

# From Symmetric to Asymmetric Asynchronous Byzantine Consensus

Christian Cachin<sup>1</sup>  
University of Bern  
cachin@inf.unibe.ch

Luca Zanolini<sup>1</sup>  
University of Bern  
luca.zanolini@inf.unibe.ch

## Abstract

Consensus is arguably one of the most important notions in distributed computing. Among asynchronous, randomized, and signature-free implementations, the protocols of Mostéfaoui *et al.* (PODC 2014 and JACM 2015) represent a landmark result, which has been extended later and taken up in practical systems. The protocols achieve optimal resilience and takes, in expectation, only a constant expected number of rounds of quadratic message complexity. Randomization is provided through a common-coin primitive.

In traditional consensus protocols, all involved processes adhere to a global, symmetric failure model, typically only defined by bounds on the number of faulty processes. Motivated by applications to blockchains, however, more flexible trust assumptions have recently been considered. In particular, with *asymmetric trust*, a process is free to choose which other processes it trusts and which ones might collude against it.

This paper revisits the optimal asynchronous protocol of Mostéfaoui *et al.* and shows how to realize it with asymmetric trust. The paper starts by pointing out in detail why some versions of this protocol may violate liveness. Then it proposes a fix for the protocol that does not affect its properties, but lets it regain the simplicity of its original version (PODC 2014). At the same time, the paper shows how to realize randomized signature-free asynchronous Byzantine consensus with asymmetric quorums. This results in an optimal consensus protocol with subjective, asymmetric trust and constant expected running time. It is suitable for applications to blockchains, for instance.

**Key words.** Consensus, Asymmetric trust, Quorums, Randomized consensus, Common coin.

## 1 Introduction

Consensus represents a fundamental abstraction in distributed systems. It captures the problem of reaching agreement among multiple processes on a common value, despite unreliable communication and the presence of faulty processes. Most protocols for consensus operate under the assumption that the *number* of faulty processes is limited. Moreover, all processes in the system share this common *trust assumption*. Traditionally, the trust assumption has been *symmetric* in this sense: all processes adhere to the global assumption about the number of faulty processes and properties of protocols are guaranteed for all correct processes, but not for the faulty ones. Since the advent blockchains systems, however, more flexible trust models have been introduced. The Ripple ([www.ripple.com](http://www.ripple.com)) and Stellar ([www.stellar.org](http://www.stellar.org)) blockchains have pioneered practical models that let each process express its *own set* of trusted processes and assumptions can be more flexible than bounding only the number of faulty processes.

Motivated by this desire to make trust assumptions more flexible, Cachin and Tackmann [5] introduced *asymmetric Byzantine quorum systems* as a generalization of Byzantine quorum systems. Originally defined by Malkhi and Reiter [15], Byzantine quorum systems capture one global, but arbitrarily complex trust relation through a so-called fail-prone system. This permits protocols, in which not only the number of faults is bounded. Since Byzantine quorum systems provide a widely used abstraction for realizing practical consensus protocols for distributed systems, asymmetric quorum systems open up the possibility to implement consensus with subjective trust. However, no consensus algorithms with asymmetric trust have been formulated so far.

---

<sup>1</sup>Institute of Computer Science, University of Bern, Neubrückstrasse 10, 3012 CH-Bern, Switzerland.

In this paper, we present the first asynchronous Byzantine consensus protocol with asymmetric trust. It uses randomization, provided by an asymmetric common-coin protocol, to circumvent the impossibility of (purely) asynchronous consensus. Our protocol takes up the randomized and signature-free implementation of consensus by Mostéfaoui *et al.* [18, 19]. This represents a landmark result because it has been praised for its simplicity, was the first to achieve optimal complexity, that is, expected quadratic cost in the number of processes, and does not use digital signatures. The protocol has been extended later and taken up in practical systems, such as “Honey Badger BFT” [17].

The protocol of Mostéfaoui *et al.*, however, comes in multiple versions. The original one, published at PODC 2014 [18] and where it also won the best-paper award, suffers from a subtle and little-known liveness problem [22]: an adversary can prevent progress among the correct processes by controlling the messages between them and by sending them values in a specific order. The subsequent version (JACM 2015) [19] resolves this issue, but requires many more communication steps and adds considerable complexity.

Our asymmetric asynchronous Byzantine consensus protocol is based on the simpler version (PODC 2014). We first revisit this and show in detail how it is possible to violate liveness. We propose a method that overcomes the problem, maintains the elegance of the protocol, and does not affect its appealing properties. Based on this insight, we show how to realize asynchronous consensus with asymmetric trust, again with a protocol that maintains the simplicity of the original approach of Mostéfaoui *et al.* [18].

Asymmetric quorum systems go back to the notion of asymmetric trust introduced by Damgård *et al.* [7]. Every process in the system subjectively selects its own fail-prone system. Depending on the choice that a correct process makes about who it trusts and who not, and considering the processes that are actually faulty during an execution, two different situations may arise. A correct process may either make a “wrong” trust assumption, for example, by trusting too many processes that turn out to be faulty or by tolerating too few faults; such a process is called *naïve*. Alternatively, when the correct process makes the “right” trust assumption, it is called *wise*. Protocols with asymmetric trust do not guarantee the same properties for naïve processes as for wise ones.

As an additional contribution, we extend our knowledge about the relation between naïve and wise processes in protocols with asymmetric trust. We show that, under certain conditions, guarantees can only be given for a subset of the wise processes that form a so-called *guild*. The existence of a guild is necessary for a protocol execution with asymmetric trust to terminate.

The remainder of this work is structured as follows. In Section 2 we discuss related work. We present our system model together with preliminaries on Byzantine quorums in Section 3. In Section 4 we recall the randomized consensus protocol as originally introduced by Mostéfaoui *et al.* [18], discuss the liveness issue [22], and show a way to prevent it. In Section 5 we recall and extend the theory behind asymmetric quorums. We define and implement asymmetric strong Byzantine consensus protocol in Section 6 by extending and improving on the randomized consensus algorithm by Mostéfaoui *et al.* [18]. Moreover, we build a common coin based on secret sharing that works in the asymmetric-trust model and it is used in our randomized protocol. Conclusions are drawn in Section 7.

## 2 Related work

Mostéfaoui *et al.* [18] present a signature-free round-based asynchronous consensus algorithm for binary values. It achieves optimal resilience and takes  $O(n^2)$  constant-sized messages. The algorithm is randomized and randomization is achieved through a common coin as defined by Rabin [21]. Their algorithm has been taken up for the construction of the HoneyBadgerBFT by Miller *et al.* [17].

Tholoniati and Gramoli [22] observe a liveness issue in the protocol by Mostéfaoui *et al.* [18] in which an adversary is able to prevent progress among the correct processes by controlling messages between them and by sending them values in a specific order.

In a later work, Mostéfaoui *et al.* [19] present a different version of their randomized consensus algorithm that does not suffer from the liveness problem anymore. The resulting algorithm offers the same asymptotic complexity in message and time as their previous algorithm [18] but it requires more communication steps.

Flexible trust structures have recently received a lot of attention [5, 7, 8, 13, 14, 16], primarily motivated by consensus protocols for blockchains, as introduced by Ripple ([www.ripple.com](http://www.ripple.com)) and Stellar ([www.stellar.org](http://www.stellar.org)). According to general idea behind these models, processes are free to express individual, *subjective* trust choices about other processes, instead of adopting a common, global view of trust.

Damgård *et al.* [7] define the basics of *asymmetric trust* for secure computation protocols. This model is strictly more powerful than the standard model with symmetric trust and abandons the traditional global failure assumption in the system. Moreover, they present several variations of their asymmetric-trust model and sketch synchronous protocols for broadcast, verifiable secret sharing, and general multi-party computation.

Mazières [16] introduces a new model for consensus called *federated Byzantine agreement* (FBA) and uses it to construct the *Stellar consensus protocol* [12]. In FBA, every process declares *quorums slices* – a collection of trusted sets of processes sufficient to convince the particular process of agreement. These slices are subsets of a *quorum*, which is a set of processes sufficient to reach agreement. More precisely, a quorum is defined as a set of processes that contains one slice for each member, and all quorums constitute a *federated Byzantine quorum system* (FBQS).

*Byzantine quorum systems* have originally been formalized by Malkhi and Reiter [15] and exist in several forms; they generalize the classical quorum systems aimed at tolerating crashes to algorithms with Byzantine failures. Byzantine quorum systems assume one global shared fail-prone system.

A link between FBQS and Byzantine quorums system has been built by García-Pérez and Gotsman [8], who implement Byzantine reliable broadcast on an FBQS. They prove that a FBQS *induces* a Byzantine quorum system.

*Asymmetric Byzantine quorum systems* have been introduced by Cachin and Tackmann [5] and generalize Byzantine quorum systems [15] to the model with asymmetric trust. This work also explores properties of asymmetric Byzantine quorum systems and differences to the model with symmetric trust. In particular, Cachin and Tackmann [5] distinguish between different classes of correct processes, depending on whether their failure assumptions in an execution are correct. The standard properties of protocols are guaranteed only to so-called *wise* processes, i.e., those that made the “right” trust choices. Protocols with asymmetric quorums are shown for Byzantine consistent broadcast, reliable broadcast, and emulations of shared memory. In contrast to FBQS, asymmetric quorum systems appear to be a natural extension of symmetric quorum systems.

Recently, Losa *et al.* [13] have formulated an abstraction of the consensus mechanism in the Stellar network by introducing *Personal Byzantine quorum systems* (PBQS). In contrast to the other notions “quorums”, their definition does not require a global intersection among quorums. This may lead to several separate *consensus clusters* such that each one satisfies agreement and liveness on its own.

Another new approach for designing Byzantine fault-tolerant (BFT) consensus protocols has been introduced by Malkhi *et al.* [14], namely *Flexible BFT*. This notion guarantees higher resilience by introducing a new *alive-but-corrupt* fault type, which denotes processes that attack safety but not liveness. Malkhi *et al.* [14] also define *flexible Byzantine quorums* that allow processes in the system to have different faults models.

## 3 System model and preliminaries

### 3.1 System model

**Processes.** We consider a system of  $n$  processes  $\mathcal{P} = \{p_1, \dots, p_n\}$  that communicate with each other. The processes interact by exchanging messages over reliable point-to-point links, specified below.

