How to Trust Strangers: Composition of Byzantine Quorum Systems

Orestis Alpos
University of Bern
orestis.alpos@inf.unibe.ch

Christian Cachin
University of Bern
cachin@inf.unibe.ch

Luca Zanolini
University of Bern
luca.zanolini@inf.unibe.ch

Abstract

Trust is the basis of any distributed, fault-tolerant, or secure system. A *trust assumption* specifies the failures that a system, such as a blockchain network, can tolerate and determines the conditions under which it operates correctly. In systems subject to Byzantine faults, the trust assumption is usually specified through sets of processes that may fail together. Trust has traditionally been *symmetric*, such that all processes in the system adhere to the same, global assumption about potential faults. Recently, *asymmetric* trust models have also been considered, especially in the context of blockchains, where every participant is free to choose who to trust.

In both cases, it is an open question how to compose trust assumptions. Consider two or more systems, run by different and possibly disjoint sets of participants, with different assumptions about faults: how can they work together? This work answers this question for the first time and offers composition rules for symmetric and for asymmetric quorum systems. These rules are static and do not require interaction or agreement on the new trust assumption among the participants. Moreover, they ensure that if the original systems allow for running a particular protocol (guaranteeing consistency and availability), then so will the joint system. At the same time, the composed system tolerates as many faults as possible, subject to the underlying consistency and availability properties.

Reaching consensus with asymmetric trust in the model of personal Byzantine quorum systems (Losa *et al.*, DISC 2019) was shown to be impossible, if the trust assumptions of the processes diverge from each other. With asymmetric quorum systems, and by applying our composition rule, we show how consensus is actually possible, even with the combination of disjoint sets of processes.

1 Introduction

Secure distributed systems rely on *trust*. A security assumption defines the failures and attacks that can be tolerated and names conditions under which the system may operate. Implicitly, this determines the trust in certain components to be correct. In fault-tolerant replicated systems, trust has traditionally been expressed globally, through a *symmetric* assumption on the number or kind of faulty processes, which is shared by all processes. An example of this is the well-known threshold fault assumption: the system tolerates up to a finite and limited number of faulty processes in the system; no guarantees can be given beyond this about the correct execution of protocols. More generally, a symmetric trust assumption is defined through a *fail-prone system*, which is a collection of subsets of processes, such that each of them contains all the processes that may at most fail together during a protocol execution.

Quorum systems [18] complement the notion of fail-prone systems and are used within distributed fault-tolerant protocols to express trust assumptions operationally.

In the classical interpretation, a quorum system is a collection of subsets of processes, called *quorums*, with two properties, formally known as *consistency* and *availability*, respectively, that any two quorums have a non-empty intersection and that in every execution, there exists a quorum made of correct processes. *Byzantine quorum systems* (BQS) have been formalized by Malkhi and Reiter [14] and

generalize classical quorum systems by tolerating Byzantine failures, i.e., where faulty processes may behave arbitrarily. They are the focus of this work and allow for building secure, trustworthy systems. A BQS assumes one global shared Byzantine fail-prone system and, because of that, use the model of symmetric trust. Consistency for a BQS demands that any two quorums intersect in a set that contains at least one correct process in every execution.

Motivated by the requirements of more flexible trust models, particularly in the context of blockchain networks, new approaches to trust have been explored. It is evident that a common trust model cannot be imposed in an open and decentralized or permissionless environment. Instead, every participant in the system should be free to choose who to trust and who not to trust. Damgård *et al.* [7], and Cachin and Tackmann [4] extend Byzantine quorum systems to permit subjective trust by introducing *asymmetric* Byzantine quorum systems. They let every process specify their own fail-prone system and quorum system. Global system guarantees can be derived from these personal assumptions. Extending traditional Byzantine quorum systems that use threshold assumptions, several recent recent suggestions [10, 9, 13] have also introduced more flexible notions of trust.

In this paper, we study the problem of composing trust assumptions, as expressed by symmetric and by asymmetric Byzantine quorum systems. Starting from two or more running distributed systems, each one with its own assumption, how can they be combined, so that their participant groups are joined and operate together? A simple, but not so intriguing solution could be to stop all running protocols and to redefine the trust structure from scratch, with full knowledge of all assumptions across the participants. With symmetric trust, a new global assumption that includes all participants would be defined. In the asymmetric-trust model, every process would specify new personal assumptions on all other participants. Subsequently, the composite system would have to be restarted. Although this solution can be effective, it requires that all members of each initial group express assumptions about the trustworthiness of the processes in the other groups. In realistic scenarios, this might not be possible, since the participants of one system lack knowledge about the members of other systems, and can therefore not express their trust about them. Moreover, one needs to ensure that the combined system satisfies the liveness and safety conditions, as expressed by the B^3-condition for quorum intersection. Since the assumptions are personal, it is not guaranteed, and in practice quite challenging, that the composite system will indeed satisfy the B^3-condition.

