Breaking the Safety and Liveness Guarantees of the Snowflake Consensus Construction with Byzantine Response Delays

Sarah Jamie Lewis

April 18, 2020

Introduction

> *"When [snowflake] is correctly parameterized for a given threshold of Byzantine nodes and a desired ϵ-guarantee, it can ensure both safety (P1) and liveness (P2, P3)"* [2]

We observe that a byzantine adversary can violate the safety guarantees of the Snowflake consensus construction described in [2] and [3] by delaying responses (within permitted bounds) to sampling requests from correct nodes at rates intended to exhaust the β confidence limits.

Through probabilistic modelling we formally verify an adversarial strategy that forces correct nodes to choose between safety and liveness even when $f < \mathcal{O}(\sqrt{n})$.

Intuition

The Snow family of leaderless Byzantine fault tolerance consensus protocols attempt to achieve metastability via network subsampling. To achieve these all the protocols sample from a random subset of network nodes and make protocol decisions based on a function of the returned samples.

If a byzantine node can delay responses to all by 1 correct peer, and the delay is long enough such that it allows a the correct node to reach a state where cnt $> \beta$ (i.e. the node settles on a color) while other correct nodes are delayed in a sampling process then a byzantine peer gains a significant advantage by effectively freezing the network and attempting to bias the decisions 1 correct node at a time.

In reality, nodes will timeout of sampling decisions and the network will be unlikely to reach a completely frozen state - but the bias introduced by adversarial delays to responses from byzantine nodes is enough to bias the protocol enough to force a choice between safety and liveness.

Formal Verification

We use PRISM [1] to model [1] and verify instantiations of the Snowflake protocol (with 4 peers, 1 byzantine) normalizing the notion that malicious nodes *may* delay their response to a sampling request indefinitely and correct nodes *may* abort a sampling process after exceeding a timeout.

We assume, as per [2] that correct nodes select sampling peers uniformly from the set of available peers - without replacement and after aborting a sampling process a correct node will start a new sampling process. [2]

Atomic Sampling and Timeouts

Our model permits correct nodes a slight additional advantage by allowing them to obtain the colors of other correct nodes in the same instance as sampling.

We model progression towards a maximum timeout t_m as $t_i < t_m$ where i is incremented when a correct node is scheduled but is unable to complete a sampling.

Results and Discussion

In the catastrophic case of $n = 4$, $k = 3$ and $\alpha = 2$ (parameters that would otherwise be optimal in terms of both safety and liveness when considering adversaries who cannot delay responses) a single byzantine node can delay consensus in a bivalent network indefinitely (by preventing 2 nodes from sampling at all, and ensuring the remaining correct node remains undecided indefinitely.

[1] `https://git.openprivacy.ca/sarah/formal-verification/src/branch/master/snowflake-4-adversary.prism`

[2] (alternatively a correct node may choose to select a replacement node to replace the node that has timed out - we do not explore this strategy, but note that it may be possible for an adversary to exploit this behaviour to bias the sampling process)

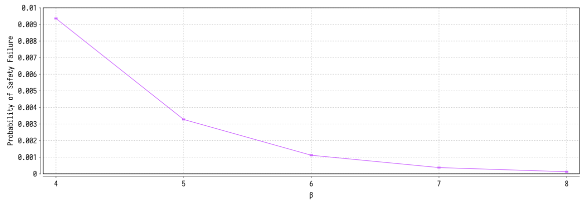

Figure 1: Verification of our model for all pairings of $t_m = 1$ and $\beta = \{4..8\}$. We note that even for $\beta = 8$ the probability of safety failure $P(F) > 0.00001$ for all timeout values. It is possible to continue to increase β to reduce the probability $\leq \epsilon$) but to do so trades off liveness.

Reducing $k = 2$ ensures that liveness-guarantee is maintained, however a correct client is forced to a set $\beta > 15$ dramatically increasing the time to consensus well beyond the time required for the ideal $\mathcal{O}(\sqrt{n})$ rounds in order to achieve a probability of failure, $P(F) \geq 1 \times 10^{-73}$ orders of magnitude above a suggested safety parameter i.e. 1×10^{-9}.

As evident from the verification, there is no parameterization of a 4-node Snowflake network (1 byzantine node) that can achieve P1 safety and P3 liveness.

Further to the above analysis, we note that any correct client which permits a timeout $t_m > 1$ will suffer greater impacts to both safety and liveness - obviously increasing the time a correct client will wait for a response from any node will increase the overall time of the protocol but, additionally, increasing tolerance to delays also increases the probability of failure in the presence of a byzantine adversary by permitting them bias nodes in greater isolation.

Applicability to Larger Networks, Snowball and Avalanche

The strategies presented above can be applied to larger instantiations of Snowflake $n > 4$ and, in particular, the result presented above should generalize in cases where $f \geq n/4$ permitting that all byzantine nodes all follow the same strategy, in such cases the probability that and correct node is delayed by a byzantine adversary is the same, and as such the proportion of bias in the network should also be the same.

Snowball introduces confidence indicators on top of Snowflake, but its safety and security properties are fundamentally grounded in those of snowflake. As such we believe that while Snowball based constructions (like Avalanche) will bound the impact of byzantine strategies, more work is needed to ensure that such strategies satisfy the desired ϵ-guarantee for both Safety and Liveness.

References

[1] M. Kwiatkowska, G. Norman, and D. Parker. PRISM 4.0: Verification of probabilistic real-time systems. In G. Gopalakrishnan and S. Qadeer, editors, *Proc. 23rd International Conference on Computer Aided Verification (CAV'11)*, volume 6806 of *LNCS*, pages 585–591. Springer, 2011.

[2] Team Rocket. Snowflake to avalanche: A novel metastable consensus protocol family for cryptocurrencies, 2018.

[3] Team Rocket, Maofan Yin, Kevin Sekniqi, Robbert van Renesse, and Emin Gün Sirer. Scalable and probabilistic leaderless bft consensus through metastability. *arXiv preprint arXiv:1906.08936*, 2019.

[3] Verified by PRISM