SC,CT,WT

[8] Borzou Rostami, André Chassein, Michael Hopf, Davide Frey, Christoph Buchheim, Federico Malucelli, and Marc Goerigk. "The quadratic shortest path problem: complexity, approximability, and solution methods". In: *European Journal of Operational Research* 268.2 (July 2018), pp. 473 –485. DOI: 10.1016/j.ejor.2018.01.054. URL: https://hal.inria.fr/hal-01781605.
Rating: A
Contribution: CT,EX,WT

[9] Brice Nédelec, Julian Tanke, Pascal Molli, Achour Mostefaoui, and Davide Frey. "An Adaptive Peer-Sampling Protocol for Building Networks of Browsers". In: *World Wide Web* 25 (2017), p. 1678. DOI: 10.1007/s11280-017-0478-5. URL: https://hal.inria.fr/hal-01619906.
Rating: A
Contribution: SC,WT

[10] Antoine Boutet, Davide Frey, Rachid Guerraoui, Arnaud Jégou, and Anne-Marie Kermarrec. "Privacy-Preserving Distributed Collaborative Filtering". In: *Computing*. Special Issue on NETYS 2014 98.8 (Aug. 2016), pp. 827–846. DOI: 10.1007/s00607-015-0451-z. URL: https://hal.inria.fr/hal-01251314.
Rating: A (Special Issue)
Contribution: SC,CT,WT

[11] Antoine Boutet, Davide Frey, Arnaud Jégou, Anne-Marie Kermarrec, and Heverson Ribeiro. "FreeRec: an Anonymous and Distributed Personalization Architecture". In: *Computing* (Dec. 2015). URL: https://hal.inria.fr/hal-00909127.
Rating: A (Special Issue)
Contribution: SC,CT,WT,EX

[12] Davide Frey, Arnaud Jégou, Anne-Marie Kermarrec, Michel Raynal, and Julien Stainer. "Trust-Aware Peer Sampling: Performance and Privacy Tradeoffs". In: *Theoretical Computer Science* (Feb. 2013). URL: https://hal.inria.fr/hal-00872996.
Rating: A (Special Issue)
Contribution: SC,CT,WT

13.2 Reviewed international conferences

[13] Timothé Albouy, Davide Frey, Michel Raynal, and François Taïani. "A Modular Approach to Construct Signature-Free BRB Algorithms under a Message Adversary". In: *CONFERENCE ON PRINCIPLES OF DISTRIBUTED SYSTEMS, OPODIS 2022*. Vol. 253. LIPIcs, 2022.
Contribution: SC,CT,EX,WT

[14] Timothé Albouy, Davide Frey, Michel Raynal, and François Taïani. "Good-Case Early-Stopping Latency of Synchronous Byzantine Reliable Broadcast: The Deterministic Case". In: *36th International Symposium on Distributed Computing, DISC 2022, October 25-27, 2022, Augusta, Georgia, USA*. Ed. by Christian Scheideler. Vol. 246. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 4:1–4:22. DOI: 10.4230/LIPIcs.DISC.2022.4. URL: https://doi.org/10.4230/LIPIcs.DISC.2022.4.
Contribution: SC,CT

[15] Yérom-David Bromberg, Quentin Dufour, Davide Frey, and Etienne Rivière. "Donar: Anonymous VoIP over Tor". In: *19th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2022, Renton, WA, USA, April 4-6, 2022*. Ed. by Amar Phanishayee and Vyas Sekar. USENIX Association, 2022, pp. 249–265. URL: https://www.usenix.org/conference/nsdi22/presentation/bromberg.
Contribution: SC,WT

[16] Matthieu Pigaglio, Joachim Bruneau-Queyreix, Yérom-David Bromberg, Davide Frey, Etienne Rivière, and Laurent Réveillère. "RAPTEE: Leveraging trusted execution environments for Byzantine-tolerant peer sampling services". In: *42nd IEEE International Conference on Distributed Computing Systems, ICDCS 2022, Bologna, Italy, July 10-13, 2022*. IEEE, 2022, pp. 603–613. DOI: 10.1109/ICDCS54860.2022.00064. URL: https://doi.org/10.1109/ICDCS54860.2022.00064.
Contribution: SC,CT,WT

[17] Timothé Albouy, Davide Frey, Michel Raynal, and Francois Taiani. "Byzantine-tolerant reliable broadcast in the presence of silent churn". In: *Invited Paper - SSS 2021. Virtual Conference*. 2021.
Contribution: SC,WT

[18] Amaury Bouchra Pilet, Davide Frey, and François Taïani. "AUCCCR: Agent Utility Centered Clustering for Cooperation Recommendation". In: *NETYS*. working paper or preprint. Marrakesh, Morocco, May 2021. URL: https://hal.archives-ouvertes.fr/hal-03181696.
Contribution: SC,WT

[19] Davide Frey, Lucie Guillou, Michel Raynal, and François Taïani. "Consensus-Free Ledgers When Operations of Distinct Processes are Commutative". In: *Parallel Computing Technologies - 16th International Conference, PaCT 2021, Kaliningrad, Russia, September 13-18, 2021, Proceedings*. Ed. by Victor Malyshkin. Vol. 12942. Lecture Notes in Computer Science. Springer, 2021, pp. 359–370. DOI: 10.1007/978-3-030-86359-3_27. URL: https://doi.org/10.1007/978-3-030-86359-3_27.
Contribution: SC,WT

