



Offer #2022-05224

PhD Position F/M Byzantine fault tolerance and novel Sybil techniques for P2P storage

Contract type : Fixed-term contract

Level of qualifications required : Graduate degree or equivalent

Fonction : PhD Position

About the research centre or Inria department

The Inria Rennes - Bretagne Atlantique Centre is one of Inria's eight centres and has more than thirty research teams. The Inria Center is a major and recognized player in the field of digital sciences. It is at the heart of a rich R&D and innovation ecosystem: highly innovative PMEs, large industrial groups, competitiveness clusters, research and higher education players, laboratories of excellence, technological research institute, etc.

Context

This PhD thesis will be in the context of a collaboration between HIVE and Wide and Coast Inria teams. The Ph.D student will be located at Inria Center of the University of Rennes and will be visiting team Coast at Inria Nancy-Grand Est and the Hive offices in Cannes.

About Hive:

Hive intends to play the role of a next generation cloud provider in the context of Web 3.0. Hive aims to exploit the unused capacity of computers to offer the general public a greener and more sovereign alternative to the existing clouds where the true power lies in the hands of the users. It relies both on distributed peer-to-peer networks, on the encryption of end-to-end data and on blockchain technology.

About Inria Center of the University of Rennes:

The Inria Center of the University of Rennes is one of Inria's eight centers and has more than thirty research teams. The Inria Center is a major and recognized player in the field of digital sciences. It is at the heart of a rich R&D and innovation ecosystem: highly innovative PMEs, large industrial groups, competitiveness clusters, research and higher education players, laboratories of excellence, technological research institutes, etc.

About Inria Nancy - Grand Est:

The Inria Nancy - Grand Est center is one of Inria's eight centers and has twenty project teams, located in Nancy, Strasbourg and Saarbrücken. Its activities occupy over 400 people, scientists and research and innovation support staff, including 45 different nationalities. The Inria Center is a major and recognized player in the field of digital sciences. It is at the heart of a rich R&D and innovation ecosystem: highly innovative PMEs, large industrial groups, competitiveness clusters, research and higher education players, laboratories of excellence, technological research institutes, etc.

Assignment

A key challenge in a distributed storage system lies in withstanding the behavior of Byzantine nodes. These can result from malicious actors or simply from bugs or hardware problems, but in all cases they can render the system unusable in the absence of proper protocols [1]. Research in distributed algorithms has led to a number of protocols for managing Byzantine behavior in a classical, permissioned, setting [2, 3, 4, 5]. Yet operation in open permissionless settings remains challenging [6]. Currently, used methods like Proof-of-Work or Proof-of-Stake each have their own limitations. PoW suffers from enormous energy consumption, while PoS relies on the presence of a cryptocurrency or other similar asset.

Recent research has shown that, albeit promoted in the context of Bitcoin, the blockchain data structure is unnecessary to implement a cryptocurrency [7]. This has led to the appearance of a variety of proposals that implement cryptocurrencies with lighter primitives such as Byzantine Reliable Broadcast [8, 9]. This simplification is likely to extend to a variety of other applications that currently fall under the Blockchain hype, including distributed storage.

In this PhD thesis, we plan to explore novel techniques to handle Sybil [10, 11, 12] and Byzantine attacks in the context of distributed storage. We will explore distributed storage solutions that, while not requiring

the presence of a classical blockchain, can offer the same or even stronger guarantees in terms of fault-tolerance, data persistence, privacy, and security.

Main activities

Programme

The Ph.D. will consider this research question in two different contexts, close and open systems. These two context represents a gradation in the algorithmic and implementation difficulty of trustless distributed storage.

— In a first stage we will investigate large-scale trustless decentralized storage under a closed model. This model assumes that an attacker may comprise existing participants, but exclude attacks that require the injection of large numbers of new nodes, or the creation of novel identities. In this context the Ph.D. will consider the design, implementation and characterization of increasingly complex data-structures.

We will first consider append-only data structures, possibly with commutative operations. This type of storage object can be implemented using FIFO Byzantine Reliable Broadcast only [13]. Introduced in the eighties [14, 15], Byzantine reliable broadcast (BRB) is a fundamental abstraction of distributed computing [16, 17, 2, 18, 3, 4, 19, 20, 21]. BRB assumes that one particular process, the sender, broadcasts a message to the rest of the system, and that correct (a.k.a. honest) processes all deliver the value initially broadcast if the sender is correct, or that, if it is not, either all agree on some value, or that none delivers any value.

BRB algorithms are lighter and computationally weaker than consensus or total-order broadcast, but most existing algorithms and implementations exhibit a high message complexity and have not been design for large-scale systems. To alleviate these obstacles, we will explore how epidemic strategies can help overcome these performance bottlenecks, in particular in intermediate synchrony models, such as eventually synchronous networks [22].

