University of Bern
Institute of Computer Science
Cryptology and Data Security

Metastable Consensus Protocols - Implementation & Evaluation

BSc, MSc Thesis

Project Description

Avalanche has become one of the leading protocols for implementing distributed ledgers (like a block chain) and the market capitalization of its associated cryptocurrency (AVAX) ranks in the top 20. While the Avalanche protocol combines several techniques on different layers of abstraction to ensure security, low latency and high throughput, the focus of this thesis is on a set of the subroutines the Avalanche protocol uses. These address a problem of any (truly) distributed data structure: forming consensus about the current state of the data structure among the participating (honest) parties (aka nodes), even if some are Byzantine (i.e., behave in unintended or even malicious ways).

In particular, this thesis will focus on the (binary) consensus problem, which is defined as follows. Given n nodes of a network, each of which initially has an input bit (0 or 1). The consensus problem is solved when all honest nodes output the same bit value, with the additional constraint that the output bit must be equal to the input bit of *at least one* node (honest or Byzantine).

The Avalanche protocol uses randomized mechanisms to solve the consensus problem. The most basic algorithm works as follows. Each node repeatedly makes queries. In each query a certain number of random nodes is asked for their current value and if a robust majority of one value (0 or 1) is observed in the responses, the querying node adopts the value as its own. The idea of this protocol is that a slight overall majority in either direction (0 or 1) will quickly converge to a full consensus in case there are no Byzantine nodes. If that is the case, then basic protocol must be amended by termination conditions, where a node outputs its current value when it is confident that this value will eventually be adopted by the whole network.

This thesis will focus on the implementation and evaluation of the basic Avalanche consensus protocols. The implementation of a testbed should focus on efficiency (allowing simulation of large networks) and modularity (for fluently switching between different consensus algorithms, modeling assumptions and adversarial behavior etc.) and can be done in the students language of choice (with the level of support depending on the supervisors own proficiency in that language). In the evaluation part, a sufficiently wide range of scenarios and parameters (e.g., number of queried nodes, number of adversarial nodes) should tested with respect to certain performance parameters (likelihood of success, number of queries until consensus). The results should be visualized and discussed in a written part.

Recommended Reading

- Sections 1,2,3 of Yin et al., Scalable and Probabilistic Leaderless BFT Consensus through Metastability
- More detailed pseudocode is provided in Amores-Sesar at al., When is Spring coming? A Security Analysis of Avalanche Consensus, OPODIS ’22

Requirements

- Moderate programming skills
- Theory: 30% Systems: 70%

Contact

- **Philipp Schneider**, philipp.schneider2@unibe.ch