[20] Amaury Bouchra Pilet, Davide Frey, and François Taïani. "Simple, Efficient and Convenient Decentralized Multi-task Learning for Neural Networks". In: *Advances in Intelligent Data Analysis XIX - 19th International Symposium on Intelligent Data Analysis, IDA 2021, Porto, Portugal, April 26-28, 2021, Proceedings*. Ed. by Pedro Henriques Abreu, Pedro Pereira Rodrigues, Alberto Fernández, and João Gama. Vol. 12695. Lecture Notes in Computer Science. Springer, 2021, pp. 37–49. DOI: 10.1007/978-3-030-74251-5_4. URL: https://doi.org/10.1007/978-3-030-74251-5_4.
Contribution: SC,WT

[21] A. B. Pilet, D. Frey, and F. Taiani. "Foiling Sybils with HAPS in Permissionless Systems: An Address-based Peer Sampling Service". In: *2020 IEEE Symposium on Computers and Communications (ISCC)*. 2020, pp. 1–6. DOI: 10.1109/ISCC50000.2020.9219606.
Contribution: SC,WT

[22] Amaury Bouchra Pilet, Davide Frey, and Francois Taiani. "Robust Privacy-Preserving Gossip Averaging". In: *Stabilization, Safety, and Security of Distributed Systems*. Ed. by Mohsen Ghaffari, Mikhail Nesterenko, Sébastien Tixeuil, Sara Tucci, and Yukiko Yamauchi. Cham: Springer International Publishing, 2019, pp. 38–52. ISBN: 978-3-030-34992-9.
Contribution: SC,WT

[23] Yérom-David Bromberg, Quentin Dufour, and Davide Frey. "Multisource Rumor Spreading with Network Coding". In: *INFOCOM 2019*. Paris, France, Apr. 2019. URL: https://hal.inria.fr/hal-01946632.
Rating: A*
Contribution: SC,CT,WT

[24] Alejandro Gómez-Boix, Davide Frey, Yérom-David Bromberg, and Benoit Baudry. "A Collaborative Strategy for Mitigating Tracking through Browser Fingerprinting". In: *Proceedings of the 6th ACM Workshop on Moving Target Defense, MTD@CCS 2019, London, UK, November 11, 2019*. 2019, pp. 67–78. DOI: 10.1145/3338468.3356828. URL: https://doi.org/10.1145/3338468.3356828.
Contribution: SC,WT

[25] Antoine Boutet, Florestan De Moor, Davide Frey, Rachid Guerraoui, Anne-Marie Kermarrec, and Antoine Rault. "Collaborative Filtering Under a Sybil Attack: Similarity Metrics do Matter!" In: *DSN 2018 - the 48th International Conference on Dependable Systems and Networks*. Luxembourg, Luxembourg, June 2018, pp. 1–12. URL: https://hal.inria.fr/hal-01787060.
Rating: A
Contribution: SC,CT,WT

[26] Stéphane Delbruel, Davide Frey, and François Taïani. "Exploring The Use of Tags for Georeplicated Content Placement". In: *IEEE IC2E'16*. Best paper award. Berlin, Germany, Apr. 2016, pp. 172–181. DOI: 10.1109/IC2E.2016.37. URL: https://hal.inria.fr/hal-01257939.
Contribution: SC,WT

[27] Stéphane Delbruel, Davide Frey, and François Taïani. "Mignon: A Fast Decentralized Content Consumption Estimation in Large-Scale Distributed Systems". In: *16th IFIP WG 6.1 International Conference on Distributed Applications and Interoperable Systems (DAIS)*. Ed. by Márk Jelasity and Evangelia Kalyvianaki. Vol. LNCS-9687. Distributed Applications and Interoperable Systems. Heraklion, Greece: Springer, June 2016, pp. 32–46. DOI: 10.1007/978-3-319-39577-7_3. URL: https://hal.inria.fr/hal-01301230.
Rating: B
Contribution: SC,CT,WT

[28] Davide Frey, Achour Mostefaoui, Matthieu Perrin, Pierre-Louis Roman, and François Taïani. "Speed for the elite, consistency for the masses: differentiating eventual consistency in large-scale distributed systems". In: *Proceedings of the 2016 IEEE 35th Symposium on Reliable Distributed Systems (SRDS 2016)*. Budapest, Hungary: IEEE Computer Society, Sept. 2016, pp. 197–206. DOI: 10.1109/SRDS.2016.032. URL: https://hal.inria.fr/hal-01344138.
Rating: A
Contribution: SC,CT,WT

[29] Hicham Lakhlef, Davide Frey, and Michel Raynal. "Optimal Collision/Conflict-Free Distance-2 Coloring in Wireless Synchronous Broadcast/Receive Tree Networks". In: *45th International Conference on Parallel Processing*. Philadelphia, PA, United States, Aug. 2016, pp. 350–359. DOI: 10.1109/ICPP.2016.47. URL: https://hal.archives-ouvertes.fr/hal-01396940.
Rating: A
Contribution: SC,WT

[30] Antoine Boutet, Davide Frey, Rachid Guerraoui, Anne-Marie Kermarrec, Antoine Rault, François Taïani, and Jingjing Wang. "Hide & Share: Landmark-based Similarity for Private KNN Computation". In: *45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. Rio de Janeiro, Brazil, June 2015, pp. 263–274. DOI: 10.1109/DSN.2015.60. URL: https://hal.archives-ouvertes.fr/hal-01171492.
Rating: A
Contribution: SC,CT,WT

[31] Davide Frey, Anne-Marie Kermarrec, Christopher Maddock, Andreas Mauthe, Pierre-Louis Roman, and François Taïani. "Similitude: Decentralised Adaptation in Large-Scale P2P Recommenders". In: *15th IFIP International Conference on Distributed Applications and Interoperable Systems (DAIS)*. Ed. by Alysson Bessani and Sara Bouchenak. Vol. LNCS-9038. Distributed Applications and Interoperable Systems. Grenoble, France: Springer International Publishing, June 2015, pp. 51–65. DOI: 10.1007/978-3-319-19129-4_5. URL: https://hal.inria.fr/hal-01138365.
Rating: B
Contribution: SC,CT,WT