This work formulates the problem of composing quorum systems and gives methods for assembling trust assumptions from different, possibly disjoint, systems to a common model. We do so by introducing composition rules for trust assumptions, in both the symmetric-trust and asymmetric-trust model. Our methods describe the resulting fail-prone systems and the corresponding quorum systems.

In a different line of work, subjective trust assumptions have also been introduced with the Stellar blockchain (`www.stellar.org`) [17, 11, 12], a cryptocurrency ranked in the top-20 by market capitalization today. In contrast to the original, well-understood notion of quorum systems, these works depart from the classical intersection requirement among quorums. Such systems may fork into separate *consensus clusters*, each one satisfying agreement and liveness on its own. This implies that consensus may hold only "locally", and a unique consensus across disjoint clusters is not possible. More specifically, Losa *et al.* prove [12, Lemma 4] that no quorum-based algorithm can guarantee agreement between two processes whose quorums do not intersect in their model. Our work overcomes this impossibility and shows that consensus can be reached even with disjoint sets of participants, whose trust assumptions do not intersect. Moreover, we use the established notion of quorums, which enables to run many well-understood protocols, such as consensus, reliable broadcast, emulations of shared memory, and more [4, 5].

Specifically, the contributions of this work are as follows:

1. We show how to join together two or more systems in a way where processes in one system do not need a complete knowledge of the trust assumptions of those in the other.

2. We allow processes in each system to maintain their trust assumptions within their original system.

3. We define a deterministic rule to extend the trust assumptions of each system by including the new

participants.

4. Our composition rules guarantee that *consistency* and *availability* will be satisfied in the composite quorum system.

Organization. The remainder of this work is structured as follows. In Section 2 we review related work. We present our system model and preliminaries on quorum systems with symmetric and asymmetric assumptions in Section 3. In Section 4 we focus on the symmetric-trust model and show different composition rules on both fail-prone systems and quorum systems. These rules achieve different properties, which we explore formally. A composition rule in the asymmetric model is presented in Section 5. For this proof, we make use of a deterministic method called *purification*, whose purpose is to streamline and improve the trust assumption of each participant in a system, making the composition between more systems possible. We then discuss the implications and the limits of this approach and offer ideas on how to implement our results. Finally, conclusions are drawn in Section 6.

2 Related work

Byzantine quorum systems (BQS) have originally been formalized by Malkhi and Reiter [14] to generalize classical quorum systems toward processes prone to Byzantine failures. They model symmetric trust, where every process in the system adheres to a global, common assumption. Many distributed protocols employ BQS as their foundation; in the area of state-machine replication, for example, they range from PBFT [6] to Tendermint [3], HotStuff [19], and other blockchain-specific protocols. Recently, also generalized BQS have been demonstrated for implementing consensus [1].

Measures of quality for classical (non-Byzantine) quorum system have been studied by Naor and Wool [18] in terms of the load, capacity, and availability properties. The load (the probability of access of the busiest process) and availability (probability of some quorum surviving failures) properties have then been considered by Malkhi *et al.* [15] in the context of the Byzantine quorum systems. They construct different types of Byzantine quorum systems with optimal load or availability.

Subsequent literature extends the BQS model, seeking to overcome some limitations and to take them into practice. To this end, *probabilistic* quorum systems have been introduced by Malkhi *et al.* [16] as a tool for ensuring consistency of replicated data with high probability despite both benign and Byzantine failure of processes. They define the ϵ-*intersecting quorum systems* by relaxing the intersection property of a quorum system in a way that every two quorums fail to intersect with some small probability ϵ. By the quality measures, these new quorums show an improvement over the classic and Byzantine ones.

Alvisi *et al.* [2] introduce *dynamic* Byzantine quorum systems in the context of quorum-based Byzantine fault-tolerant data services. They present protocols for dynamically changing the threshold of the system. In this this way, they solve an intrinsic limitation of standard Byzantine quorums, which is their dependence on *a-priori* defined resilience thresholds.