In a second strand of research we will look at mutable data with possibly conflicting operations. In this context we will focus on practical means to detect and equivocation in replication protocols. Replication protocols are responsible for reconciling the state of copies to converge to a state that integrates the modifications of each one. The convergence of the copies conditions the liveliness and therefore the success of a collaboration. If two contributors have copies that are not unable to converge, their views of the content will be different. Some malicious peers may try to permanently diverge the views of other honest peers. To do this, they use the weaknesses of certain replication protocols: they transmit different modifications to different honest peers that are perceived as identical by the protocol. This is an equivocation.

The research question we want to address is how to ensure the convergence of copies of honest peers in the presence of malicious peers and to preserve the properties offered by the peer-to-peer collaboration?

We propose to use tamper-proof logs of operations. Each operation is signed by its author and references operations on which it depends. As two distinct operations have a different signature, pairs can identify equivocations.

— In a second stage we will move to an open trustless system, in which an attacker may create additional identities and launch Sybil attacks. Current approaches typically use some form of Proof-of-X (such as proof-of-work [23, 24] or proof-of-stake [25]) to address this problem. Proof-of-work (in the form of crypto puzzle) comes with a high energy cost [26], and limits the system's scalability. Proof-of-stake limits the openness of the system as only nodes with significant stake can participate in the consensus, leading to an oligopoly situation, and introduces a complex interdependency between the consensus algorithm and the cryptocurrency built upon it.

In this Ph.D. we plan to address these limitations by exploring how techniques related to storage, such as Proof-of-storage, could be used in the project's context, to address the challenges of open systems, while avoiding the (stark) limitations of existing solutions.

References:

[1] L. Lamport, R. Shostak, and M. Pease. "The Byzantine Generals Problem". In: ACM TOPLAS (1982). doi: 10.1145/357172.357176.

[2] G. Bracha. "Asynchronous Byzantine Agreement Protocols". In: Information and Computation 75.2 (1987), pp. 130–143. issn: 0890-5401. doi: 10.1016/0890-5401(87)90054-X.

[3] D. Imbs and M. Raynal. "Trading off t-Resilience for Efficiency in Asynchronous Byzantine Reliable Broadcast". In: Parallel Processing Letters 26.04 (2016). doi: 10.1142/ S0129626416500171.

[4] A. Mosteřfaoui, H. Moumen, and M. Raynal. "Signature-Free Asynchronous Binary Byzantine Consensus with $t < n/3$, $O(n^2)$ Messages, and $O(1)$ Expected Time". In: Journal of the ACM 62 (4 Aug. 2015). doi: 10.1145/2785953.

[5] M. O. Rabin. "Randomized Byzantine Generals". In: 2013 IEEE 54th Annual Symposium on Foundations of Computer Science. Los Alamitos, CA, USA: IEEE Computer Society, Nov. 1983, pp. 403–409. doi: 10.1109/SFCS.1983.48.

[6] R. Guerraoui, P. Kuznetsov, M. Monti, M. Pavlovic, and D.-A. Seredinschi. "Scalable Byzantine Reliable

Broadcast". In: 33rd International Symposium on Distributed Computing, DISC 2019. Oct. 2019. doi: 10.4230/LIPIcs.DISC.2019.22.

[7] R. Guerraoui, P. Kuznetsov, M. Monti, M. Pavlovic, and D.-A. Seredinschi. "The Consensus Number of a Cryptocurrency". In: Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing. ACM, 2019, pp. 307–316. isbn: 9781450362177. doi: 10.1145/3293611.3331589.

[8] A. Auvolat, D. Frey, M. Raynal, and F. Taïani. "Money Transfer Made Simple: a Specification, a Generic Algorithm and its Proof". In: Bulletin of EATCS 132 (2020). hal: hal-02861511v3.

[9] M. Baudet, G. Danezis, and A. Sonnino. "FastPay: High-Performance Byzantine Fault Tolerant Settlement". In: 2nd ACM Conference on Advances in Financial Technologies. ACM, 2020, pp. 163–177. doi: 10.1145/3419614.3423249.

[10] A. Bouchra Pilet, D. Frey, and F. Taïani. "Foiling Sybils with HAPS in Permissionless Systems: An Address-based Peer Sampling Service". In: IEEE Symposium on Computers and Communications. IEEE, 2020. doi: 10.1109/ISCC50000.2020.9219606.

[11] J. R. Douceur. "The Sybil Attack". In: International Workshop on Peer-to-Peer Systems. Springer Berlin Heidelberg, 2002, pp. 251–260. doi: 10.1007/3-540-45748-8_24.

[12] G. Urdaneta, G. Pierre, and M. V. Steen. "A Survey of DHT Security Techniques". In: ACM Computing Surveys 43.2 (Feb. 2011). issn: 0360-0300. doi: 10.1145/1883612.1883615.

[13] D. Frey, L. Guillou, M. Raynal, and F. Taïani. "Consensus-Free Ledgers When Operations of Distinct Processes are Commutative". In: Parallel Computing Technologies. Vol. 12942. Lecture Notes in Computer Science. Springer International Publishing, 2021, pp. 359–370. isbn: 978-3-030-86359-3. doi: 10.1007/978-3-030-86359-3_27.