[32] Borzou Rostami, Federico Malucelli, Davide Frey, and Christoph Buchheim. "On the Quadratic Shortest Path Problem". In: *14th International Symposium on Experimental Algorithms*. 14th International Symposium on Experimental Algorithms. Paris, France, June 2015. DOI: 10.1007/978-3-319-20086-6_29. URL: https://hal.inria.fr/hal-01251438.
Rating: B
Contribution: CT,EX,WT

[33] Antoine Boutet, Davide Frey, Rachid Guerraoui, Arnaud Jégou, and Anne-Marie Kermarrec. "Privacy-Preserving Distributed Collaborative Filtering". In: *NETYS*. Marrakech, Morocco, May 2014, pp. 169–184. DOI: 10.1007/978-3-319-09581-3_12. URL: https://hal.inria.fr/hal-00975137.
Contribution: SC,CT,WT

[34] Antoine Boutet, Davide Frey, Rachid Guerraoui, Anne-Marie Kermarrec, and Rhicheek Patra. "HyRec: Leveraging Browsers for Scalable Recommenders". In: *Middleware 2014*. Bordeaux, France, Dec. 2014, pp. 85–96. DOI: 10.1145/2663165.2663315. URL: https://hal.inria.fr/hal-01080016.

Rating: A
Contribution: SC,WT

[35] Davide Frey, Mathieu Goessens, and Anne-Marie Kermarrec. "Behave: Behavioral Cache for Web Content". In: *4th International Conference on Distributed Applications and Interoperable Systems (DAIS)*. Ed. by Kostas Magoutis and Peter Pietzuch. Vol. LNCS 8460. Lecture Notes in Computer Science. Berlin, Germany: Springer, June 2014, pp. 89–103. DOI: 10.1007/978-3-662-43352-2_8. URL: https://hal.inria.fr/hal-01079976.
Rating: B
Contribution: SC,CT,WT

[36] Antoine Boutet, Davide Frey, Rachid Guerraoui, Arnaud Jégou, and Anne-Marie Kermarrec. "WhatsUp Decentralized Instant News Recommender". In: *IPDPS 2013*. Boston, United States, May 2013. URL: https://hal.inria.fr/hal-00769291.
Rating: A
Contribution: CT,WT,EX

[37] Antoine Boutet, Davide Frey, Arnaud Jégou, Anne-Marie Kermarrec, and Heverson Borba Ribeiro. "FreeRec: an Anonymous and Distributed Personalization Architecture". In: *NETYS*. Marrakesh, Morocco, May 2013. URL: https://hal.inria.fr/hal-00820377.
Contribution: SC,CT,WT,EX

[38] Davide Frey, Anne-Marie Kermarrec, and Konstantinos Kloudas. "Probabilistic Deduplication for Cluster-Based Storage Systems". In: *ACM Symposium on Cloud Computing*. San Jose, CA, United States, Oct. 2012, p. 17. DOI: 10.1145/2391229.2391246. URL: https://hal.inria.fr/hal-00728215.
Contribution: SC,WT

[39] Davide Frey, Arnaud Jégou, and Anne-Marie Kermarrec. "Social Market: Combining Explicit and Implicit Social Networks". In: *International Symposium on Stabilization, Safety, and Security of Distributed Systems*. Grenoble, France: LNCS, Oct. 2011, pp. 193–207. DOI: 10.1007/978-3-642-24550-3_16. URL: https://hal.inria.fr/inria-00624129.
Rating: C
Contribution: SC,CT,WT

[40] Marin Bertier, Davide Frey, Rachid Guerraoui, Anne-Marie Kermarrec, and Vincent Leroy. "The Gossple Anonymous Social Network". In: *ACM/IFIP/USENIX 11th International Middleware Conference (MIDDLEWARE)*. Ed. by Indranil Gupta; Cecilia 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.
Rating: A
Contribution: CT,EX,WT

[41] Davide Frey, Rachid Guerraoui, Anne-Marie Kermarrec, and Maxime Monod. "Boosting Gossip for Live Streaming". In: *P2P 2010*. Delft, Netherlands, Aug. 2010, pp. 1–10. DOI: 10.1109/P2P.2010.5569962. URL: https://hal.inria.fr/inria-00517384.
Rating: C
Contribution: CT,EX,WT

[42] Davide Frey, Rachid Guerraoui, Anne-Marie Kermarrec, Maxime Monod, Koldehofe Boris, Mogensen Martin, and Vivien Quéma. "Heterogeneous Gossip". In: *Middleware 2009*. Urbana-Champaign, IL, United States, Dec. 2009, pp. 42–61. DOI: 10.1007/978-3-642-10445-9_3. URL: https://hal.inria.fr/inria-00436125.
Rating: A
Contribution: CT,EX,WT

[43] Davide Frey, Rachid Guerraoui, Anne-Marie Kermarrec, Maxime Monod, and Vivien Quéma. "Stretching Gossip with Live Streaming". In: *DSN 2009*. Estoril, Portugal, June 2009, pp. 259–264. DOI: 10.1109/DSN.2009.5270330. URL: https://hal.inria.fr/inria-00436130.
Rating: A
Contribution: CT,EX,WT

[44] Davide Frey and Amy L. Murphy. "Failure-Tolerant Overlay Trees for Large-Scale Dynamic Networks". In: *8th International Conference on Peer-to-Peer Computing 2008 (P2P'08)*. Aachen, Germany, Sept. 2008, pp. 351–361. DOI: 10.1109/P2P.2008.30. URL: https://hal.inria.fr/inria-00337054.
Rating: C
Contribution: CT,EX,WT