A protocol for  $\mathcal{P}$  consists of a collection of programs with instructions for all processes. Protocols are presented in a modular way using the event-based notation of Cachin *et al.* [3].

**Failures.** A process that follows its protocol during an execution is called *correct*. On the other hand, a *faulty* process may crash or deviate arbitrarily from its specification, e.g., when *corrupted* by an adversary; such processes are also called *Byzantine*. We consider only Byzantine faults here and assume for simplicity that the faulty processes fail right at the start of an execution.

**Functionalities and modularity.** A *functionality* is an abstraction of a distributed computation, either used as a primitive available to the processes or defining a service that a protocol run by the processes will provide. Functionalities may be composed in a modular way. Every functionality in the system is specified through its *interface*, containing the *events* that it exposes applications that may call it, and through a number of *properties* that define its behavior. There are two kinds of events in an interface: *input events* that the functionality receives from other abstractions, typically from an application that invokes its services, and *output events*, through which the functionality delivers information or signals a condition.

Multiple functionalities may be composed together modularly. In a modular protocol implementation, in particular, every process executes the program instructions of the protocol implementations for all functionalities in which it participates.

**Links.** We assume there is a low-level functionality for sending messages over point-to-point links between each pair of processes. In a protocol, this functionality is accessed through the events of “sending a message” and “receiving a message.” Point-to-point messages are authenticated and delivered reliably among correct processes.

Moreover, at certain points, we assume FIFO ordering on the reliable point-to-point links for every pair of processes. This means that if a correct process has “sent” a message  $m_1$  and subsequently “sent” a message  $m_2$ , then every correct process does not “receive”  $m_2$  unless it has earlier also “received”  $m_1$ . Protocols that guarantee FIFO order on top of (unordered) reliable point-to-point links are well-known and simple to implement [3, 9]. We remark that there is only one FIFO-ordered reliable point-to-point link functionality in the model; hence, FIFO order holds among the messages exchanged by the implementations for all functionalities used by a protocol.

**Time and randomization.** In this work we consider an asynchronous system, where processes have no access to any kind of physical clock, and there is no bound on processing or communication delays. The randomized consensus algorithm delegates probabilistic choices to a *common coin* abstraction [21]; this is a functionality that delivers the same sequence of random binary values to each process, where each binary value has the value 0 or 1 with probability  $\frac{1}{2}$ .

### 3.2 Byzantine quorum systems

Let us recall Byzantine quorums as originally introduced [15]. We refer to them as *symmetric* Byzantine quorums.

**Definition 1 (Fail-prone system).** Let  $\mathcal{P}$  be a set of processes. A *fail-prone system*  $\mathcal{F}$  is a collection of subsets of  $\mathcal{P}$ , none of which is contained in another, such that some  $F \in \mathcal{F}$  with  $F \subseteq \mathcal{P}$  is called a *fail-prone set* and contains all processes that may at most fail together in some execution.

**Definition 2 (Symmetric Byzantine quorum system).** Let  $\mathcal{P}$  be a set of processes and let  $\mathcal{F} \subseteq 2^{\mathcal{P}}$  be a *fail-prone system*. A *symmetric Byzantine quorum system* for  $\mathcal{F}$  is a collection of sets of processes  $\mathcal{Q} \subseteq 2^{\mathcal{P}}$ , where each  $Q \in \mathcal{Q}$  is called a *quorum*, such that

**Consistency:**

$$\forall Q_1, Q_2 \in \mathcal{Q}, \forall F \in \mathcal{F} : Q_1 \cap Q_2 \not\subseteq F.$$

**Availability:**

$$\forall F \in \mathcal{F} : \exists Q \in \mathcal{Q} : F \cap Q = \emptyset.$$

For example, under the common threshold failure model, the quorums are all sets of at least  $\lceil \frac{n+f+1}{2} \rceil$  processes, where  $f$  is the number of processes that may fail. In particular, if  $n = 3f + 1$ , quorums have  $2f + 1$  or more processes.

Malkhi and Reiter [15] refer to the above definition as *Byzantine dissemination quorum system*. They also define other variants of Byzantine quorum systems.

Notice that in our notion of a quorum system, one quorum can be contained in another.

We say that a set system  $\mathcal{T}$  *dominates* another set system  $\mathcal{S}$  if for each  $S \in \mathcal{S}$  there is some  $T \in \mathcal{T}$  such that  $S \subseteq T$ . In this sense, a quorum system for  $\mathcal{F}$  is *minimal* whenever it does not dominate any other quorum system for  $\mathcal{F}$ .

**Definition 3 ( $Q^3$ -condition [10, 15]).** Let  $\mathcal{F}$  be a fail-prone system. We say that  $\mathcal{F}$  satisfies the  $Q^3$ -condition, abbreviated as  $Q^3(\mathcal{F})$ , if it holds

$$\forall F_1, F_2, F_3 \in \mathcal{F} : \mathcal{P} \not\subseteq F_1 \cup F_2 \cup F_3.$$

This is the generalization of the threshold condition  $n > 3f$  for Byzantine quorum systems. Let  $\bar{S} = \{\mathcal{P} \setminus S \mid S \in \mathcal{S}\}$  be the *bijective complement* of a set  $\mathcal{S} \subseteq 2^{\mathcal{P}}$ .

**Lemma 1 (Quorum system existence [15]).** Let  $\mathcal{F}$  be a fail-prone system. A Byzantine quorum system for  $\mathcal{F}$  exists if and only if  $Q^3(\mathcal{F})$ . In particular, if  $Q^3(\mathcal{F})$  holds, then  $\bar{\mathcal{F}}$ , the bijective complement of  $\mathcal{F}$ , is a Byzantine quorum system called *canonical quorum system of  $\mathcal{F}$* .

Note that the canonical quorum system is not always minimal. The canonical quorum system will play a role in Section 6 for implementing a common-coin functionality with asymmetric quorums.

Given a symmetric Byzantine quorum system  $\mathcal{Q}$ , we define a *kernel*  $K$  as a set of processes that overlaps with every quorum. A kernel generalizes the notion of a *core set* [11].

**Definition 4 (Kernel system).** A set  $K \subseteq \mathcal{P}$  is a *kernel* of a quorum system  $\mathcal{Q}$  whenever it holds

$$\forall Q \in \mathcal{Q} : K \cap Q \neq \emptyset.$$

This can be viewed as a *consistency* property.

We also define the *kernel system*  $\mathcal{K}$  of  $\mathcal{Q}$  to be the set of all kernels of  $\mathcal{Q}$ . Given this, the *minimal kernel system* is a kernel system for which every kernel  $K$  satisfies

$$\forall K' \subsetneq K, \exists Q \in \mathcal{Q} : K' \cap Q = \emptyset.$$

For example, under a threshold failure assumption where any  $f$  processes may fail, every set of  $\lfloor \frac{n-f+1}{2} \rfloor$  processes is a kernel. In particular,  $n = 3f + 1$  if and only if every kernel has  $f + 1$  processes.

**Lemma 2.** *For every  $F \in \mathcal{F}$  and for every quorum  $Q \in \mathcal{Q}$  there exists a kernel  $K \in \mathcal{K}$  such that  $K \subseteq Q$ .*

*Proof.* Let  $\mathcal{Q}$  be a quorum system for  $\mathcal{F}$  and let  $F \in \mathcal{F}$ . From the consistency property of a quorum system we have that for all  $Q_1, Q_2 \in \mathcal{Q}$  it holds  $Q_1 \cap Q_2 \not\subseteq F$ . Then, the set  $K = Q_1 \setminus F \subseteq Q_1$  intersects all quorums in  $\mathcal{Q}$  and is a kernel of  $\mathcal{Q}$ .  $\square$

## 4 Revisiting signature-free asynchronous Byzantine consensus

In 2014, Mostéfaoui *et al.* [18] introduced a round-based asynchronous randomized consensus algorithm for binary values. It had received considerable attention because it was the first protocol with optimal resilience, tolerating up to  $f < \frac{n}{3}$  Byzantine processes, that did not use digital signatures. Hence, this protocol needs only authenticated channels and remains secure against a computationally unbounded adversary. Moreover, it takes  $O(n^2)$  constant-sized messages in expectation and has a particularly simple structure. This description excludes the necessary cost for implementing randomization, for which the protocol relies on an abstract common-coin primitive, as defined by Rabin [21].

This protocol, which we call the *PODC-14* version [18] in the following, suffers from a subtle and little-known problem. It may violate liveness, as has been explicitly mentioned by Tholoniati and Gramoli [22]. The corresponding journal publication by Mostéfaoui *et al.* [19], to which we refer as the *JACM-15* version, touches briefly on the issue and goes on to present an extended protocol. This fixes the problem, but requires also many more communication steps and adds considerable complexity.

In this section, we revisit the PODC-14 protocol, point out in detail how it may fail, and introduce a compact way for fixing it. We discovered these issues while extending the algorithm to asymmetric quorums. In Section 6, we present the corresponding fixed asymmetric randomized Byzantine consensus protocol and prove it secure. Our protocol changes the PODC-14 version in a crucial way and thereby regains the simplicity of the original approach.

Before addressing randomized consensus, we recall the key abstraction introduced in the PODC-14 paper, a protocol for broadcasting binary values.

### 4.1 Binary-value broadcast

The *binary validated broadcast* primitive has been introduced in the PODC-14 version [18] under the name *binary-value broadcast*.<sup>1</sup> In this primitive, every process may broadcast a bit  $b \in \{0, 1\}$  by invoking *bv-broadcast*( $b$ ). The broadcast primitive outputs at least one value  $b$  and possibly also both binary values through a *bv-deliver*( $b$ ) event, according to the following notion.

**Definition 5 (Binary validated broadcast).** A protocol for *binary validated broadcast* satisfies the following properties:

**Validity:** If at least  $(f + 1)$  correct processes *bv-broadcast* the same value  $b \in \{0, 1\}$ , then every correct process eventually *bv-delivers*  $b$ .

**Integrity:** A correct process *bv-delivers* a particular value  $b$  at most once and only if  $b$  has been *bv-broadcast* by some correct process.

**Agreement:** If a correct process *bv-delivers* some value  $b$ , then every correct process eventually *bv-delivers*  $b$ .