Malkhi *et al.* [13] define *flexible Byzantine quorums* that allow processes in the system to have different faults models. This work presents a new approach for designing Byzantine fault-tolerant consensus protocols which guarantees higher resilience by introducing a new *alive-but-corrupt* fault type, which denotes processes that attack safety but not liveness.

Recent work has explored frameworks that loosen the global model of trust, allowing processes to choose in a subjective way who to trust. Damgård *et al.* [7] define the basics of *asymmetric trust* for secure computation protocols. Under this model, processes are free to make their personal assumptions regarding other processes, resulting in a broader and richer trust structure, compared to the symmetric model. They introduce a wider class of correct processes, differentiated according to their trust choices. Moreover, they show protocols for synchronous broadcast, verifiable secret sharing, and other primitives. Properties of these protocol can be guaranteed only to a specific subset of correct processes.

Asymmetric Byzantine quorum systems have been introduced by Cachin and Tackmann [4] as a natural extension of symmetric Byzantine quorum systems [14] to the model with asymmetric trust. They

present protocols for asynchronous Byzantine consistent broadcast, reliable broadcast, and emulations of shared memory with asymmetric quorums. Their work gives rise to a new structure called a *guild*, which is a subset of processes that are called *wise* because they correctly anticipated the actual faults. Some protocol guarantees can only be ensured for wise processes or only for those in a guild. Asymmetric Byzantine consensus protocols have been described as well [5].

With the rise of blockchains, protocols using flexible trust structures have been deployed in practice as well. Ripple (www.ripple.com) and Stellar (www.stellar.org) do not base their resilience guarantees on a global threshold, but allow participants to express their own beliefs. However, their formalization is not a generalization of the BQS model. In this work, we explore compositions of symmetric and asymmetric BQS that are based on the well-studied notions.

A related form of recursive composition of (Byzantine) quorum systems has been explored and utilized in the literature. The idea is that, given two systems, each occurrence of a process in the first is *replaced* by a copy of the second system. Malkhi *et al.* [15] construct and study composite BQS, such as *recursive threshold* BQS, using this idea. Hirt and Maurer [8] use this technique to reason about multiparty computation over access structures. Our approach is orthogonal to these works, in the sense that it places the two original systems on the same level. In other words, we explore the failures that two systems can tolerate when they are joined together, as opposed when one is inserted into the other.

3 System model and preliminaries

3.1 System model

Processes. We consider a system $\mathcal{P}$ with an arbitrary number of *processes* p_i, also called *participants*, that communicate with each other. A protocol for $\mathcal{P}$ consists of a collection of programs with instructions for all processes.

Executions and faults. An *execution* starts with all processes in a special initial state; subsequently the processes repeatedly change their state through computation steps. Every execution is fair in the sense that, informally, processes do not halt prematurely when there are still steps to be taken.

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.

3.2 Preliminaries

We start by presenting definitions and main results in the symmetric-trust model. These will be used in the next section to construct and prove our composition rules.

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.

A complementary structure to the fail-prone system is given by a Byzantine quorum system [14], defined as follows.

Definition 2 (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 *Byzantine quorum system* (BQS) 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:

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

A link between the above two definition is given by the following results.

Definition 3 (Q^3-condition). 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

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

Lemma 4 (Symmetric quorum system existence [14]). *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 $\overline{\mathcal{F}}$, the bijective complement of $\mathcal{F}$, is a Byzantine quorum system called* canonical quorum system *of $\mathcal{F}$.*

Finally, we present the asymmetric-trust model as introduced by Damgård *et al.* [7] and Cachin and Tackmann [4].

Definition 5 (Asymmetric fail-prone system). An asymmetric fail-prone system $\mathbb{F} = [\mathcal{F}_1, \ldots, \mathcal{F}_n]$ consists of an array of fail-prone systems, where $\mathcal{F}_i \subseteq 2^{\mathcal{P}}$ denotes the trust assumption of p_i. We assume $p_i \notin \mathcal{F}_i$

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 $\mathcal{A}^* = \{A' | A' \subseteq A, A \in \mathcal{A}\}$ denote the collection of all subsets of the sets in $\mathcal{A}$.