[45] Davide Frey and Gruia-Catalin Roman. "Context-Aware Publish Subscribe in Mobile ad Hoc Networks". In: *Coordination*. Paphos, Cyprus, June 2007, pp. 37–55. DOI: 10.1007/978-3-540-72794-1_3. URL: https://hal.inria.fr/hal-00739641.
Rating: B
Contribution: CT,WT,EX

[46] Gianpaolo Cugola, Davide Frey, Amy L. Murphy, and Gian Pietro Picco. "Minimizing the Reconfiguration Overhead in Content-Based Publish-Subscribe". In: *Symposium on Applied Computing*. Nicosia, Cyprus, Mar. 2004. URL: https://hal.inria.fr/hal-00739607.
Rating: B
Contribution: CT,WT,EX

13.3 Books and book chapters

[47] Benoit Baudry, Yérom-David Bromberg, Davide Frey, Alejandro Gómez-Boix, Pierre Laperdrix, and François Taïani. "Profilage de navigateurs : état de l'art et contre-mesures". In: *Le profilage en ligne : entre libéralisme et régulation*. Ed. by Sandrine Turgis, Alexandra Bensamoun, and Maryline Boizard. Mare et Martin, Oct. 2020. URL: https://hal.inria.fr/hal-03043187.
Contribution: SC,WT

13.4 Other international publications (posters, short papers)

Workshops

[48] Tristan Allard, Davide Frey, George Giakkoupis, and Julien 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.
Contribution: SC,CT,WT

[49] Davide Frey, Marc X. Makkes, Pierre-Louis Roman, François Taïani, and Spyros Voulgaris. "Bringing secure Bitcoin transactions to your smartphone". In: *Proceedings of the 15th International Workshop on Adaptive and Reflective Middleware (ARM 2016)*. Trento, Italy: ACM, Dec. 2016, 3:1–3:6. DOI: 10.1145/3008167.3008170. URL: https://hal.inria.fr/hal-01384461.
Contribution: SC,CT,WT

[50] Stéphane Delbruel, Davide Frey, and François Taïani. "Decentralized view prediction for global content placement". In: *Middleware 2015 : ARM Workshop*. Vancouver, Canada, Dec. 2015. DOI: 10.1145/2834965.2834974. URL: https://hal.inria.fr/hal-01247159.
Contribution: SC,CT,WT

[51] Davide Frey, Rachid Guerraoui, Anne-Marie Kermarrec, and Antoine Rault. "Collaborative Filtering Under a Sybil Attack: Analysis of a Privacy Threat". In: *Eighth European Workshop on System Security EuroSec 2015*. Bordeaux, France, Apr. 2015. DOI: 10.1145/2751323.2751328. URL: https://hal.inria.fr/hal-01158723.
Contribution: SC,CT,WT

[52] Davide Frey, Anne-Marie Kermarrec, Christopher Maddock, Andreas Mauthe, and François Taïani. "Adaptation for the Masses: Towards Decentralized Adaptation in Large-Scale P2P Recommenders". In: *Workshop on Adaptive and Reflective Middleware ARM 2014*. Bordeaux, France, Dec. 2014, 4:1–4:6. DOI: 10.1145/2677017.2677021. URL: https://hal.inria.fr/hal-01080030.
Contribution: SC,WT

[53] Davide Frey, Jérôme Royan, Romain Piegay, Anne-Marie Kermarrec, Emmanuelle Anceaume, and Fabrice Le Fessant. "Solipsis: A Decentralized Architecture for Virtual Environments". In: *1st International Workshop on Massively Multiuser Virtual Environments*. Reno, NV, United States, Mar. 2008. URL: https://hal.inria.fr/inria-00337057.
Contribution: CT,WT

[54] Davide Sormani, Gabriele Turconi, Paolo Costa, Davide Frey, Matteo Migliavacca, and Luca Mottola. "Towards lightweight information dissemination in inter-vehicular networks". In: *Proceedings of the 3rd international workshop on Vehicular ad hoc networks*. Los Angeles, United States, Sept. 2006. URL: https://hal.inria.fr/hal-00739632.
Contribution: SC,CT,WT

[55] Davide Frey and Paolo Costa. "Publish-subscribe tree maintenance over a DHT". In: *DEBS 2005 Workshop colocated with ICDCS*. Columbus, OHIO, United States, June 2005. URL: https://hal.inria.fr/hal-00739617.
Contribution: CT,EX,WT

Demos

[56] Davide Frey, Marc X. Makkes, Pierre-Louis Roman, François Taïani, and Spyros Voulgaris. *Dietcoin: Hardening Bitcoin Transaction Verification Process For Mobile Devices*. 2019. DOI: 10.14778/3352063.3352106. URL: http://www.vldb.org/pvldb/vol12/p1946-frey.pdf.
Rating: A*
Contribution: SC,CT,WT

[57] Raziel Carvajal-Gómez, Davide Frey, Matthieu Simonin, and Anne-Marie Kermarrec. *WebGC Gossiping on Browsers without a Server [Live Demo/Poster]*. Web Information System Engineering. Nov. 2015. URL: https://hal.inria.fr/hal-01251787.
Contribution: SC,CT,WT

[58] Raziel Carvajal-Gómez, Davide Frey, Matthieu Simonin, and Anne-Marie Kermarrec. *WebGC: Browser-Based Gossiping [Live Demo/Poster]*. Middleware 2014. Dec. 2014. DOI: 10.1145/2678508.2678515. URL: https://hal.inria.fr/hal-01080032.
Contribution: SC,CT,WT