**Termination:** Every correct process eventually *bv-delivers* some value  $b$ .

The implementation given by Mostéfaoui *et al.* [18] works as follows. When a correct process  $p_i$  invokes *bv-broadcast*( $b$ ) for  $b \in \{0, 1\}$ , it sends a VALUE message containing  $b$  to all processes. Afterwards, whenever a correct process receives VALUE messages containing  $b$  from at least  $f + 1$  processes and has not itself sent a VALUE message containing  $b$ , then it sends such message to every process. Finally, once a correct process receives VALUE messages containing  $b$  from at least  $2f + 1$  processes, it delivers  $b$  through *bv-deliver*( $b$ ). Notice that a process may *bv-deliver* up to two values. A formal description (in the more general asymmetric model) of this protocol appears in Algorithm 3 in Section 6.

<sup>1</sup>Compared to their work, we adjusted some conditions to standard terminology and chose to call the primitive “binary validated broadcast” to better emphasize its aspect of validating that a delivered value was broadcast by a correct process.

## 4.2 Randomized consensus

We recall the notion of *randomized Byzantine consensus* here and its implementation by Mostéfaoui *et al.* [18].

In a consensus primitive, every correct process proposes a value  $v$  by invoking  $propose(v)$ , which typically triggers the start of the protocol among processes; it obtains as output a decided value  $v$  through a  $decide(v)$  event. There are no assumptions made about the faulty processes.

Observe that when working with randomized consensus, one has to formulate the termination property probabilistically. In round-based consensus algorithms, the termination property is formulated with respect to the round number  $r$  that a process executes. It requires that the probability that a correct process decides after executing infinitely many rounds approaches 1.

**Definition 6 (Strong Byzantine consensus).** A protocol for asynchronous *strong Byzantine consensus* satisfies:

**Probabilistic termination:** Every correct process decides with probability 1, in the sense that

$$\lim_{r \rightarrow +\infty} \text{P}[\text{a correct process } p_i \text{ decides by round } r] = 1.$$

**Strong validity:** A correct process only decides a value that has been proposed by some correct process.

**Integrity:** No correct process decides twice.

**Agreement:** No two correct processes decide differently.

The probabilistic termination and integrity properties together imply that every correct process decides exactly once, while the agreement property ensures that the decided values are equal. Strong validity asks that if all correct processes propose the same value  $v$ , then no correct process decides a value different from  $v$ . Otherwise, a correct process may only decide a value that was proposed by some correct process [3]. In a *binary* consensus protocol, as considered here, only 0 and 1 may be proposed. In this case, strong validity is equivalent to the more commonly used property of weak validity.

The implementation of the randomized consensus algorithm Mostéfaoui *et al.* [18] delegates its probabilistic choices to a *common coin* abstraction [3,21], a random source observable by all processes but unpredictable for an adversary. A common coin is invoked at every process by triggering a *release-coin* event. We say that a process *releases* a coin because its value is unpredictable before the first correct process invokes the coin. The value  $s \in \mathcal{B}$  of the coin with tag  $r$  is output through an event *output-coin*.

**Definition 7 (Common coin).** A protocol for *common coin* satisfies the following properties:

**Termination:** Every correct process eventually outputs a coin value.

**Unpredictability:** Unless at least one correct process has released the coin, no process has any information about the coin output by a correct process.

**Matching:** With probability 1 every correct process outputs the same coin value.

**No bias:** The distribution of the coin is uniform over  $\mathcal{B}$ .

Common-coin primitives may be realized directly by distributed protocols or with the help of a trusted entity using secret sharing [21] or threshold cryptography [4].

In the remainder of this section, we recall the implementation of strong Byzantine consensus according to Mostéfaoui *et al.* [18] in the PODC-14 version, shown in Algorithm 1. A correct process *proposes* a binary value  $b$  by invoking  $rbc-propose(b)$ ; the consensus abstraction *decides* for  $b$  through an  $rbc-decide(b)$  event.

The algorithm proceeds in rounds. In each round, an instance of *bv-broadcast* is invoked. A correct process  $p_i$  executes *bv-broadcast* and waits for a value  $b$  to be *bv-delivered*, identified by a tag characterizing the current round. When such a bit  $b$  is received,  $p_i$  adds  $b$  to *values* and broadcasts  $b$  through an AUX message to all processes. Whenever a process receives an AUX message containing  $b$  from  $p_j$ , it stores  $b$  in a local set  $aux[j]$ . Once  $p_i$  has received a set  $B \subseteq \text{values}$  of values such that every  $b \in B$  has been delivered in AUX messages from at least  $n - f$  processes, then  $p_i$  releases the coin for the round. Subsequently, the process waits for the coin protocol to output a binary value  $s$  through  $output-coin(s)$ , tagged with the current round number.

Process  $p_i$  then checks if there is a single value  $b$  in  $B$ . If so, and if  $b = s$ , then it decides for value  $b$ . The process then proceeds to the next round with proposal  $b$ . If there is more than one value in  $B$ , then  $p_i$  changes its proposal to  $s$ . In any case, the process starts another round and invokes a new instance of *bv-broadcast* with its proposal. Notice that the protocol appears to execute rounds forever.

---

**Algorithm 1** Randomized binary consensus according to Mostéfaoui *et al.* [18] (code for  $p_i$ ).

---

```
1: State
2:    $round \leftarrow 0$ : current round
3:    $values \leftarrow \{\}$ : set of bv-delivered binary values for the round
4:    $aux \leftarrow [\{\}]^n$ : stores sets of values that have been received in AUX messages in the round
5:
6: upon event rbc-propose( $b$ ) do
7:   invoke bv-broadcast( $b$ ) with tag  $round$ 
8:
9: upon bv-deliver( $b$ ) with tag  $r$  such that  $r = round$  do
10:   $values \leftarrow values \cup \{b\}$ 
11:  send message [AUX,  $round, b$ ] to all  $p_j \in \mathcal{P}$ 
12:
13: upon receiving a message [AUX,  $r, b$ ] from  $p_j$  such that  $r = round$  do
14:   $aux[j] \leftarrow aux[j] \cup \{b\}$ 
15:
16: upon exists  $B \subseteq values$  such that  $B \neq \{\}$  and  $|\{p_j \in \mathcal{P} \mid B = aux[j]\}| \geq n - f$  do
17:  release-coin with tag  $round$ 
18:  wait for output-coin( $s$ ) with tag  $round$ 
19:   $round \leftarrow round + 1$ 
20:  if exists  $b$  such that  $|B| = 1 \wedge B = \{b\}$  then
21:    if  $b = s$  then
22:      output rbc-decide( $b$ )
23:      invoke bv-broadcast( $b$ ) with tag  $round$  // propose  $b$  for the next round
24:    else
25:      invoke bv-broadcast( $s$ ) with tag  $round$  // propose coin value  $s$  for the next round
26:   $values \leftarrow [\perp]^n$ 
27:   $aux \leftarrow [\{\}]^n$ 
```

---

### 4.3 A liveness problem

Tholoniati and Gramoli [22] mention a liveness issue with the randomized algorithm in the PODC-14 version [18], as presented in the previous section. They sketch a problem that may prevent progress by the correct processes when the messages between them are received in a specific order. In the JACM-15 version, Mostéfaoui *et al.* [19] appear to be aware of the issue and present a different, more complex consensus protocol.


We give a detailed description of the problem in Algorithm 1. Recall the implementation of binary-value broadcast, which disseminates bits in VALUE messages. According to our model, the processes communicate by exchanging messages through an asynchronous reliable point-to-point network. Messages may be reordered, as in the PODC-14 version.

Let us consider a system with  $n = 4$  processes and  $f = 1$  Byzantine process. Let  $p_1, p_2$  and  $p_3$  be correct processes with input values 0, 1, 1, respectively, and let  $p_4$  be a Byzantine process with control over the network. Process  $p_4$  aims to cause  $p_1$  and  $p_3$  to release the coin with  $B = \{0, 1\}$ , so that they subsequently propose the coin value for the next round. If messages are scheduled depending on knowledge of the round's coin value  $s$ , it is possible, then, that  $p_2$  releases the coin with  $B = \{-s\}$ . Subsequently,  $p_2$  proposes also  $\neg s$  for the next round, and this may continue forever. We now work out the details, as illustrated in Figures 1–2.


First,  $p_4$  may cause  $p_1$  to receive  $2f + 1$  [VALUE, 1] messages, from  $p_2, p_3$  and  $p_4$ , and to *bv-deliver* 1 sent at the start of the round. Then,  $p_4$  sends [VALUE, 0] to  $p_3$ , so that  $p_3$  receives value 0 twice (from  $p_1$  and  $p_4$ ) and also broadcasts a [VALUE, 0] message itself. Process  $p_4$  also sends 0 to  $p_1$ , hence,  $p_1$  receives 0 from  $p_3, p_4$ , and itself and therefore *bv-delivers* 0. Furthermore,  $p_4$  causes  $p_3$  to *bv-deliver* 0 by making it receive [VALUE, 0] messages from  $p_1, p_4$ , and itself. Hence,  $p_3$  *bv-delivers* 0. Finally, process  $p_3$  receives three [VALUE, 1] messages (from itself,  $p_2$ , and  $p_4$ ) and *bv-delivers* also 1.

Recall that a process may broadcast more than one AUX message. In particular, it broadcasts an AUX message containing a bit  $b$  whenever it has *bv-delivered*  $b$ . Thus,  $p_1$  broadcasts first [AUX, 1] and subsequently [AUX, 0], whereas  $p_3$  first broadcasts [AUX, 0] and then [AUX, 1]. Process  $p_4$  then sends to  $p_1$  and  $p_3$  AUX messages containing 1 and 0. After delivering all six AUX messages, both  $p_1$  and  $p_3$  finally obtain  $B = \{0, 1\}$  in line 16 and see that  $|B| \neq 1$  in line 20. Processes  $p_1, p_3$  and  $p_4$  invoke the common coin.

