# Modeling and optimizing BFT in the presence of network failures and anomolies

Martin Nischwitz [1], Marko Esche [1], and Florian Tschorsch [2]

**Abstract:** The most challenging fault within a distributed system is the so called Byzantine Fault, explained and named by Leslie Lamport et al. [LSP82]. The first practical implementation of Byzantine Fault Tolerance (BFT) was published in 1999 and fittingly named Practical Byzantine Fault Tolerance [CL99]. This was followed by a stream of BFT implementations, each improving upon the consensus process for specific use cases or emphasizing certain properties of it. One important aspect is the deployment in networks with unreliable communication. Examples of protocols that were developed for such use-cases are CheapBFT [Ka12] or the works of Fathollahnejad et al. [FBK17].

Our research presented here will highlight two aspects of BFT systems, namely modeling and optimizing existing protocols for certain setups with unreliable communication without the need for costly test implementations. Communication properties such as reliability or delay can have a profound impact on the performance of the distributed system. The scope of different networks in which distributed systems are deployed, ranges from a static global scale, e. g., the service of a distributed ledger implemented across continents, down to wireless dynamic networks, e. g., car-to-car communication. Furthermore, the devices in dynamic networks can often only supply a limited amount of power to the system, which makes energy efficiency an important aspect of such algorithms. It is therefore advised to assess the impact of communication or link failures on the efficiency and performance of BFT algorithms and design and adapt the algorithms accordingly.

Evaluating the effect of link failures on different implementations of fault tolerant systems may prove harder than one might initially assume. First of, the transport layer protocol will drastically change the impact of link failures. Since fault tolerant protocols are designed to withstand faulty processes, and a link failure can be interpreted as a faulty process for both communication partners, it is not absolutely necessary to provide reliable communication. For example, TCP as well as UDP can be used for fault tolerant systems but the impact of link failures will differ for both protocols. Secondly, implementing a fault tolerant system requires a lot of work. Many publications of algorithms offer only proof-of-concept implementations that were tested for very specific use cases. Deploying and testing such an algorithm in a new environment to observe and measure the effect of failures in order to determine its suitability is time consuming.

To this end, the first part of our presentation is a modeling process that predicts the impact of link failures within dynamic networks, based on the communication patterns of fault tolerant algorithms [NET20]. In particular, we will investigate a procedure how to create and utilize an atomic building block to model the impact of link failures for any fault tolerant protocol that is based on quorum collection.

In addition to the general assessment of link failures in dynamic networks, another important aspect is efficiency and optimization. Since the communication links in dynamic networks are typically not

[1] Physikalisch-Technische Bundesanstalt, Abbestr. 2-12, 10587 Berlin, Deutschland, martin.nischwitz@ptb.de, https://orcid.org/0009-0002-5488-5820; marko.esche@ptb.de, https://orcid.org/0009-0001-7110-5665

[2] Technische Universität Dresden, Privacy and Security, 01062 Berlin, Deutschland, florian.tschorsch@tu-dresden.de, https://orcid.org/0000-0001-6716-7225

stable, latency and throughput of individual links will change over time. Coupled with the fact that most fault tolerant systems do not equally distribute the communication effort on all processes, and elect a temporary leader to coordinate the system instead, even changing the properties of a single link has the potential to affect the performance of the whole system.

Therefore, the second part of our research considers the effects of dynamic network changes in a fault tolerant system and explores possibilities how to detect and react to those changes in order to optimize the performance of the system [NET24]. Emphasis is put on algorithms with limited information on the network state, such as HotStuff [Yi18].

In conclusion, we present a model for the behavior of quorum-based fault tolerant systems in dynamic networks and offer means for assessment and optimization of those systems in presence of delays, link and crash failures.

**Keywords:** distributed systems, fault tolerance, dynamic networks

## Bibliography

[CL99] Castro, Miguel; Liskov, Barbara: Practical Byzantine Fault Tolerance. In: Proceedings of the Third Symposium on Operating Systems Design and Implementation. OSDI '99, USENIX Association, Berkeley, CA, USA, pp. 173–186, 1999.

[FBK17] Fathollahnejad, N.; Barbosa, R.; Karlsson, J.: A Probabilistic Analysis of a Leader Election Protocol for Virtual Traffic Lights. In: 2017 IEEE 22$^{nd}$ Pacific Rim International Symposium on Dependable Computing (PRDC). pp. 311–320, January 2017.

[Ka12] Kapitza, Rüdiger; Behl, Johannes; Cachin, Christian; Distler, Tobias; Kuhnle, Simon; Mohammadi, Seyed Vahid; Schröder-Preikschat, Wolfgang; Stengel, Klaus: CheapBFT: Resource-efficient Byzantine Fault Tolerance. In: Proceedings of the 7$^{th}$ ACM European Conference on Computer Systems. EuroSys '12, ACM, New York, NY, USA, pp. 295–308, 2012.

[LSP82] Lamport, Leslie; Shostak, Robert; Pease, Marshall: The Byzantine Generals Problem. ACM Trans. Program. Lang. Syst., 4(3):382–401, July 1982.

[NET20] Nischwitz, Martin; Esche, Marko; Tschorsch, Florian: Bernoulli Meets PBFT: Modeling BFT Protocols in the Presence of Dynamic Failures, 2020.

[NET24] Nischwitz, Martin; Esche, Marko; Tschorsch, Florian: StarReact: Detecting Important Network Changes in BFT Protocols with Star-Based Communication. In: 2024 IEEE 49th Conference on Local Computer Networks (LCN). pp. 1–7, 2024.

[Yi18] Yin, Maofan; Malkhi, Dahlia; Reiter, Michael K.; Gueta, Guy Golan; Abraham, Ittai: HotStuff: BFT Consensus in the Lens of Blockchain, 2018.