[59] Antoine Boutet, Davide Frey, Rachid Guerraoui, and Anne-Marie Kermarrec. *WhatsUp: news from, for, through everyone*. Delft, Netherlands, Aug. 2010. DOI: 10.1109/P2P.2010.5569981. URL: https://hal.inria.fr/inria-00515420.
Contribution: SC,CT,EX,WT

13.5 Research reports and publications under review

[60] Timothé Albouy, Davide Frey, Michel Raynal, and François Taïani. "Asynchronous Byzantine Reliable Broadcast With a Message Adversary". working paper or preprint. May 2022. URL: https://hal.inria.fr/hal-03671451.
Contribution: SC,WT

[61] Alex Auvolat, Yérom-David Bromberg, Davide Frey, and François Taïani. "*BASALT*: A Rock-Solid Foundation for Epidemic Consensus Algorithms in Very Large, Very Open Networks". working paper. Feb. 2021. URL: https://hal.inria.fr/hal-03131734.
Contribution: SC,WT

[62] Carlos Baquero et al. “Measuring Icebergs: Using Different Methods to Estimate the Number of COVID-19 Cases in Portugal and Spain”. 2020. DOI: 10.1101/2020.04.20.20073056. eprint: https://www.medrxiv.org/content/early/2020/04/23/2020.04.20.20073056.full.pdf. URL: https://www.medrxiv.org/content/early/2020/04/23/2020.04.20.20073056.
Contribution: EX,WT

[63] Davide Frey, Marc X. Makkes, Pierre-Louis Roman, François Taïani, and Spyros Voulgaris. “Dietcoin: shortcutting the Bitcoin verification process for your smartphone”. Research Report. Mar. 2018. URL: https://hal.inria.fr/hal-01743995.
Contribution: SC,CT,WT

[64] Davide Frey. “Epidemic Protocols: From Large Scale to Big Data”. Habilitation à diriger des recherches. Université De Rennes 1, June 2019. URL: https://hal.inria.fr/tel-02375909.
Contribution: SC,CT,EX,WT

[65] Davide Frey. “Publish Subscribe on Large-Scale Dynamic Topologies: Routing and Overlay Management”. Theses. Politecnico di Milano, May 2006. URL: https://tel.archives-ouvertes.fr/tel-00739652.

14 Software

YALPS

APP: IDDN.FR.001.500003.000.S.P.2013.000.10000 on 09/12/2013
Contribution: 35% (design, implementation)

YALPS is an open-source Java library designed to facilitate the development, deployment, and testing of distributed applications. Applications written using YALPS can be run both in simulation and in real-world mode without changing a line of code or even recompiling the sources. A simple change in a configuration file will load the application in the proper environment. A number of feature make YALPS useful both for the design and evaluation of research prototypes and for the development of applications to be released to the public.

GossipLib

APP: IDDN.FR.001.500001.000.S.P.2013.000.10000 on 09/12/2013
Contribution: 50% (design implementation)

GossipLib is a library built on top of YALPS and consisting of a set of Java classes aimed to facilitate the development of gossip-based application in a large-scale setting. The current version of GossipLib provides the implementation of a peer-sampling protocol, as well as a demo application enabling the visualization of the execution of the protocols. The architecture of GossipLib is designed to facilitate code-reuse. Each gossip-based component may be used as a building block to develop new and more complex protocols. We employed GossipLib in a number of projects in the ASAP team as detailed in Section 11.

HEAP: Heterogeneity-Aware Gossip Protocol

Contribution: 50% (design, implementation)

HEAP is the implementation of the protocol described in[42, 41]. It provides a video-streaming platform particularly suited for environment characterized by heterogeneous bandwidth capabilities such as those comprising ADSL edge nodes. HEAP is, in fact, able to dynamically leverage the most capable nodes and increase their contribution to the protocol, while decreasing by the same proportion that of less capable nodes.

Peer-to-peer AllYours/WhatsUp: distributed News Recommender

APP: IDDN.FR.001.500002.000.S.P.2013.000.30000 on 09/12/2013
Contribution: 30% (design, implementation)

AllYours-P2P is the the implementation of WhatsUp [36] refined in the context of the AllYours EIT/ICT-Labs project. It consists of a peer-to-peer based news recommender system that organizes users into an implicit social network based on their explicit opinions. In AllYours-P2P the recommendation process is collaboratively performed by connected users. The AllYours-P2P software consists of two parts, running on each peer: an embedded application server, based on Jetty, and a web interface accessible from any web browser. The back-end is written in Java, while the user interface comprises HTML and Javascript code. AllYours-P2P is currently available in three different platforms: Mac OSx (10.5 or latter), Windows (Vista and Windows 7) and Linux (Ubuntu 10.4 or latter). We tested Allyours-p2p in a real life environment with a set of invited users in Italy in Autumn 2013 and Spring 2014, these test were a part of joint project between ASAP Team and its Italian partner Trentorise.

HyRec: A hybrid recommender system

APP: IDDN.FR.001.500007.000.S.P.2013.000.30000 on 09/12/2013
Contribution: 10% (scientific advice, design)

HyRec implements the hybrid recommender system described in [34]. The motivation of this work is to explore solutions that could in some sense democratize personalization by making it accessible to any content provider company without generating huge investments. HyRec implements a user-based collaborative filtering

scheme and offloads CPU-intensive recommendation tasks to front-end client browsers, while retaining storage and orchestration tasks within back-end servers. HyRec seeks to provide the scalability of p2p approaches without forcing content providers to give up the control of the system.

WebGC: Web-based Gossip Communication

APP:IDDN.FR.001.120008.000.S.P.2017.000.10600 in May 2017
Contribution: 20% (scientific advice, architecture)

WebGC is a library for gossip-based communication between web-browsers. It has been developed in collaboration with Mathieu Simonin in the context of the Brow2Brow ADT project. WebGC builds on the recent WebRTC standard as well as on PeerJS, an open-source project that provides primitives for data transfer on top of WebRTC.