[14] D. Dolev and H. R. Strong. "Authenticated algorithms for Byzantine agreement". In: SIAM Journal on Computing 12.4 (1983), pp. 656–666. doi: 10.1137/0212045.

[15] L. Lamport, R. Shostak, and M. Pease. "The Byzantine Generals Problem". In: ACM TOPLAS (1982). doi: 10.1145/357172.357176.

[16] H. Attiya and J. L. Welch. Distributed computing - fundamentals, simulations, and advanced topics (2. ed.) Wiley series on parallel and distributed computing. Wiley, 2004.

[17] A. Auvolat, D. Frey, M. Raynal, and F. Taïani. "Byzantine-tolerant causal broadcast". In: Theoretical Computer Science 885 (2021), pp. 55–68. doi: 10.1016/j.tcs.2021.06.021. hal:hal-03346710.

[18] C. Cachin, R. Guerraoui, and L. E. T. Rodrigues. Introduction to Reliable and Secure Distributed Programming (2. ed.) Springer, 2011.

[19] K. Nayak, L. Ren, E. Shi, N. H. Vaidya, and Z. Xiang. "Improved Extension Protocols for Byzantine Broadcast and Agreement". In: DISC. Vol. 179. LIPIcs. 2020, 28:1–28:17. doi: 10.4230/LIPIcs.DISC.2020.28.

[20] M. Raynal. Fault-Tolerant Message-Passing Distributed Systems - An Algorithmic Approach. Springer, 2018. doi: 10.1007/978-3-319-94141-7.

[21] J. Wan, H. Xiao, E. Shi, and S. Devadas. "Expected Constant Round Byzantine Broadcast Under Dishonest Majority". In: Proc. 18th Theory of Cryptography Conference. LNCS 12550. Springer, 2020, pp. 381–411. doi: 10.1007/978-3-030-64375-1_14.

[22] C. Dwork, N. A. Lynch, and L. J. Stockmeyer. "Consensus in the presence of partial synchrony". In: Journal of the ACM 35.2 (1988), pp. 288–323. doi: 10.1145/42282.42283.

[23] J. Garay, A. Kiayias, and N. Leonardos. "The Bitcoin Backbone Protocol: Analysis and Applications". In: Advances in Cryptology - EUROCRYPT 2015. Ed. by E. Oswald and M. Fischlin. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2015, pp. 281–310. isbn: 978-3-662-46803-6. doi: 10.1007/978-3-662-46803-6_10.

[24] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. 2009.

[25] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich. "Algorand: Scaling byzantine agreements for cryptocurrencies". In: Proceedings of the 26th Symposium on Operating Systems Principles. ACM, 2017, pp. 51–68. doi: 10.1145/3132747.3132757.

[26] Bitcoin Energy Consumption Index - Digiconomist. <https://digiconomist.net/bitcoin-energy-consumption>. Accessed: 2020-03-05.

Skills

- Engineering and/or Master 2 degree in Computer science / Applied mathematics with an experience in computer networks.

- Theoretical expertise: distributed systems, P2P networks, security
- Good collaborative and networking skills, excellent written and oral communication in English
- Good programming skills
- Strong analytical skills

Benefits package

- Subsidized meals
- Partial reimbursement of public transport costs
- Possibility of teleworking (90 days per year) and flexible organization of working hours
- partial payment of insurance costs

Remuneration

Monthly gross salary amounting to 2051 euros for the first and second years and 2158 euros for the third year

General Information

- **Theme/Domain** : Distributed programming and Software engineering System & Networks (BAP E)
- **Town/city** : Rennes
- **Inria Center** : [Centre Inria de l'Université de Rennes](#)
- **Starting date** : 2022-12-01
- **Duration of contract** : 3 years
- **Deadline to apply** : 2022-09-30

Contacts

- **Inria Team** : [WIDE](#)
- **PhD Supervisor** :
Taiani Francois / Francois.Taiani@irisa.fr

About Inria

Inria is the French national research institute dedicated to digital science and technology. It employs 2,600 people. Its 200 agile project teams, generally run jointly with academic partners, include more than 3,500 scientists and engineers working to meet the challenges of digital technology, often at the interface with other disciplines. The Institute also employs numerous talents in over forty different professions. 900 research support staff contribute to the preparation and development of scientific and entrepreneurial projects that have a worldwide impact.

Warning : you must enter your e-mail address in order to save your application to Inria. Applications must be submitted online on the Inria website. Processing of applications sent from other channels is not guaranteed.

Instruction to apply

Please submit online : your resume, cover letter and letters of recommendation eventually

Defence Security :

This position is likely to be situated in a restricted area (ZRR), as defined in Decree No. 2011-1425 relating to the protection of national scientific and technical potential (PPST). Authorisation to enter an area is granted by the director of the unit, following a favourable Ministerial decision, as defined in the decree of 3 July 2012 relating to the PPST. An unfavourable Ministerial decision in respect of a position situated in a ZRR would result in the cancellation of the appointment.

Recruitment Policy :

As part of its diversity policy, all Inria positions are accessible to people with disabilities.