


Data Plane based Optimization for Byzantine Fault Tolerant Consensus

Motivation

State-machine-replication (SMR) is used to build fault tolerant systems such as in airplanes, industrial control systems, and data centers. SMR uses multiple redundant machines (replicas), which agree on a common value even in case of faulty behavior of a minority. Byzantine Fault Tolerance (BFT) allows for a limited number of replicas to behave arbitrarily (drop/send wrong values/messages).



Due to the overhead of BFT-SMR system operation, a plethora of optimization approaches emerged. This includes moving parts of consensus to dedicated networking hardware. P4 [1] is a domain-specific language for defining packet processing on data plane devices.

In this thesis, we work on the seminal PBFT [2] consensus protocol. The main objective of this thesis is to analyze the impact of data plane acceleration on PBFT, implement PBFT using P4 functionality, and measure the resulting performance impact.

Your Tasks

- Familiarize yourself with the PBFT protocol and our codebase
- Familiarize yourself with the P4 language and ecosystem
- Extend the PBFT codebase to utilize P4 acceleration
- Measure and evaluate the results

References

- [1] Open Networking Foundation. "P4 Open Source Programming Language." <https://p4.org/>. 2023.
- [2] Castro, Miguel, and Barbara Liskov. "Practical Byzantine Fault Tolerance." *OsDI*. Vol. 99. No. 1999. 1999.

Prerequisites

- Experience with Linux-based operating systems and networking
- Experience with C/C++/Python programming
- Bonus: Experience with P4

Contact

Richard von Seck seck@net.in.tum.de
Manuel Simon simonm@net.in.tum.de