The library currently includes the implementation of two peer sampling protocols, Cyclon [111] and the generic peer-sampling protocol from [98], as well as a clustering protocol [40]. All protocols implement a common GossipProtocol "interface"—since Javascript does not natively support interfaces, we adopt the interface pattern. A Coordinator makes it possible to stack these protocols on top of each other to implement applications.

MediEgo

APP: IDDN.FR.001.490030.001.S.A.2013.000.30000 in September 2015
Contribution: 5% (Specification)

MediEgo is a solution for content recommendation based on the users navigation history. The solution 1) collects the usages of the Web users and store them in a profile, 2) uses this profile to associate to each user her most similar users, 3) leverages this implicit network of close users in order to infer their preferences and recommend advertisements and recommendations. MediEgo achieves scalability using a sampling method, which provides very good results at a drastically reduced cost.

References

[66] Gianni Pasolini, Davide Dardari, and Michel Kieffer. "Exploiting the Agent's Memory in Asymptotic and Finite-Time Consensus Over Multi-Agent Networks". In: *IEEE transactions on Signal and Information Processing over Networks* 6 (2020), pp. 479–490. DOI: 10.1109/TSIPN.2020.3002613. URL: https://hal-centralesupelec.archives-ouvertes.fr/hal-03097195.

[67] Rachid Guerraoui, Petr Kuznetsov, Matteo Monti, Matej Pavlovič, and Dragos-Adrian Seredinschi. "Scalable Byzantine Reliable Broadcast". In: *DISC*. 2019.

[68] Rachid Guerraoui, Petr Kuznetsov, Matteo Monti, Matej Pavlovič, and Dragos-Adrian Seredinschi. "The Consensus Number of a Cryptocurrency". In: *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*. PODC '19. Toronto ON, Canada: Association for Computing Machinery, 2019, 307–316. ISBN: 9781450362177. DOI: 10.1145/3293611.3331589. URL: https://doi.org/10.1145/3293611.3331589.

[69] ITU. *G.1028: End-to-end quality of service for voice over 4G mobile networks*. 2019.

[70] Elli Androulaki, Artem Barger, Vita Bortnikov, et al. "Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains". In: *CoRR* abs/1801.10228 (2018).

[71] Simon Bouget, Yérom-David Bromberg, Adrien Luxey, and François Taïani. "Pleiades: Distributed Structural Invariants at Scale". In: *DSN 2018*. Luxembourg, Luxembourg: IEEE, June 2018, pp. 1–12.

[72] Rachid Guerraoui, Petr Kuznetsov, Matteo Monti, Matej Pavlovic, and Dragos-Adrian Seredinschi. "AT2: Asynchronous Trustworthy Transfers". In: *CoRR* abs/1812.10844 (2018). arXiv: 1812.10844. URL: http://arxiv.org/abs/1812.10844.

[73] Lucas Nunes Barbosa, Jonathan Gemmell, Miller Horvath, and Tales Heimfarth. "Distributed User-Based Collaborative Filtering on an Opportunistic Network". In: *2018 IEEE 32nd International Conference on Advanced Information Networking and Applications (AINA)*. 2018, pp. 266–273. DOI: 10.1109/AINA.2018.00049.

[74] J. Sousa, A. Bessani, and M. Vukolic. "A Byzantine Fault-Tolerant Ordering Service for the Hyperledger Fabric Blockchain Platform". In: *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 2018, pp. 51–58. DOI: 10.1109/DSN.2018.00018.

[75] Team Rocket. *Snowflake to avalanche: A novel metastable consensus protocol family for cryptocurrencies*. 2018.

[76] Daniel Valcarce, Alejandro Bellogín, Javier Parapar, and Pablo Castells. "On the Robustness and Discriminative Power of Information Retrieval Metrics for top-N Recommendation". In: *Proceedings of the 12th ACM Conference on Recommender Systems*. RecSys '18. Vancouver, British Columbia, Canada: ACM, 2018, pp. 260–268. ISBN: 978-1-4503-5901-6. DOI: 10.1145/3240323.3240347. URL: http://doi.acm.org/10.1145/3240323.3240347.

[77] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. *Practical Secure Aggregation for Privacy Preserving Machine Learning*. Cryptology ePrint Archive, Report 2017/281. https://ia.cr/2017/281. 2017.

[78] Armon Dadgar, James Phillips, and Jon Currey. "Lifeguard : SWIM-ing with Situational Awareness". In: *CoRR* abs/1707.00788 (2017).

[79] Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. "Algorand: Scaling Byzantine Agreements for Cryptocurrencies". In: *Proceedings of the 26th Symposium on Operating Systems Principles*. SOSP '17. Shanghai, China: ACM, 2017, pp. 51–68. ISBN: 978-1-4503-5085-3. DOI: 10.1145/3132747.3132757. URL: http://doi.acm.org/10.1145/3132747.3132757.

[80] Manami Kawasaki and Takashi Hasuike. "A recommendation system by collaborative filtering including information and characteristics on users and items". In: *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*. 2017, pp. 1–8. DOI: 10.1109/SSCI.2017.8280983.

[81] Rui Zhu, Bang Liu, Di Niu, et al. "Network Latency Estimation for Personal Devices: A Matrix Completion Approach". In: *IEEE/ACM Trans. Netw.* 25.2 (2017), pp. 724–737.

[82] Damien Imbs and Michel Raynal. "Trading off t-Resilience for Efficiency in Asynchronous Byzantine Reliable Broadcast". In: *Parallel Processing Letters* 26.4 (2016), 1650017:1–1650017:8.

