PRD: Byzantine Fault-tolerance in dynamic networks

Advisor: Sébastien Tixeuil, Professor at Sorbonne University, LIP6 NPA team
sebastien.tixeuil@lip6.fr

April 2, 2021

1 Context

Following the setting of the seminal paper of Lamport et al. [13], many subsequent papers focusing on Byzantine tolerance [2, 14, 15, 17, 18] study agreement and reliable communication primitives using cryptography-free protocols in networks that are both *static* and *fully connected*. An important line of research assumes the existence of $2k + 1$ node-disjoint paths from source to destination, in order to provide reliable communication in the presence of up to k Byzantine failure [10, 9, 19]. However, these results rely on Menger's theorem [3], which can be informally expressed as follows: we have x disjoint paths between two nodes if and only if x nodes must be removed to disconnect these two nodes. This theorem only applies to *static* networks. None of the aforementioned papers considers genuinely dynamic networks, *i.e.*, where the topology evolves while the protocol executes.

In a seminal paper [16], Maurer *et al.* considered the following problem: two nodes want to reliably communicate in a *dynamic multi-hop network* where a subset of the nodes are Byzantine. As Menger's theorem does not extend to dynamic networks [11], Maurer *et al.* [16] proved a new necessary and sufficient condition (in other words, the weakest possible condition) for enabling reliable communication in a dynamic multi-hop network where a subset of the nodes are Byzantine. Their proof is constructive, as they provide a Byzantine-resilient algorithm for reliable communication that is optimal with respect to their impossibility result. However, their algorithm *requires an exponential computation at every node upon receipt of every message*, hindering its practical relevance.

Scientific lock: Overall, to date, for dynamic networks, no tractable (that is, polynomial) solution exists for reliable communication with Byzantine attacks.

2 Methodology and Organization

The first challenge to solve is a tractable Byzantine broadcast algorithm. Although the condition provided by Maurer et al. [16] is the best possible from a theoretical point of view, it lacks practicality. That is, verifying their condition requires to run at each node upon each message reception an exponential time algorithm to compute the minimal node-cut of the received paths attached to received messages. As our preliminary results suggest that computing the minimal dynamic cut in dynamic networks using received paths attached to messages is at least NP-complete, one option is to investigate approximation algorithms. However, as we do not want to compromise safety, the approximation should always be lower than the actual minimal dynamic cut. One trivial such approximation is *the number of node disjoint-paths* among the received paths, but there exist dynamic networks where no two dynamic paths are node-disjoint [11], so the computed dynamic minimal cut would be 1, *resulting in no genuine message being delivered even if at most one Byzantine node is in the network*, that is, a loss of liveness.

Designing a polynomial approximation for the dynamic minimal cut problem that compromises neither safety nor liveness of the resulting broadcast algorithm is an obvious first goal.

The second goal is to investigate the orthogonal approach of using a purely local Byzantine-resilient broadcast algorithm, that is, an algorithm that does not store traversed paths with each message. In static networks, this approach was promoted as the Certified Propagation Algorithm (CPA) [12, 20, 21]: if sufficiently many neighbors broadcast the same message, then it is safe to rebroadcast it. For example, if at most k Byzantine nodes can be neighbors of a particular node, then it is safe to rebroadcast it once received $k + 1$ copies of the same message from different neighbors. Of course, the necessary and sufficient conditions induced by local algorithms are different from the global ones (we go from a global requirement on the number of Byzantine attackers to a local requirement about the neighborhood of each node), but not necessarily less useful as they are intrinsically polynomial locally. A first step in this direction was initiated by Bonomi et al. [5] for a limited set of (unpractical) dynamic networks. Another aproach by Dobrev et al. [8] consider the broadcast problem in the context of dynamic edge faults, that are a strict subset of Byzantine failures. So, our second goal in this task will be to fully characterize the CPA approach in dynamic networks. This is likely to promote a new set of lower and upper bounds, expressed with local metrics, for polynomial broadcast algorithms.

We expect the outputs of the first two goals (a global approximation approach, and a local exact approach) to be incomparable from a theoretical point of view, as their constraints (on the number and placement of Byzantine) are not comparable. However, it makes sense to assess their relevance from a practical point of view, using data from real dynamic networks. We plan to devise a set of implementation-level optimizations for algorithms developped in goals 1 and 2, following the line of work initiated by our recent work [6, 4], but also general optimizations for long lived executions [7]. Furthermore, we will first use the traces of dynamic networks available at the CRAWDAD [1] archive: *e.g.* mobility and connectivity traces extracted from GPS traces from the regional Fire Department of Asturias, social interactions and propinquity based on wireless and bluetooth, mobility traces of taxis in Rome, etc. Then, we will benchmark the algorithms obtained from Task 1.1 using the sames set of traces for each algorithm, and measure the following metrics: *(i)* number of safety violations (that is, the number of fake messages delivered), *(ii)* number of liveness violations (that is, the number of nodes that do not deliver a genuine message), *(iii)* time to complete a single broadcast, *(iv)* time to complete a series of m broadcasts.

