When $3f + 1$ is not Enough: Tradeoffs for Decentralized Asynchronous Byzantine Consensus[⋆]

Alysson Neves Bessani, Miguel Correia, Henrique Moniz,
Nuno Ferreira Neves, and Paulo Verissimo

LaSIGE, Faculdade de Ciências da Universidade de Lisboa, Portugal

1 Context and Motivation

Recently, we challenged the belief that randomized Byzantine agreement protocols are inefficient, by designing, implementing and assessing the performance of a stack of protocols of that type [3]. That assessment lead us to a set of properties desirable for Byzantine asynchronous binary consensus protocols: (1) Strong validity – if all correct processes propose the same value v, the decision is v (values proposed by Byzantine processes are often useless); (2) Asynchrony – no time assumptions are made (systems are often prone to arbitrary delays); (3) Decentralization – there is no leader (leader elections have a great impact on performance); (4) Optimal resilience – $n \geq 3f+1$ processes to tolerate f Byzantine (extra processes are costly); (5) Optimal message complexity – $O(n^2)$ (high impact on throughput); (6) Signature freedom (high impact of signatures based on public-key cryptography on the performance); (7) Early decision – in "nice" runs the protocol should decide in a few communication steps (good latency in the "normal" case).

The main characteristic of the decentralized protocols we are interested in this paper is that they can not require any reliable certificate from a process p_i, obtained in phase k or less, in order to justify a message sent in phase $k+1$. This is the case because, in our system model, this kind of certificate can only be build with digital signatures (violating signature freedom) or reliable multicast (that can not be executed by all processes maintaining a message complexity $O(n^2)$). Moreover, given the validity condition we stipulated (1), we require that all correct processes communicate their proposals to each other (a process can not trust another process to correctly communicate its value to a third process, since there are no signatures).

2 The Tradeoff

Is it possible to design such a Byzantine asynchronous binary consensus protocol? The main results in the paper are given by the following theorems:

[⋆] Work supported by project IST-4-027513-STP (CRUTIAL) and LaSIGE.

Theorem 1 (Impossibility result). *There is no decentralized algorithm that solves asynchronous binary Byzantine consensus with $n \leq 5f$, $O(n^2)$ message complexity and without signatures.*

Given this impossibility and several other results and protocols already described in the literature, it is possible to define in which conditions a binary decentralized Byzantine consensus protocol can exist:

Theorem 2 (Tradeoff). *Decentralized algorithms that solve asynchronous binary Byzantine consensus can be build with and only with:*

1. ***More Processes:*** *$n \geq 5f + 1$, $O(n^2)$ message complexity and signature freedom;*
2. ***More Messages:*** *$n \geq 3f + 1$, $O(o)$ message complexity ($n^2 < o \leq n^2 f$) and signature freedom;*
3. ***Signatures:*** *$n \geq 3f + 1$, $O(n^2)$ message complexity and using signatures.*

Notice that the bound established by Theorem 2 regarding *more messages* is not tight: we do not know if it is possible to solve Byzantine consensus without signatures and optimal resilience with message complexity lower than $O(n^2 f)$, but greater than $O(n^2)$.

3 Discussion

An interesting consequence of the theorems above is that decentralized protocols are inherently more costly in terms of the three properties considered (resilience, message complexity, signature) than leader-based Byzantine consensus protocols. For instance, the CL-BFT state machine replication protocol, that can be trivially adapted to solve consensus, is not subject to the tradeoff in Theorem 2 [2]. However, this protocol does not ensure the strong validity condition that we are interested in and requires synchrony to be able to terminate.

Theorem 1 implies that a consensus protocol with all the desired properties listed above can not be designed. However, we developed an improved protocol based on Bracha's Byzantine consensus [1], an algorithm that we believe is close enough to the desired characteristics that we envisage. This protocol improves the original Bracha's protocol in two main points: *(1.)* its message complexity is $O(n^2 f)$ instead of $O(n^3)$; and *(2.)* it can terminate in one communication step if some optimistic conditions hold (no faults and unanimity).

References

1. G. Bracha. An asynchronous $\lfloor (n-1)/3 \rfloor$-resilient consensus protocol. In *Proceedings of the 3rd ACM Symposium on Principles of Distributed Computing - PODC'84*, pages 154–162, Aug. 1984.
2. M. Castro and B. Liskov. Practical Byzantine fault-tolerance and proactive recovery. *ACM Transactions Computer Systems*, 20(4):398–461, Nov. 2002.
3. H. Moniz, N. F. Neves, M. Correia, and P. Veríssimo. Randomized intrusion-tolerant asynchronous services. In *Proceedings of the International Conference on Dependable Systems and Networks – DSN'06*, pages 568–577, June 2006.