The Byzantine process  $p_4$  may learn the coin value as soon as  $p_1$  or  $p_3$  have released the common coin,



**Figure 1.** The execution of Algorithm 1, where processes  $p_1$  and  $p_3$  execute line 16 with  $B = \{0, 1\}$ .



**Figure 2.** Continuing the execution for the case  $s = 0$ : Process  $p_2$  executes line 16 with  $B = \{1\}$ . Processes  $p_1$  and  $p_3$  have already proposed the coin value  $s = 0$  for the next round, but  $p_2$  proposes  $\neg s = 1$ .

according to unpredictability. Let  $s$  the coin output. We distinguish two cases:

**Case  $s = 0$ :** Process  $p_2$  receives now three  $[\text{VALUE}, 1]$  messages, from  $p_3$ ,  $p_4$  and itself, as shown in Figure 2. It *bv-delivers* 1 and broadcasts an  $[\text{AUX}, 1]$  message. Subsequently,  $p_2$  delivers three AUX messages containing 1, from  $p_1$ ,  $p_4$  and itself, but no  $[\text{AUX}, 0]$  message. It follows that  $p_2$  obtains  $B = \{1\}$  and proposes 1 for the next round in line 23. On the other hand,  $p_1$  and  $p_3$  adopt 0 as their new proposal for the next round, according to line 25. This means that no progress was made within this round. The three correct processes start the next round again with differing values, again two of them propose one bit and the remaining one proposes the opposite.

**Case  $s = 1$ :** Process  $p_4$  sends  $[\text{VALUE}, 0]$  to  $p_2$ , so that it delivers two VALUE messages containing 0 (from  $p_1$  and  $p_4$ ) and thus also broadcast  $[\text{VALUE}, 0]$  (this execution is not shown). Recall that  $p_3$  has already sent  $[\text{VALUE}, 0]$  before. Thus,  $p_2$  receives  $n - f$   $[\text{VALUE}, 0]$  messages, *bv-delivers* 0, and also broadcasts an AUX message containing 0. Subsequently,  $p_2$  may receive  $n - f$  messages  $[\text{AUX}, 0]$ , from  $p_3$ ,  $p_4$ , and itself. It follows that  $p_2$  executes line 16 with  $B = \{0\}$  and chooses 0 as its proposal for the next round (in line 23). On the other hand, also here,  $p_1$  and  $p_3$  adopt the coin value  $s = 1$  and propose 1 for the next round in line 25. Hence, no progress has been made in this round, as the three correct processes enter the next round with differing values.

The execution may continue like this forever, producing an infinite execution with no termination.



## 4.4 Fixing the problem

We show how the problem can be prevented with two small changes to the original protocol. We do this by recalling the example just presented. A formal proof is given in Section 6, using the more general model of asymmetric quorums.

Observe that in the problematic execution, the network reorders messages between correct processes. Our first change, therefore, is to assume FIFO ordering on the reliable point-to-point links. This may be implemented over authenticated links, by adding sequence numbers to messages and maintaining a buffer at the receiver [3]. Consider  $p_2$  in the example and the messages it receives from the other correct processes,  $p_1$  and  $p_3$ . W.l.o.g. any protocol implementing a common coin requires an additional message exchange, where a correct process sends at least one message to every other process, say, a COIN message with arbitrary content (to be specific, see Algorithm 2, Section 6).

When  $p_2$  waits for the output of the coin, it needs to receive, again w.l.o.g., a COIN message from  $n - f$  processes. Since the other two correct processes ( $p_1$  and  $p_3$ ) have sent two VALUE messages and AUX messages each before releasing the coin, then  $p_2$  receives these messages from at least one of them before receiving enough COIN messages, according to the overlap among Byzantine quorums.

This means that  $p_2$  cannot satisfy the condition in line 16 with  $|B| = 1$ . Thus the adversary may no longer exploit its knowledge of the coin value to prevent termination. (Mostéfaoui *et al.* [19] (JACM-15) remark in retrospect about the PODC-14 version that a “fair scheduler” is needed. However, this comes without any proof and thus remains open, especially because the JACM-15 version introduces a much more complex version of the protocol.)

Our second change is to allow the set  $B$  to dynamically change while the coin protocol executes. In this way, process  $p_2$  may find a suitable  $B$  according to the received AUX messages while concurrently running the coin protocol. Eventually,  $p_2$  will have output the coin *and* its set  $B$  will contain the same values as the sets  $B$  of  $p_1$  and  $p_3$ . Observe that this dynamicity is necessary; process  $p_2$  could start to release the coin after receiving  $n - f$  AUX messages containing only the value 1. However, following our example, due to the assumed FIFO order, it will receive from another correct process also an AUX message containing the value 0, before the COIN message. If we do not ask for the dynamicity of the set  $B$ , process  $p_2$ , after outputting the coin, will still have  $|B| = 1$ . Mostéfaoui *et al.* in the PODC-14 version (Figure 2, line 5 [18]) seem to rule this out.

In Section 6, we implement these changes in a generalization of the PODC-14 version, show that the liveness problem presented in this section does no longer applies (Lemma 8), and give a formal proof for the correctness of our protocol (Theorem 9). The generalization works in the asymmetric-trust model, as defined in the following section.

## 5 Asymmetric trust

In this section, we first review and extend the model of asymmetric trust, as introduced by Damgård *et al.* [7] and by Cachin and Tackmann [5]. We first recall asymmetric quorums. Then we focus on a *maximal guild*, which is needed for ensuring liveness and consistency in protocols, and we prove that the maximal guild is unique. We also characterize it in relation to *wise* processes, which are those correct processes whose a failure assumption turns out to be right.

In the asymmetric-trust model, every process is free to make its own trust assumption, expressing it through a subjective fail-prone system.

**Definition 8 (Asymmetric fail-prone system).** An asymmetric fail-prone system  $\mathbb{F} = [\mathcal{F}_1, \dots, \mathcal{F}_n]$  consists of an array of fail-prone systems, where  $\mathcal{F}_i \subseteq 2^{\mathcal{P}}$  denotes the trust assumption of  $p_i$ .

One often assumes that  $\forall F \in \mathcal{F}_i : p_i \notin F$  for practical reasons, but this is not necessary. For a system  $\mathcal{A} \subseteq 2^{\mathcal{P}}$ , let  $\mathcal{A}^* = \{A' \mid A' \subseteq A, A \in \mathcal{A}\}$  denote the collection of all subsets of the sets in  $\mathcal{A}$ .

**Definition 9 (Asymmetric Byzantine quorum system).** Let  $\mathbb{F} = [\mathcal{F}_1, \dots, \mathcal{F}_n]$  be an asymmetric fail-prone system. An *asymmetric Byzantine quorum system* for  $\mathbb{F}$  is an array of collections of sets  $\mathcal{Q} = [\mathcal{Q}_1, \dots, \mathcal{Q}_n]$ , where  $\mathcal{Q}_i \subseteq 2^{\mathcal{P}}$  for  $i \in [1, n]$ . The set  $\mathcal{Q}_i \subseteq 2^{\mathcal{P}}$  is called the *quorum system of  $p_i$*  and any set  $Q_i \in \mathcal{Q}_i$  is called a *quorum (set) for  $p_i$*  whenever the following conditions hold:

**Consistency:**  $\forall i, j \in [1, n]$

$$\forall Q_i \in \mathcal{Q}_i, \forall Q_j \in \mathcal{Q}_j, \forall F_{ij} \in \mathcal{F}_i^* \cap \mathcal{F}_j^* : Q_i \cap Q_j \not\subseteq F_{ij}.$$

**Availability:**  $\forall i \in [1, n]$

$$\forall F_i \in \mathcal{F}_i : \exists Q_i \in \mathcal{Q}_i : F_i \cap Q_i = \emptyset.$$

In other words, the intersection of two quorums for any two processes contains at least one process for which neither process assumes that it may fail. Furthermore, for all fail-prone sets of every process, there exists a disjoint quorum for this process.

The following property generalizes the  $Q^3$ -condition from Definition 3 to the asymmetric-trust model.

**Definition 10 ( $B^3$ -condition [5, 7]).** Let  $\mathbb{F}$  be an asymmetric fail-prone system. We say that  $\mathbb{F}$  satisfies the  $B^3$ -condition, abbreviated as  $B^3(\mathbb{F})$ , whenever it holds for all  $i, j \in [1, n]$  that

$$\forall F_i \in \mathcal{F}_i, \forall F_j \in \mathcal{F}_j, \forall F_{ij} \in \mathcal{F}_i^* \cap \mathcal{F}_j^* : \mathcal{P} \not\subseteq F_i \cup F_j \cup F_{ij}.$$

An asymmetric fail-prone system satisfying the  $B^3$ -condition is sufficient for the existence of a corresponding asymmetric quorum system [5].

**Theorem 3.** An asymmetric fail-prone system  $\mathbb{F}$  satisfies  $B^3(\mathbb{F})$  if and only if there exists an asymmetric quorum system for  $\mathbb{F}$ .

For implementing consensus, we also need the notion of an asymmetric kernel system.

**Definition 11 (Asymmetric kernel system).** Let  $\mathbb{Q} = [Q_1, \dots, Q_n]$  be an asymmetric quorum system. An *asymmetric kernel system*  $\mathbb{K}$  is an array of collections of sets  $[\mathcal{K}_1, \dots, \mathcal{K}_n]$  such that each  $\mathcal{K}_i$  is a kernel system of  $Q_i$ . We call a set  $K_i \in \mathcal{K}_i$  a *kernel* for  $p_i$ .

In traditional Byzantine quorum systems, under a symmetric-trust assumption, every process in the system adheres to a global fail-prone system  $\mathcal{F}$  and the set  $F$  of faults or corruptions occurring in a protocol execution is in  $\mathcal{F}$ . Given this common trust assumption, properties of a protocol are guaranteed at each correct process, while they are not guaranteed for faulty ones. With asymmetric quorums, there is a distinction among correct processes with respect to  $F$ , namely the correct processes that consider  $F$  in their trust assumption and those who do not. Given a protocol execution, the processes are classified in three different types:

**Faulty:** A process  $p_i \in F$  is *faulty*.

**Naïve:** A correct process  $p_i$  for which  $F \notin \mathcal{F}_i^*$  is called *naïve*.

**Wise:** A correct process  $p_i$  for which  $F \in \mathcal{F}_i^*$  is called *wise*.

Recall that all processes are wise under a symmetric-trust assumption. Protocols for asymmetric quorums cannot guarantee the same properties for naïve processes as for wise ones.