[83] Brice Nédelec, Pascal Molli, and Achour Mostefaoui. "CRATE: Writing Stories Together with Our Browsers". In: *Proceedings of the 25th International Conference Companion on World Wide Web*. 2016, pp. 231–234. ISBN: 978-1-4503-4144-8.

[84] Paul Vanhaesebrouck, Aurélien Bellet, and Marc Tommasi. "Decentralized Collaborative Learning of Personalized Models over Networks". In: *CoRR* abs/1610.05202 (2016). arXiv: 1610.05202. URL: http://arxiv.org/abs/1610.05202.

[85] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. "Model Inversion Attacks That Exploit Confidence Information and Basic Countermeasures". In: *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security*. CCS '15. Denver, Colorado, USA: ACM, 2015, pp. 1322–1333. ISBN: 978-1-4503-3832-5. DOI: 10.1145/2810103.2813677. URL: http://doi.acm.org/10.1145/2810103.2813677.

[86] Miguel Matos, Hugues Mercier, Pascal Felber, Rui Oliveira, and José Pereira. "EpTO: An Epidemic Total Order Algorithm for Large-Scale Distributed Systems". In: *Proceedings of the 16th Annual Middleware Conference*. Middleware '15. Vancouver, BC, Canada: Association for Computing Machinery, 2015, 100–111. ISBN: 9781450336185. DOI: 10.1145/2814576.2814804. URL: https://doi.org/10.1145/2814576.2814804.

[87] Van Gegel. *TORFone: secure VoIP tool*. 2013.

[88] Pascal Felber, Anne-Marie Kermarrec, Lorenzo Leonini, et al. "Pulp: An adaptive gossip-based dissemination protocol for multi-source message streams". In: *Peer-to-Peer Networking and Applications* 5.1 (2012), pp. 74–91.

[89] J.A. Calandrino, A. Kilzer, A. Narayanan, E.W. Felten, and V. Shmatikov. ""You Might Also Like:" Privacy Risks of Collaborative Filtering". In: *SP*. IEEE, 2011. DOI: 10.1109/SP.2011.40.

[90] Armando Castañeda, Sergio Rajsbaum, and Michel Raynal. "The renaming problem in shared memory systems: An introduction". In: *Computer Science Review* 5.3 (2011), pp. 229–251.

[91] Bernhard Haeupler. "Analyzing network coding gossip made easy". In: *Proceedings of the forty-third annual ACM symposium on Theory of computing*. ACM, 2011, pp. 293–302.

[92] Avinash Lakshman and Prashant Malik. "Cassandra: A Decentralized Structured Storage System". In: *SIGOPS Oper. Syst. Rev.* 44.2 (Apr. 2010), pp. 35–40. ISSN: 0163-5980.

[93] Elica Campochiaro, Riccardo Casatta, Paolo Cremonesi, and Roberto Turrin. "Do Metrics Make Recommender Algorithms?" In: *2009 International Conference on Advanced Information Networking and Applications Workshops*. 2009, pp. 648–653. DOI: 10.1109/WAINA.2009.127.

[94] Mary-Luc Champel, Anne-Marie Kermarrec, and Nicolas Le Scouarnec. "FoG: Fighting the Achilles' Heel of Gossip Protocols with Fountain Codes". In: *SSS 2009, Lyon, France, November 3-6, 2009. Proceedings*. 2009, pp. 180–194.

[95] C. Fragouli, J. Widmer, and J. Y. Le Boudec. "Efficient Broadcasting Using Network Coding". In: *IEEE/ACM Transactions on Networking* 16.2 (Apr. 2008), pp. 450–463. ISSN: 1063-6692.

[96] ITU. *E.800 : Definitions of terms related to quality of service*. 2008. URL: https://www.itu.int/rec/T-REC-E.800-200809-I.

[97] Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, et al. "Dynamo: Amazon's Highly Available Key-value Store". In: *Proceedings of Twenty-first ACM SIGOPS*. SOSP '07. Stevenson, Washington, USA: ACM, 2007, pp. 205–220. ISBN: 978-1-59593-591-5.

[98] M. Jelasity, S. Voulgaris, R. Guerraoui, A.-M. Kermarrec, and M.v. Steen. "Gossip-based peer sampling". In: *TOCS* (2007).

[99] Márk Jelasity, Spyros Voulgaris, Rachid Guerraoui, Anne-Marie Kermarrec, and Maarten van Steen. "Gossip-based Peer Sampling". In: *TOCS* 25.3 (2007).

[100] Sujay Sanghavi, Bruce E. Hajek, and Laurent Massoulié. "Gossiping With Multiple Messages". In: *IEEE Trans. Information Theory* 53.12 (2007), pp. 4640–4654.

[101] N. Santoro and P. Widmayer. "Agreement in synchronous networks with ubiquitous faults". In: *Theoretical Computer Science* 384.2-3 (2007), pp. 232–249.

[102] Supratim Deb, Muriel Médard, and Clifford Choute. "Algebraic gossip: A network coding approach to optimal multiple rumor mongering". In: *IEEE/ACM Transactions on Networking (TON)* 14.SI (2006), pp. 2486–2507.

[103] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. "Calibrating Noise to Sensitivity in Private Data Analysis". In: *Proceedings of the Third Conference on Theory of Cryptography*. TCC'06. Berlin, Heidelberg: Springer-Verlag, 2006, pp. 265–284. ISBN: 3-540-32731-2 978-3-540-32731-8. DOI: 10.1007/11681878_14. URL: http://dx.doi.org/10.1007/11681878_14 (visited on 07/18/2014).

[104] Christina Fragouli, Jean-Yves Le Boudec, and Jörg Widmer. "Network Coding: An Instant Primer". In: *SIGCOMM Comput. Commun. Rev.* 36.1 (Jan. 2006), pp. 63–68. ISSN: 0146-4833.