Obviously, a difficult setting in the evaluation is the placement of Byzantine attackers, as it may have a dramatic impact on the performance of the protocols. We expect to start with a random placement and run sufficiently many evaluations to have high confidence intervals. However, we also assume that the first results will mandate additional iterations of experiments are we gain insight about the worst case Byzantine placement.

3 Practicalities

The PhD will be located in LIP6 laboratory in Paris, France. Part of the research will be done in cooperation with colleagues in Japan (Xavier Défago at TokyoTech, Toshimitsu Masuzawa at Osaka University).

References

[1] CRAWDAD archive. `http://www.crawdad.org`.

[2] H. Attiya and J. Welch. *Distributed Computing: Fundamentals, Simulations, and Advanced Topics*. McGraw-Hill Publishing Company, New York, May 1998. 6.

[3] T. Böhme, F. Göring, and J. Harant. Menger's theorem. *Journal of Graph Theory*, 37(1):35–36, 2001.

[4] Silvia Bonomi, Jérémie Découchant, Giovanni Farina, Vincent Rahli, and Sébastien Tixeuil. Practical byzantine reliable broadcast on partially-connected networks. In *Prooceedings of IEEE ICDCS 2021, Washington, USA*, 2021.

[5] Silvia Bonomi, Giovanni Farina, and Sébastien Tixeuil. Reliable broadcast in dynamic networks with locally bounded byzantine failures. In Taisuke Izumi and Petr Kuznetsov, editors, *Stabilization, Safety, and Security of Distributed Systems - 20th International Symposium, SSS 2018, Tokyo, Japan, November 4-7, 2018, Proceedings*, volume 11201 of *Lecture Notes in Computer Science*, pages 170–185. Springer, 2018.

[6] Silvia Bonomi, Giovanni Farina, and Sébastien Tixeuil. Multi-hop byzantine reliable broadcast with honest dealer made practical. *J. Braz. Comp. Soc.*, 25(1):9:1–9:23, 2019.

[7] Silvia Bonomi, Giovanni Farina, and Sébastien Tixeuil. Boosting the efficiency of byzantine-tolerant reliable communication. In *Proceedings of the 22nd International Symposium on Stabiliation, Safety, and Security of Distributed Systems (SSS 2020)*, Lecture Notes in Computer Science (LNCS), Austin, Texas, November 2020. Springer Berlin / Heidelberg.

[8] Stefan Dobrev, Rastislav Kralovic, Richard Královic, and Nicola Santoro. On fractional dynamic faults with thresholds. *Theor. Comput. Sci.*, 399(1-2):101–117, 2008.

[9] D. Dolev. The Byzantine generals strike again. *Journal of Algorithms*, 3(1):14–30, 1982.

[10] Danny Dolev. Unanimity in an unknown and unreliable environment. In *22nd Annual Symposium on Foundations of Computer Science, Nashville, Tennessee, USA, 28-30 October 1981*, pages 159–168. IEEE Computer Society, 1981.

[11] D. Kempe, J. Kleinberg, and A. Kumar. Connectivity and inference problems for temporal networks. *Journal of Computer and System Sciences*, 64(4):820–842, 2002.

[12] C.-Y. Koo. Broadcast in radio networks tolerating Byzantine adversarial behavior. In Soma Chaudhuri and S. Kutten, editors, *PODC*, pages 275–282. ACM, 2004.

[13] L. Lamport, R. E. Shostak, and M. C. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.

[14] D. Malkhi, Y. Mansour, and M.K. Reiter. Diffusion without false rumors: on propagating updates in a Byzantine environment. *Theoretical Computer Science*, 299(1–3):289–306, April 2003.

[15] D. Malkhi, M. Reiter, O. Rodeh, and Y. Sella. Efficient update diffusion in Byzantine environments. In *The 20th IEEE Symposium on Reliable Distributed Systems (SRDS '01)*, pages 90–98, Washington - Brussels - Tokyo, October 2001. IEEE.

[16] A. Maurer, S. Tixeuil, and X. Défago. Communicating reliably in multihop dynamic networks despite byzantine failures. In *34th IEEE Symposium on Reliable Distributed Systems, SRDS 2015, Montreal, QC, Canada, September 28 - October 1, 2015*, pages 238–245. IEEE Computer Society, 2015.

[17] Y. Minsky and F.B. Schneider. Tolerating malicious gossip. *Distributed Computing*, 16(1):49–68, 2003.

[18] Achour Mostéfaoui and Michel Raynal. Signature-free asynchronous byzantine systems: from multivalued to binary consensus with t< n/3, o(n^2) messages, and constant time. *Acta Informatica*, 54(5):501–520, 2017.

[19] M. Nesterenko and S. Tixeuil. Discovering network topology in the presence of Byzantine faults. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 20(12):1777–1789, December 2009.

[20] A. Pelc and D. Peleg. Broadcasting with locally bounded Byzantine faults. *Inf. Process. Lett.*, 93(3):109–115, 2005.

[21] L. Tseng, N.H. Vaidya, and V. Bhandari. Broadcast using certified propagation algorithm in presence of byzantine faults. *Inf. Process. Lett.*, 115(4):512–514, 2015.