A useful notion for ensuring liveness and consistency for protocols is that of a *guild*. This is a set of wise processes that contains at least one quorum for each member.

**Definition 12 (Guild).** Given a fail-prone system  $\mathbb{F}$ , an asymmetric quorum system  $\mathbb{Q}$  for  $\mathbb{F}$ , and a protocol execution with faulty processes  $F$ , a *guild*  $\mathcal{G}$  for  $F$  satisfies two properties:

**Wisdom:**  $\mathcal{G}$  consists of wise processes,

$$\forall p_i \in \mathcal{G} : F \in \mathcal{F}_i^*.$$

**Closure:**  $\mathcal{G}$  contains a quorum for each of its members,

$$\forall p_i \in \mathcal{G}, \exists Q_i \in \mathbb{Q}_i : Q_i \subseteq \mathcal{G}.$$

The following lemma shows that every two guilds intersect.

**Lemma 4.** In any execution with a guild  $\mathcal{G}$ , there cannot exist two disjoint guilds.

*Proof.* Let  $\mathcal{P}$  be a set of processes,  $\mathcal{G}$  be a guild and  $F$  be the set of actually faulty processes. Furthermore, suppose that there is another guild  $\mathcal{G}'$ , with  $\mathcal{G} \cap \mathcal{G}' = \emptyset$ . Let  $p_i \in \mathcal{G}$  and  $p_j \in \mathcal{G}'$  be two processes and consider a quorum  $Q_i \subseteq \mathcal{G}$  for  $p_i$  and a quorum  $Q_j \subseteq \mathcal{G}'$  for  $p_j$ . From the definition of an asymmetric quorum system it must hold  $Q_i \cap Q_j \not\subseteq F$ , with  $Q_i \cap Q_j \neq \emptyset$  and  $F \in \mathcal{F}_i^* \cap \mathcal{F}_j^*$ . It follows that there exists a wise process  $p_k \in Q_i \cap Q_j$  with  $p_k \in \mathcal{G}$  and  $p_k \in \mathcal{G}'$ . Notice also that  $\mathcal{G}$  and  $\mathcal{G}'$  both contain a quorum for  $p_k$ .  $\square$

Observe that the union of two guilds is again a guild. It follows that every execution with a guild contains a unique *maximal guild*  $\mathcal{G}_{\max}$ . Analogously to the other asymmetric notions, for a given asymmetric fail-prone system, we call the list of canonical quorum systems of all processes an *asymmetric canonical quorum system*.

The following lemma shows that if a guild exists, then there cannot be a quorum for any process  $p_j$  containing only faulty processes.

**Lemma 5.** Let  $\mathcal{G}_{\max}$  be the maximal guild for a given execution and let  $\mathbb{Q}$  be the canonical asymmetric quorum system. Then, there cannot be a quorum  $Q_j \in \mathcal{Q}_j$  for any process  $p_j$  consisting only of faulty processes.

*Proof.* Given an execution with  $F$  as set of faulty processes, suppose there is a guild  $\mathcal{G}_{\max}$ . This means that for every process  $p_i \in \mathcal{G}_{\max}$ , a quorum  $Q_i \subseteq \mathcal{G}_{\max}$  exists such that  $Q_i \cap F = \emptyset$ . It follows that for every  $p_i \in \mathcal{G}_{\max}$ , there is a set  $F_i \in \mathcal{F}_i$  such that  $F \subseteq F_i$ . Recall that since  $\mathbb{Q}$  is a quorum system,  $B^3(\mathbb{F})$  holds. From Definition 10, we have that for all  $i, j \in [1, n]$ ,  $\forall F_i \in \mathcal{F}_i, \forall F_j \in \mathcal{F}_j, \forall F_{ij} \in \mathcal{F}_i^* \cap \mathcal{F}_j^* : \mathcal{P} \not\subseteq F_i \cup F_j \cup F_{ij}$ .

Towards a contradiction, assume that there is a process  $p_j$  such that there exists a quorum  $Q_j \in \mathcal{Q}_j$  for  $p_j$  with  $Q_j = F$ . This implies that there exists  $F_j \in \mathcal{F}_j$  such that  $F_j = \mathcal{P} \setminus F$ .

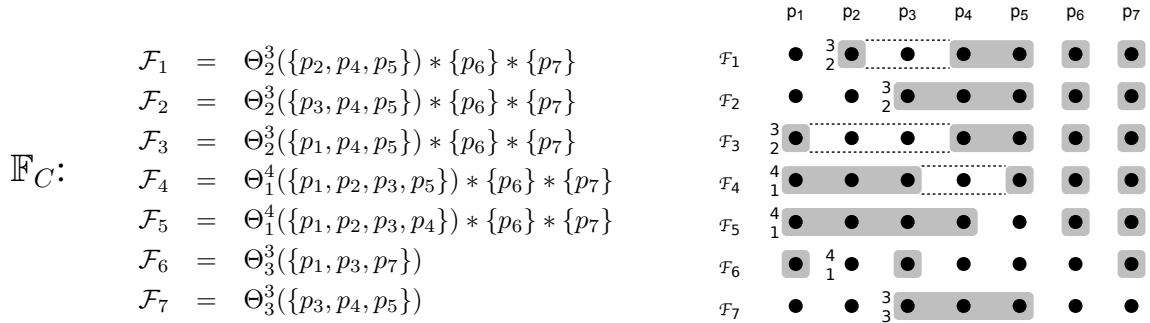
Let  $F_i$  be the fail-prone system of  $p_i \in \mathcal{G}_{\max}$  such that  $F \subseteq F_i$  and let  $F_j = \mathcal{P} \setminus F$  as just defined. Then,  $F_i \cup F_j \cup F_{ij} = \mathcal{P}$ . This follows from the fact that  $F_i$  contains  $F$  and  $F_j = \mathcal{P} \setminus F$ . This contradicts the  $B^3$ -condition for  $\mathbb{F}$ .  $\square$

**Lemma 6.** Let  $\mathcal{G}_{\max}$  be the maximal guild for a given execution and let  $p_i$  be any correct process. Then, every quorum for  $p_i$  contains at least one process from the maximal guild.

*Proof.* The proof naturally derives from the consistency property of an asymmetric quorum system. Consider any correct process  $p_i$  and one of its quorums,  $Q_i \in \mathcal{Q}_i$ . For any process  $p_j \in \mathcal{G}_{\max}$ , let  $Q_j$  be a quorum of  $p_j$  such that  $Q_j \subseteq \mathcal{G}_{\max}$ , which exists because  $\mathcal{G}_{\max}$  is a guild. Then, the quorum consistency property implies that  $Q_i \cap Q_j \neq \emptyset$ . Thus,  $Q_i$  contains a (correct) process in the maximal guild.  $\square$

Finally, we show with an example that it is possible for a wise process to be outside the maximal guild.

**Example 1.** Let us consider a seven-process asymmetric quorum system  $\mathbb{Q}_C$ , defined through its fail-prone system  $\mathbb{F}_C$ . The notation  $\Theta_k^n(S)$  for a set  $S$  with  $n$  elements denotes the *threshold* combination operator and enumerates all subsets of  $S$  of cardinality  $k$ . The diagram below shows fail-prone sets as shaded areas and the notation  $\binom{n}{k}$  in front of a fail-prone set stands for  $k$  out of the  $n$  processes in the set. The operator  $*$  for two sets satisfies  $A * B = \{A \cup B : A \in \mathcal{A}, B \in \mathcal{B}\}$ .



One can verify that  $B^3(\mathbb{F}_C)$  holds; hence, let  $\mathbb{Q}_C$  be the canonical quorum system of  $\mathbb{F}_C$ . With  $F = \{p_4, p_5\}$ , for instance, processes  $p_1, p_2, p_3$  and  $p_7$  are wise,  $p_6$  is naïve, and the maximal guild is  $\mathcal{G}_{\max} = \{p_1, p_2, p_3\}$ . It follows that process  $p_7$  is wise but outside the guild  $\mathcal{G}_{\max}$ , because quorum  $Q_7 \in \mathcal{Q}_7$  contains the naïve process  $p_6$ .

$$\begin{aligned}
\mathbb{Q}_C: \quad Q_1 &= \{\{p_1, p_3, p_5\}, \{p_1, p_3, p_4\}, \{p_1, p_2, p_3\}\} \\
Q_2 &= \{\{p_1, p_2, p_5\}, \{p_1, p_2, p_4\}, \{p_1, p_2, p_3\}\} \\
Q_3 &= \{\{p_2, p_3, p_5\}, \{p_2, p_3, p_4\}, \{p_1, p_2, p_3\}\} \\
Q_4 &= \{\{p_1, p_2, p_3, p_4\}, \{p_1, p_2, p_4, p_5\}, \{p_1, p_3, p_4, p_5\}, \{p_2, p_3, p_4, p_5\}\} \\
Q_5 &= \{\{p_1, p_2, p_3, p_5\}, \{p_1, p_2, p_4, p_5\}, \{p_1, p_3, p_4, p_5\}, \{p_2, p_3, p_4, p_5\}\} \\
Q_6 &= \{\{p_2, p_4, p_5, p_6\}\} \\
Q_7 &= \{\{p_1, p_2, p_6, p_7\}\}
\end{aligned}$$

## 6 Asymmetric randomized Byzantine consensus

In this section we define asymmetric Byzantine consensus. Then we implement it by a randomized algorithm, which is based on the protocol of Mostéfaoui *et al.* [18] as introduced in Section 4. Our implementation also fixes the problem described there.

Our notion of Byzantine consensus uses strong validity in the asymmetric model. Furthermore, it restricts the safety properties of consensus from all correct ones to *wise* processes. For implementing asynchronous consensus, we use a system enriched with randomization. In the asymmetric model, the corresponding probabilistic termination property is guaranteed only for processes in the maximal guild  $\mathcal{G}_{\max}$ .

**Definition 13 (Asymmetric strong Byzantine consensus).** A protocol for asynchronous *asymmetric strong Byzantine consensus* satisfies:

**Probabilistic termination:** In all executions with a guild, every process in the maximal guild decides with probability 1, in the sense that

$$\lim_{r \rightarrow +\infty} (\mathbb{P}[p_i \in \mathcal{G}_{\max} \text{ decides by round } r]) = 1.$$

**Strong validity:** In all executions with a guild, a wise process only decides a value that has been proposed by some processes in the maximal guild.