Definition 6 (Asymmetric Byzantine quorum system). Let $\mathbb{F} = [\mathcal{F}_1, \ldots, \mathcal{F}_n]$ be an asymmetric fail-prone system. An *asymmetric Byzantine quorum system* (ABQS) for $\mathbb{F}$ is an array of collections of sets $\mathbb{Q} = [\mathcal{Q}_1, \ldots, \mathcal{Q}_n]$, 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: $\forall i \in [1, n]$

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

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

Definition 7 (B^3-condition [7, 4]). 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

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

Lemma 8 (Asymmetric quorum system existence [4]). *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 a given asymmetric fail-prone system, we call the list of canonical quorum systems of all processes an *asymmetric canonical quorum system*.

Given a protocol execution with asymmetric Byzantine quorum systems, where F is the actual failed set, 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 9 (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, i.e.,

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

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

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

Observe that the union of two guild is again a guild [5]. Every execution with a guild contains a unique *maximal guild*.

Lemma 10 ([5]). *Let $\mathcal{G}$ be the 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 guild.*

4 Composition of symmetric BQS

Given two Byzantine quorum systems $\mathcal{Q}_1$ defined on processes $\mathcal{P}_1$ with fail-prone system $\mathcal{F}_1$, and $\mathcal{Q}_2$ defined on processes $\mathcal{P}_2$ with fail-prone system $\mathcal{F}_2$, we want to provide a *composition* rule between the two that allows the resulting BQS $\mathcal{Q}_3$ defined on processes $\mathcal{P}_3 = \mathcal{P}_1 \cup \mathcal{P}_2$ with fail-prone system $\mathcal{F}_3$ to run a distributed protocol together. The resulting system should satisfy the consistency and availability properties of a BQS, that is, it should remain consistent and live in any execution where a fail-prone set in $\mathcal{F}_3$ fails.

In this work we explore the composition of two BQS as a means to allow them jointly run a protocol, *without* requiring the processes in one BQS to make new trust assumptions about the processes in the other. This is useful in practice because remodeling trust from scratch would be a manual and uncertain process. We do not consider the composition as a way to increase their resilience. For example, joining four singleton BQS will result in a system with four processes, none of which is expected to fail. This makes sense if one starts from the trust assumptions of singleton BQS; by definition, the single process it contains never fails. There could be other ways to compose the BQS, but they would require changing the assumptions of each individual BQS and it is subject of future work.

According to the previous discussion, we now state properties that we expect any form of composition for BQS should satisfy. Characterizing the failures the composite BQS can tolerate now becomes the challenge because multiple definitions of $\mathcal{F}_3$ are plausible. We want to ensure the following *properties*:

1. Any $B \in \mathcal{F}_3$ satisfies $B|_{\mathcal{P}_1} \in \mathcal{F}_1^*$, i.e., the failure of B is tolerated in the first system.
2. Any $B \in \mathcal{F}_3$ satisfies $B|_{\mathcal{P}_2} \in \mathcal{F}_2^*$, i.e., the failure of B is tolerated in the second system.
3. $\mathcal{F}_3$ satisfies the Q^3-condition.
4. For any $B \in \mathcal{F}_3$, there exists a $Q \in \mathcal{Q}_3$, a quorum system in the composite system, such that $B \cap Q = \emptyset$, i.e., there is always a quorum consisting only of correct processes.

In the text above, the notation $\mathcal{X}|_{\mathcal{P}}$ denotes the restriction of a set $\mathcal{X}$ to $\mathcal{P}$.

We need properties 1 and 2 because, as we shall see next, they imply Property 3, and, hence, ensure consistency for the composite BQS against any fail-prone set in $\mathcal{F}_3$. Moreover, they enable a composition

by using the existing assumptions, without requiring a redesign of the two systems. One might also desire that the inverse of properties 1 and 2 be satisfied, i.e., that any fail-prone set in $\mathcal{F}_1$ and $\mathcal{F}_2$ be tolerated in $\mathcal{F}_3$. However, we will later see that this does not always result in a BQS (i.e., in a fail-prone system that satisfies the Q^3-condition). Thus, the objective of a composition rule is to satisfy these properties, thus ensuring safety, while producing a maximal fail-prone system $\mathcal{F}_3$ (in the sense that it contains the largest fail-prone sets that could be created without having to redefine the trust assumptions within the original systems). Finally, the composition rule should also satisfy Property 4, which ensures liveness in the composite system.