[105] Sachin Katti, Hariharan Rahul, Wenjun Hu, et al. "XORs in the air: Practical wireless network coding". In: *ACM SIGCOMM computer communication review*. Vol. 36. ACM, 2006, pp. 243–254.

[106] Sebastian Michel, Matthias Bender, Nikos Ntarmos, Peter Triantafillou, Gerhard Weikum, and Christian Zimmer. "Discovering and Exploiting Keyword and Attribute-Value Co-Occurrences to Improve P2P Routing Indices". In: *Proceedings of the 15th ACM International Conference on Information and Knowledge Management*. CIKM '06. Arlington, Virginia, USA: Association for Computing Machinery, 2006, 172–181. ISBN: 1595934332. DOI: 10.1145/1183614.1183643. URL: https://doi.org/10.1145/1183614.1183643.

[107] Sebastian Michel, Matthias Bender, Nikos Ntarmos, Peter Triantafillou, Gerhard Weikum, and Christian Zimmer. "Discovering and Exploiting Keyword and Attribute-Value Co-occurrences to Improve P2P Routing Indices". In: *CIKM*. 2006.

[108] Daniel Stutzbach and Reza Rejaie. "Understanding Churn in Peer-to-Peer Networks". In: *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement*. IMC '06. Rio de Janeriro, Brazil: Association for Computing Machinery, 2006, 189–202. ISBN: 1595935614. DOI: 10.1145/1177080.1177105. URL: https://doi.org/10.1145/1177080.1177105.

[109] Márk Jelasity and Özalp Babaoglu. "T-Man: Gossip-Based Overlay Topology Management". In: *ESOA 2005, Utrecht, The Netherlands, July 25, 2005, Revised Selected Papers*. 2005, pp. 1–15.

[110] Márk Jelasity, Alberto Montresor, and Ozalp Babaoglu. "Gossip-based Aggregation in Large Dynamic Networks". In: *ACM Trans. Comput. Syst.* 23.3 (Aug. 2005), pp. 219–252. ISSN: 0734-2071.

[111] S. Voulgaris, D. Gavidia, and M. v. Steen. "CYCLON: inexpensive membership management for unstructured P2P overlays". In: *Journal of Network and Systems Management* (2005).

[112] Patrick Euster, Rachid Guerraoui, Anne-Marie Kermarrec, and Laurent Maussoulie. "From epidemics to distributed computing". In: *IEEE Computer* 37.5 (2004), pp. 60–67.

[113] Boris Koldehofe. "Simple gossiping with balls and bins." In: *Stud. Inform. Univ.* 3.1 (2004), pp. 43–60.

[114] Ranjita Bhagwan, Stefan Savage, and Geoffrey M. Voelker. "Understanding Availability". In: *Proceedings of the 2nd International Workshop on Peer-to-Peer Systems (IPTPS'03)*. 2003.

[115] Philip A. Chou, Yunnan Wu, and Kamal Jain. "Practical Network Coding". In: *Allerton Conference on Communication, Control, and Computing*. Oct. 2003.

[116] ITU. *ITU-T Recommendation G.114, "One way transmission time"*. 2003.

[117] Abhinandan Das, An Das, Indranil Gupta, and Ashish Motivala. "SWIM: Scalable Weakly-consistent Infection-style Process Group Membership Protocol". In: *In Proc. 2002 Intnl. Conf. DSN*. 2002, pp. 303–312.

[118] N. Santoro and P. Widmayer. "Time is not a healer". In: *Proc. 6th Annual Symposium on Theoretical Aspects of Computer Science (STACS'89)*. Springer, 1989, pp. 304–316.

[119] Alan J. Demers, Daniel H. Greene, Carl Hauser, et al. "Epidemic Algorithms for Replicated Database Maintenance". In: *Operating Systems Review* 22.1 (1988), pp. 8–32.

[120] Gabriel Bracha. "Asynchronous Byzantine Agreement Protocols". In: *Information & Computation* 75.2 (1987), pp. 130–143.

[121] David Chaum. "Security without Identification: Transaction Systems to Make Big Brother Obsolete". In: *Commun. ACM* 28.10 (Oct. 1985), 1030–1044. ISSN: 0001-0782. DOI: 10.1145/4372.4373.

[122] P. Flajolet and G. N. Martin. "Probabilistic counting algorithms for data base applications." In: *Journal of Computer and System Sciences* 31 (1985).

[123] Philippe Flajolet and G. Nigel Martin. "Probabilistic Counting Algorithms for Data Base Applications". In: *J. Comput. Syst. Sci.* 31.2 (1985), 182–209. ISSN: 0022-0000. DOI: 10.1016/0022-0000(85)90041-8. URL: https://doi.org/10.1016/0022-0000(85)90041-8.

[124] Leslie Lamport, Robert Shostak, and Marshall Pease. "The Byzantine Generals Problem". In: *ACM TOPLAS* (1982). DOI: 10.1145/357172.357176. URL: https://doi.org/10.1145/357172.357176.

[125] Leslie Lamport. "Time, clocks, and the ordering of events in a distributed system". In: *Communications of the ACM* 21.7 (1978), pp. 558–565.

[126] Team Rocket 0x 8a5d 2d32 e68b c500 36e4 d086 0446 17fe 4a0a 0296 b274 999b a568 ea92 da46 d533. "Snowflake to Avalanche: A Novel Metastable Consensus Protocol Family for Cryptocurrencies". https://ipfs.io/ipfs/QmUy4jh5mGNZvLkjies1RWM4YuvJh5o2FYopNPVYwrRVGV.

[127] Charaf Hassan and Frank HP Fitzek. "Network coding on the GPU". In: ().