**Integrity:** No correct process decides twice.

**Agreement:** No two wise processes decide differently.

**Common coin.** Our randomized consensus algorithm delegates its probabilistic choices to a *common coin* abstraction [3,21]. We define this in the asymmetric-trust model.

**Definition 14 (Asymmetric common coin).** A protocol for *asymmetric common coin* satisfies the following properties:

**Termination:** In all executions with a guild, every process in the maximal guild eventually outputs a coin value.

**Unpredictability:** Unless at least one correct process has released the coin, no process has any information about the coin output by a wise process.

**Matching:** In all executions with a guild, with probability 1 every process in the maximal guild outputs the same coin value.

**No bias:** The distribution of the coin is uniform over  $\mathcal{B}$ .

An asymmetric common coin has an output domain  $\mathcal{B}$ . Here we consider binary consensus and  $\mathcal{B} = \{0, 1\}$ . The *termination* property guarantees that every process in the maximal guild eventually output a coin value that is ensured to be the same for each of them by the *matching* property. The *unpredictability* property ensures that the coin value is kept secret until a correct process releases the coin. Finally, the *no bias* property specifies the probability distribution of the coin output. The bias and matching properties may be weakened using well-known methods.

**Implementing an asymmetric common coin.** Our implementation of asymmetric common coin relies on the scheme of Benaloh and Leichter [1] and is shown in Algorithm 2. Furthermore, following the approach started by Rabin [21], we assume that coins are *predistributed* by an ideal dealer using secret sharing, in a way that for every round  $r$  there is exactly one coin with value  $s \in \{0, 1\}$ . Specifically, given an asymmetric quorum system  $\mathcal{Q}$ , the dealer creates random shares  $s_{i_1}, \dots, s_{i_{m-1}}$  for one random coin value  $s \in \{0, 1\}$  per round and for each  $Q_i = \{p_{i_1}, \dots, p_{i_m}\} \in \mathcal{Q}_i$  of  $\mathcal{Q}$  with  $|Q_i| = m$ . Then the dealer sets  $s_{i_m} = s + \sum_{j=1}^{m-1} s_{i_j} \pmod{2}$  and gives the shares to every process in  $Q_i$ . This ensures that any quorum can reconstruct the secret by computing the sum modulo 2 of the shares. This procedure is done for each quorum in every quorum system of  $\mathcal{Q}$ . Furthermore, the dealer authenticates the shares, preventing Byzantine processes to send inconsistent bits to other processes in the quorum (omitted from Algorithm 2). Correctness of this protocol follows easily. Given a quorum of correct processes, every wise process can thus reconstruct the secret and all of them will output the same coin value. On the other hand, in every execution with a guild, Byzantine processes cannot recover the secret without receiving at least one share from a correct process by Lemma 5.

This implementation is expensive because the number of shares for one particular coin held by a process  $p_i$  is equal to the number of quorums in which  $p_i$  is contained. In practical systems, one may also implement an asymmetric coin “from scratch” according to the direction taken by Canetti and Rabin [6] or recently by Patra *et*

---

**Algorithm 2** Asymmetric common coin for round  $round$  (code for  $p_i$ )

---

```
1: State
2:    $coin[k] \leftarrow [\perp]^n$ : for  $k \in [1, |\mathcal{Q}_i|]$ ,  $coin[k][j]$  holds the share received from  $p_j$  for quorum  $Q_{i,k}$  of  $p_i$ 
3:
4: upon event release-coin do
5:   for all  $p_j \in \mathcal{P}$  do                                     // send all shares to all process  $p_j$ 
6:     for all  $Q_{j,k} \in \mathcal{Q}_j$  do
7:       if  $p_i \in Q_{j,k}$  then
8:         send message  $[COIN, s_i, Q_{j,k}, r]$  to  $p_j$            //  $s_i$  is share of  $p_i$  for quorum  $Q_{j,k}$  of  $p_j$ 
9:
10: upon receiving a message  $[COIN, s_j, Q_{i,k}, r]$  from  $p_j$  such that  $r = round \wedge Q_{i,k} \in \mathcal{Q}_i \wedge p_j \in Q_{i,k}$  do
11:   if  $coin[k][j] = \perp$  then
12:      $coin[k][j] \leftarrow s_j$ 
13:
14: upon exists  $k$  such that  $Q_i^* = \{p_j \in \mathcal{P} \mid coin[k][j] \neq \perp\} \in \mathcal{Q}_i$  do                               // a quorum for  $p_i$ 
15:    $s \leftarrow \sum_{p_j \in Q_i^*} coin[k][j]$ 
16:   output output-coin( $s$ )
```

---

al. [20]. Alternatively, distributed cryptographic implementations appear to be possible, for example, as introduced by Cachin *et al.* [4].

## 6.1 Asymmetric binary validated broadcast

We generalize the binary validated broadcast as introduced in Section 4 to the asymmetric-trust model. All safety properties are restricted to wise processes, and a guild is required for liveness. Recall that every process may broadcast a binary value  $b \in \{0, 1\}$  by invoking *abv-broadcast*( $b$ ). The broadcast primitive outputs at least one value  $b$  and possibly also both binary values through an *bv-deliver*( $b$ ) event, according to the following notion.

**Definition 15 (Asymmetric binary validated broadcast).** A protocol for *asymmetric binary validated broadcast* satisfies the following properties:

**Validity:** In all executions with a guild, let  $K_i$  be a kernel for a process  $p_i$  in the maximal guild. If every process in  $K_i$  is correct and has *abv-broadcast* the same value  $b \in \{0, 1\}$ , then every wise process eventually *abv-delivers*  $b$ .

**Integrity:** In all executions with a guild, if a wise process *abv-delivers* some  $b$ , then  $b$  has been *abv-broadcast* by some process in the maximal guild.

**Agreement:** If a wise process *abv-delivers* some value  $b$ , then every wise process eventually *abv-delivers*  $b$ .

**Termination:** Every wise process eventually *abv-delivers* some value.

Notice that it guarantees properties only for processes that are wise or even in the maximal guild. Liveness properties also assume there exists a guild.

Algorithm 3 works in the same way as the binary validated broadcast by Mostéfaoui *et al.* [18] (Section 4), but differs in the use of asymmetric quorums. The condition of receiving VALUE messages containing  $b$  from at least  $f + 1$  processes is replaced by receiving such messages from a kernel  $K_i$  for process  $p_i$ . Furthermore, a quorum  $Q_i$  for  $p_i$  is needed instead of  $2f + 1$  messages for *abv-delivering* a bit.

**Theorem 7.** *Algorithm 3 implements asymmetric binary validated broadcast.*

*Proof.* To prove the *validity* property, let us consider a kernel  $K_i$  of correct processes for a process  $p_i$  in the maximal guild  $\mathcal{G}_{\max}$ . Observe that from the consistency property of asymmetric quorum systems, there cannot be two quorums  $Q_i \in \mathcal{Q}_i$  and  $Q_j \in \mathcal{Q}_j$  for  $p_i$  and  $p_j$ , respectively, such that every correct process in  $Q_i$  has *abv-broadcast*  $b$  and every correct process in  $Q_j$  has *abv-broadcast*  $b' \neq b$ . Let us assume w.l.o.g. that every quorum  $Q_i$  for a process  $p_i \in \mathcal{G}_{\max}$  contains some correct processes that have *abv-broadcast*  $b$ . Then, the set  $K$  containing only the correct processes that have *abv-broadcast*  $b$  intersects every  $Q_i$  for  $p_i$  and is a kernel for  $p_i$ . According to the protocol,  $p_i$  therefore sends  $[VALUE, b]$  unless *sentvalue*[ $b$ ] = TRUE. However, if *sentvalue*[ $b$ ] = TRUE,  $p_i$  has already sent  $[VALUE, b]$ . Hence, every process in  $\mathcal{G}_{\max}$  eventually sends  $[VALUE, b]$ . All these processes together therefore are a kernel for every correct process by Lemma 6. Let  $p_c$  be a correct process; process  $p_c$  therefore

---

**Algorithm 3** Asymmetric binary validated broadcast (code for  $p_i$ )

---

```
1: State
2:    $sentvalue \leftarrow [FALSE]^2$ :  $sentvalue[b]$  indicates whether  $p_i$  has sent  $[VALUE, b]$ 
3:    $values \leftarrow [\emptyset]^n$ : list of sets of received binary values
4:
5: upon event  $abv\text{-broadcast}(b)$  do
6:    $sentvalue[b] \leftarrow TRUE$ 
7:   send message  $[VALUE, b]$  to all  $p_j \in \mathcal{P}$ 
8:
9: upon receiving a message  $[VALUE, b]$  from  $p_j$  do
10:  if  $b \notin values[j]$  then
11:     $values[j] \leftarrow values[j] \cup \{b\}$ 
12:
13: upon exists  $b \in \{0, 1\}$  such that  $\{p_j \in \mathcal{P} \mid b \in values[j]\} \in \mathcal{K}_i$  and  $\neg sentvalue[b]$  do // a kernel for  $p_i$ 
14:    $sentvalue[b] \leftarrow TRUE$ 
15:   send message  $[VALUE, b]$  to all  $p_j \in \mathcal{P}$ 
16:
17: upon exists  $b \in \{0, 1\}$  such that  $\{p_j \in \mathcal{P} \mid b \in values[j]\} \in \mathcal{Q}_i$  do // a quorum for  $p_i$ 
18:  output  $abv\text{-deliver}(b)$ 
```

---

sends  $[VALUE, b]$  unless  $sentvalue[b] = TRUE$ . Then, every wise process receives a quorum for itself of values  $b$  and  $abv\text{-delivers}$   $b$ . Observe that validity property requires a kernel for a process in the maximal guild, otherwise one cannot obtain the same guarantees. Recall Example 1, for instance:  $\{p_7\}$  is a kernel for the wise process  $p_7$  but not for any other wise process. If one would relax the requirement and permit that  $K_i$  is a kernel for a wise process *outside* the maximal guild, validity could not be ensured.

For the *integrity* property, let us assume an execution with a maximal guild  $\mathcal{G}_{\max}$ . Suppose first that only Byzantine processes  $abv\text{-broadcast}$   $b$ . Then, the set consisting of only these processes cannot be a kernel for any wise process. It follows that line 13 of Algorithm 3 cannot be satisfied. If only naive processes  $abv\text{-broadcast}$   $b$ , by the definition of a quorum system and by the assumed existence of a maximal guild, there is at least one quorum for every process in  $\mathcal{G}_{\max}$  that does not contain any naive processes (e.g., as in Example 1). All naive processes together cannot be a kernel for processes in  $\mathcal{G}_{\max}$ . Again, line 13 of Algorithm 3 cannot be satisfied. Finally, let us assume that a wise process  $p_i$  outside the maximal guild  $abv\text{-broadcasts}$   $b$ . Then,  $p_i$  cannot be a kernel for every wise process: it is not part of the quorums inside  $\mathcal{G}_{\max}$ . It follows that if a wise process  $abv\text{-delivers}$  some  $b$ , then  $b$  has been  $abv\text{-broadcast}$  by some processes in the maximal guild.

To show *agreement*, suppose that a wise process  $p_i$  has  $abv\text{-delivered}$   $b$ . Then it has obtained  $[VALUE, b]$  messages from the processes in some quorum  $Q_i \in \mathcal{Q}_i$  and before from a kernel  $K_i$  for itself. Each correct process in  $K_i$  has sent  $[VALUE, b]$  message to all other processes. Consider any other wise process  $p_j$ . Since  $p_i$  and  $p_j$  are both wise, we have  $F \in \mathcal{F}_i^*$  and  $F \in \mathcal{F}_j^*$ , which implies  $F \in \mathcal{F}_i^* \cap \mathcal{F}_j^*$ . It follows that  $K_i$  is also a kernel for  $p_j$ . Thus,  $p_j$  sends a  $[VALUE, b]$  message to every process. This implies that all wise processes eventually receive a quorum for themselves of  $[VALUE, b]$  messages and  $abv\text{-deliver}$   $b$ .

For the *termination* property, notice that in any execution, every correct process  $abv\text{-broadcasts}$  some binary values. Termination then follows from the validity property.  $\square$

## 6.2 Asymmetric randomized consensus

In the following primitive, a correct process may *propose* a binary value  $b$  by invoking  $arbc\text{-propose}(b)$ ; the consensus abstraction *decides* for  $b$  through an  $arbc\text{-decide}(b)$  event.

Algorithm 4 differs from Algorithm 1 in some aspects. Recall that both algorithms use a system where messages are authenticated and delivered reliably. Importantly, we assume for Algorithm 4 that all messages among correct processes are also delivered in FIFO order, even when they do not originate from the same protocol module.

Moreover, Algorithm 4 allows the set  $B$  to change while reconstructing the common coin (line 33). This step is necessary in order to prevent the problem described in Section 4. We prove this statement in Lemma 8.

Finally, our protocol may disseminate DECIDE messages in parallel to ensure termination. When  $p_i$  receives a DECIDE message from a kernel of processes for itself containing the same value  $b$ , then it broadcasts a DECIDE message itself containing  $b$  to every processes, unless it has already done so. Once  $p_i$  receives a DECIDE message from a quorum of processes for itself with the same value  $b$ , it  $arbc\text{-decides}(b)$  and halts. This “amplification” step

---

**Algorithm 4** Asymmetric randomized binary consensus (code for  $p_i$ ).

---

```
1: State
2:    $round \leftarrow 0$ : current round
3:    $values \leftarrow \{\}$ : set of abv-delivered binary values for the round
4:    $aux \leftarrow [\{\}]^n$ : stores sets of values that have been received in AUX messages in the round
5:    $decided \leftarrow [\perp]^n$ : stores binary values that have been reported as decided by other processes
6:    $sentdecide \leftarrow \text{FALSE}$ : indicates whether  $p_i$  has sent a DECIDE message
7:
8: upon event arbc-propose( $b$ ) do
9:   invoke abv-broadcast( $b$ ) with tag  $round$ 
10:
11: upon abv-deliver( $b$ ) with tag  $r$  such that  $r = round$  do
12:    $values \leftarrow values \cup \{b\}$ 
13:   send message [AUX,  $round$ ,  $b$ ] to all  $p_j \in \mathcal{P}$ 
14:
15: upon receiving a message [AUX,  $r$ ,  $b$ ] from  $p_j$  such that  $r = round$  do
16:    $aux[j] \leftarrow aux[j] \cup \{b\}$ 
17:
18: upon receiving a message [DECIDE,  $b$ ] from  $p_j$  such that  $decided[j] = \perp$  do
19:    $decided[j] = b$ 
20:
21: upon exists  $b \neq \perp$  such that  $\{p_j \in \mathcal{P} \mid decided[j] = b\} \in \mathcal{K}_i$  do // a kernel for  $p_i$ 
22:   if  $\neg sentdecide$  then
23:     send message [DECIDE,  $b$ ] to all  $p_j \in \mathcal{P}$ 
24:      $sentdecide \leftarrow \text{TRUE}$ 
25:
26: upon exists  $b \neq \perp$  such that  $\{p_j \in \mathcal{P} \mid decided[j] = b\} \in \mathcal{Q}_i$  do // a quorum for  $p_i$ 
27:   arbc-decide( $b$ )
28:   halt
29:
30: upon exist  $\{p_j \in \mathcal{P} \mid aux[j] \subseteq values\} \in \mathcal{Q}_i$  do // a quorum for  $p_i$ 
31:   release-coin with tag  $round$ 
32:
33: upon event output-coin( $s$ ) with tag  $round$  and exists  $B \neq \{\}$  such that  $\forall p_j \in \mathcal{Q}_i, B = aux[j]$  do
34:    $round \leftarrow round + 1$ 
35:   if exists  $b$  such that  $|B| = 1 \wedge B = \{b\}$  then
36:     if  $b = s \wedge \neg sentdecide$  then
37:       send message [DECIDE,  $b$ ] to all  $p_j \in \mathcal{P}$ 
38:        $sentdecide \leftarrow \text{TRUE}$ 
39:       invoke abv-broadcast( $b$ ) with tag  $round$  // propose  $b$  for the next round
40:   else
41:     invoke abv-broadcast( $s$ ) with tag  $round$  // propose coin value  $s$  for the next round
42:    $values \leftarrow [\perp]^n$ 
43:    $aux \leftarrow [\{\}]^n$ 
```

---

is reminiscent of Bracha’s reliable broadcast protocol [2]. Hence, the protocol does not execute rounds forever, in contrast to the original formulation of Mostéfaoui *et al.* [18] (Algorithm 1).

The following result shows that the problem described in Section 4 no longer occurs in our protocol.

**Lemma 8.** *If a wise process  $p_i$  outputs the coin with  $B = \{0, 1\}$ , then every other wise process that outputs the coin has also  $B = \{0, 1\}$ .*

*Proof.* Let us assume that a wise process  $p_i$  outputs the coin with  $B = \{0, 1\}$ . This means that  $p_i$  has received COIN messages from a quorum  $Q_i^{\text{COIN}}$  for itself and has  $B = \text{aux}[j] = \{0, 1\}$  for all  $p_j \in Q_i$ . Consider another wise process  $p_j$  that also outputs the coin. It follows that  $p_j$  has received COIN messages from a quorum  $Q_j^{\text{COIN}}$  for itself as well. Observe that  $p_i$  and  $p_j$ , before receiving the COIN messages from every process in  $Q_i^{\text{COIN}}$  and  $Q_j^{\text{COIN}}$ , respectively, receive all AUX messages that the correct processes in these quorums have sent before the COIN messages. This follows from the assumption of FIFO reliable point-to-point links. Quorum consistency implies that  $Q_i^{\text{COIN}}$  and  $Q_j^{\text{COIN}}$  have some correct processes in common. So,  $p_i$  and  $p_j$  receive some AUX messages from the same correct process before they may output the coin. This means that if  $p_i$  has  $B = \{0, 1\}$  after the *output-coin* event, then every quorum  $Q_j$  for  $p_j$  will contain a process  $p_k$  such that  $\text{aux}[k] = \{0, 1\}$  for  $p_j$ . Every wise process therefore must have  $B = \{0, 1\}$  before it can proceed.  $\square$

The problem shown in Section 4 arose from messages between correct processes that were reordered in such a way that knowledge of the common coin value  $s$  was able to influence another correct process and cause it to deliver  $\neg s$  alone. Lemma 8 above implies that our Algorithm 4 prevents this because all wise processes arrive at the same set  $B$  when they output the coin.

**Theorem 9.** *Algorithm 4 implements asymmetric strong Byzantine consensus.*

*Proof.* To prove the *strong validity* property, assume that a wise process  $p_i$  has decided a value  $b$  in round  $r$ . This means that  $p_i$  has received  $[\text{DECIDE}, b]$  messages from a quorum  $Q_i$  for itself. Furthermore, this also means that  $B = \{b\}$  and  $b$  is the same as the coin value in round  $r$ . Then,  $p_i$  has received  $b$  from a quorum  $Q_i$  for itself. Every process in  $Q_i$  has received a  $[\text{AUX}, r, b]$  message and  $b \in \text{values}$  has been *abv-delivered* from *abv-broadcast* instance. From the integrity property of *abv-broadcast* instance,  $b$  has been *abv-broadcast* by a process in the maximal guild and, specifically, *values* contains only values *abv-broadcast* by processes in the maximal guild. It follows that  $b$  has been proposed by some processes in the maximal guild.

For the *agreement* property, suppose that a wise process has received  $[\text{AUX}, r, b]$  messages from a quorum  $Q_i$  for itself. Consider any other wise process  $p_j$  that has received a quorum  $Q_j$  for itself of  $[\text{AUX}, r, \bar{b}]$  messages. If at the end of round  $r$  there is only one value in  $B$ , then from consistency property of quorum systems, it follows  $b = \bar{b}$ . Furthermore, if  $b = s$  then  $p_i$  and  $p_j$  broadcast a  $[\text{DECIDE}, b]$  message to every process and decide for  $b$  after receiving a quorum of  $[\text{DECIDE}, b]$  messages for themselves, otherwise they both *abv-broadcast*( $b$ ) and they continue to *abv-broadcast*( $b$ ) until  $b = s$ . If  $B$  contains more than one value, then  $p_i$  and  $p_j$  proceed to the next round and invoke a new instance of *abv-broadcast* with  $s$ . Therefore, at the beginning of the next round, the proposed values of all wise processes are equal. The property easily follows.

The *integrty* property is easily derived from the algorithm.

The *probabilistic termination* property follows from two observations. First, from the termination and the agreement properties of *abv-broadcast* it follows that every wise process *abv-delivers* the same binary value from the *abv-broadcast* instance and this value has been *abv-broadcast* by some processes in the maximal guild. Second, we show that with probability 1, there exists a round at the end of which all processes in  $\mathcal{G}_{\max}$  have the same proposal  $b$ . If at the end of round  $r$ , every process in  $\mathcal{G}_{\max}$  has proposed the coin value (line 41, Algorithm 4), then all of them start the next round with the same value. Similarly, if every process in  $\mathcal{G}_{\max}$  has executed line 39 (Algorithm 4) they adopt the value  $b$  and start the next round with the same value.

However, it could be the case that some wise processes in the maximal guild proposes  $b$  and another one proposes the coin output  $s$ . Observe that the properties of the common coin abstraction guarantee that the coin value is random and independently chosen. So, the random value  $s$  is equal to the proposal value  $b$  with probability  $\frac{1}{2}$ . The probability that there exists a round  $r'$  in which the coin equals the value  $b$  proposed by all processes in  $\mathcal{G}_{\max}$  during round  $r'$  approaches 1 when  $r$  goes to infinity.

Let  $r$  thus be some round in which every process in  $\mathcal{G}_{\max}$  *abv-broadcasts* the same bit  $b$ ; then, none of them will ever change their proposal again. This is due to the fact that every wise process invokes an *abv-broadcast* instance with the same proposal  $b$ . According to the validity and agreement properties of asymmetric binary validated broadcast, every process in the maximal guild then delivers the same, unique value  $b$ . Hence, the proposal of every process in the maximal guild is set to  $b$  and never changes. Finally, the properties of common coin guarantee that with probability 1 the processes eventually reach a round in which the coin outputs  $b$ . Therefore, with probability 1 every process in the maximal guild sends a DECIDE message with value  $b$  to every process. This implies that



it exists a quorum  $Q_i \subseteq \mathcal{G}_{\max}$  for a process  $p_i \in \mathcal{G}_{\max}$  such that every process in  $Q_i$  has sent a DECIDE message with value  $b$  to every process. Moreover, the set  $K_i = Q_i \setminus F$  is a kernel for  $p_i$  and for every other wise process  $p_j$ . If  $p_j \in \mathcal{G}_{\max}$  receives a DECIDE message with value  $b$  from a kernel for itself, it sends a DECIDE message with value  $b$  to every process unless it has already done so. It follows that every process in the maximal guild receives DECIDE messages with the value  $b$  from a quorum for itself and *arbc-decides* for  $b$ .  $\square$

## 7 Conclusion

As we show in this work, consensus protocols with asymmetric trust can be obtained by starting from existing, well-known protocols with symmetric trust. But we have no assurance that this is also the case for further traditional algorithms. Understanding subjective trust and implementing the corresponding protocols remains an interesting open problem, especially with respect to cryptographic constructions.

Moreover, while it has been shown that the existence of asymmetric quorums is characterized by the  $B^3$ -condition, it remains open how such a structure might arise spontaneously in a dynamic system, where processes join and leave without knowledge of each other. One particular issue to consider is a Byzantine process that declares a fail-prone system with the sole aim of sabotaging the  $B^3$ -condition. Achieving a complete characterization of asymmetric quorum systems and the corresponding protocols would allow to have systems with open and dynamic membership, in which the participants do not need to know each other from the start but still benefit from strong consistency guarantees.

## Acknowledgments

The authors thank Orestis Alpos and Giorgia Azzurra Marson for interesting discussions.

This work has been funded by the Swiss National Science Foundation (SNSF) under grant agreement Nr. 200021\_188443 (Advanced Consensus Protocols).

## References

- [1] J. C. Benaloh and J. Leichter, “Generalized secret sharing and monotone functions,” in *Advances in Cryptology - CRYPTO '88, 8th Annual International Cryptology Conference, Santa Barbara, California, USA, August 21-25, 1988, Proceedings* (S. Goldwasser, ed.), vol. 403 of *Lecture Notes in Computer Science*, pp. 27–35, Springer, 1988.
- [2] G. Bracha, “Asynchronous byzantine agreement protocols,” *Inf. Comput.*, vol. 75, no. 2, pp. 130–143, 1987.
- [3] C. Cachin, R. Guerraoui, and L. E. T. Rodrigues, *Introduction to Reliable and Secure Distributed Programming (2. ed.)*. Springer, 2011.
- [4] C. Cachin, K. Kursawe, and V. Shoup, “Random oracles in constantinople: Practical asynchronous byzantine agreement using cryptography,” *J. Cryptology*, vol. 18, no. 3, pp. 219–246, 2005.
- [5] C. Cachin and B. Tackmann, “Asymmetric distributed trust,” in *23rd International Conference on Principles of Distributed Systems, OPODIS 2019, December 17-19, 2019, Neuchâtel, Switzerland* (P. Felber, R. Friedman, S. Gilbert, and A. Miller, eds.), vol. 153 of *LIPIcs*, pp. 7:1–7:16, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [6] R. Canetti and T. Rabin, “Fast asynchronous byzantine agreement with optimal resilience,” in *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA* (S. R. Kosaraju, D. S. Johnson, and A. Aggarwal, eds.), pp. 42–51, ACM, 1993.
- [7] I. Damgård, Y. Desmedt, M. Fitzi, and J. B. Nielsen, “Secure protocols with asymmetric trust,” in *Advances in Cryptology - ASIACRYPT 2007, 13th International Conference on the Theory and Application of Cryptology and Information Security, Kuching, Malaysia, December 2-6, 2007, Proceedings* (K. Kurosawa, ed.), vol. 4833 of *Lecture Notes in Computer Science*, pp. 357–375, Springer, 2007.
- [8] Á. García-Pérez and A. Gotsman, “Federated byzantine quorum systems,” in *22nd International Conference on Principles of Distributed Systems, OPODIS 2018, December 17-19, 2018, Hong Kong, China* (J. Cao, F. Ellen, L. Rodrigues, and B. Ferreira, eds.), vol. 125 of *LIPIcs*, pp. 17:1–17:16, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

- [9] V. Hadzilacos and S. Toueg, “Fault-tolerant broadcasts and related problems,” in *Distributed Systems (2nd Ed.)* (S. J. Mullender, ed.), pp. 97–145, New York: ACM Press/Addison-Wesley Publishing Co., 1993.
- [10] M. Hirt and U. M. Maurer, “Player simulation and general adversary structures in perfect multiparty computation,” *J. Cryptology*, vol. 13, no. 1, pp. 31–60, 2000.
- [11] F. P. Junqueira and K. Marzullo, “Synchronous consensus for dependent process failure,” in *23rd International Conference on Distributed Computing Systems (ICDCS 2003), 19-22 May 2003, Providence, RI, USA*, pp. 274–283, IEEE Computer Society, 2003.
- [12] M. Lokhava, G. Losa, D. Mazières, G. Hoare, N. Barry, E. Gafni, J. Jove, R. Malinowsky, and J. McCaleb, “Fast and secure global payments with stellar,” in *Proceedings of the 27th ACM Symposium on Operating Systems Principles, SOSP 2019, Huntsville, ON, Canada, October 27-30, 2019* (T. Brecht and C. Williamson, eds.), pp. 80–96, ACM, 2019.
- [13] G. Losa, E. Gafni, and D. Mazières, “Stellar consensus by instantiation,” in *33rd International Symposium on Distributed Computing, DISC 2019, October 14-18, 2019, Budapest, Hungary* (J. Suomela, ed.), vol. 146 of *LIPICs*, pp. 27:1–27:15, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [14] D. Malkhi, K. Nayak, and L. Ren, “Flexible byzantine fault tolerance,” in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS 2019, London, UK, November 11-15, 2019* (L. Cavallaro, J. Kinder, X. Wang, and J. Katz, eds.), pp. 1041–1053, ACM, 2019.
- [15] D. Malkhi and M. K. Reiter, “Byzantine quorum systems,” *Distributed Comput.*, vol. 11, no. 4, pp. 203–213, 1998.
- [16] D. Mazières, “The Stellar consensus protocol: A federated model for Internet-level consensus.” Stellar, available online, <https://www.stellar.org/papers/stellar-consensus-protocol.pdf>, 2016.
- [17] A. Miller, Y. Xia, K. Croman, E. Shi, and D. Song, “The honey badger of BFT protocols,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016* (E. R. Weippl, S. Katzenbeisser, C. Kruegel, A. C. Myers, and S. Halevi, eds.), pp. 31–42, ACM, 2016.
- [18] A. Mostéfaoui, M. Hamouma, and M. Raynal, “Signature-free asynchronous byzantine consensus with  $t < n/3$  and  $o(n^2)$  messages,” in *ACM Symposium on Principles of Distributed Computing, PODC ’14, Paris, France, July 15-18, 2014* (M. M. Halldórsson and S. Dolev, eds.), pp. 2–9, ACM, 2014.
- [19] A. Mostéfaoui, H. Moumen, and M. Raynal, “Signature-free asynchronous binary byzantine consensus with  $t < n/3$ ,  $o(n^2)$  messages, and  $O(1)$  expected time,” *J. ACM*, vol. 62, no. 4, pp. 31:1–31:21, 2015.
- [20] A. Patra, A. Choudhury, and C. P. Rangan, “Asynchronous byzantine agreement with optimal resilience,” *Distributed Comput.*, vol. 27, no. 2, pp. 111–146, 2014.
- [21] M. O. Rabin, “Randomized byzantine generals,” in *Proceedings of the 24th Annual Symposium on Foundations of Computer Science, SFCS’83, (USA)*, pp. 403–409, IEEE Computer Society, 1983.
- [22] P. Tholoniati and V. Gramoli, “Formal verification of blockchain byzantine fault tolerance,” in *6th Workshop on Formal Reasoning in Distributed Algorithms (FRIDA’19)*, Oct 2019.