Lemma 11. *Properties 1 and 2 above imply Property 3.*

Proof. Let us assume that $Q^3(\mathcal{F}_1)$ and $Q^3(\mathcal{F}_2)$. Towards a contradiction, let $F_A, F_B, F_C \in \mathcal{F}_3$ such that $F_A \cup F_B \cup F_C = \mathcal{P}_3$. Now consider the restriction of F_A, F_B and F_C to $\mathcal{P}_1$ (and similarly to $\mathcal{P}_2$). We have that $F_A|_{\mathcal{P}_1} \cup F_B|_{\mathcal{P}_1} \cup F_C|_{\mathcal{P}_1} = \mathcal{P}_1$. However, from Property 1, the sets $F_A|_{\mathcal{P}_1}$, $F_B|_{\mathcal{P}_1}$, and $F_C|_{\mathcal{P}_1}$ are each (subsets of) fail-prone sets in $\mathcal{F}_1$. We thus have found three fail-prone sets that cover $\mathcal{P}_1$, a contradiction to $\mathcal{F}_1$ satisfying the Q^3-condition. □

With this list of goals, we now proceed to specific constructions. In the following, we present three composition methods of increasing suitability and give examples to show their weaknesses and strengths.

Construction 12 (Union composition). Let $\mathcal{Q}_1$ be a BQS defined on processes $\mathcal{P}_1$ with fail-prone system $\mathcal{F}_1$, and $\mathcal{Q}_2$ a BQS defined on processes $\mathcal{P}_2$ with fail-prone system $\mathcal{F}_2$, where $\mathcal{P}_1 \cap \mathcal{P}_2 = \emptyset$. The *union composition* of $\mathcal{Q}_1$ and $\mathcal{Q}_2$ is a system defined on processes $\mathcal{P}_3 = \mathcal{P}_1 \cup \mathcal{P}_2$ with fail-prone system

$$\mathcal{F}_3 = \mathcal{F}_1 \cup \mathcal{F}_2.$$

We can easily verify that the previous definition, given that $\mathcal{P}_1 \cap \mathcal{P}_2 = \emptyset$, fulfills Properties 1 and 2. Thus, $\mathcal{F}_3$ satisfies the Q^3-condition and there exists a BQS $\mathcal{Q}_3$ with fail-prone system $\mathcal{F}_3$.

Lemma 13. *Given $\mathcal{F}_3$ as in Construction 12, a BQS $\mathcal{Q}_3$ is*

$$\mathcal{Q}_3 = \{Q_i \cup Q_j \mid Q_i \in \mathcal{Q}_1, Q_j \in \mathcal{Q}_2\},$$

with $\mathcal{Q}_1$ and $\mathcal{Q}_2$ BQS.

Proof. We first show that consistency property holds. So, for every $Q_1, Q_2 \in \mathcal{Q}_3$ such that $Q_1 = Q_i \cup Q_j$ and $Q_2 = Q_i^{'} \cup Q_j^{'}$, with $Q_i, Q_i^{'} \in \mathcal{Q}_1$ and $Q_j, Q_j^{'} \in \mathcal{Q}_2$, and for every $F \in \mathcal{F}_3$, with $F \in \mathcal{F}_1$ or $F \in \mathcal{F}_2$, we have $Q_1 \cap Q_2 = (Q_i \cup Q_j) \cap (Q_i^{'} \cup Q_j^{'})$, which equals $(Q_i \cap Q_i^{'}) \cup (Q_j \cap Q_j^{'})$, because $\mathcal{P}_1 \cap \mathcal{P}_2 = \emptyset$. By assumption, both $\mathcal{Q}_1$ and $\mathcal{Q}_2$ are BQS. This means that, if $F \in \mathcal{F}_1$, then $\mathcal{Q}_i \cap Q_i^{'} \not\subseteq F$, and if $F \in \mathcal{F}_2$, then $\mathcal{Q}_j \cap Q_j^{'} \not\subseteq F$. The property then follows. Finally, the availability property follows from the fact that $\mathcal{P}_1$ and $\mathcal{P}_2$ are disjoint and $\mathcal{Q}_1$ and $\mathcal{Q}_2$ are BQS. □

However, the fail-prone system obtained by Construction 12 results in a fail-prone system that tolerates only a few failures, namely those tolerated in each of the two original systems, and not any combination of them. Moreover, it would not work if $\mathcal{P}_1$ and $\mathcal{P}_2$ had any processes in common. The next notion moves towards a composition that tolerates any combination of failures that would be tolerated in the original systems.

Construction 14 (Cartesian composition on disjoint sets). Let $\mathcal{Q}_1$ be a BQS defined on processes $\mathcal{P}_1$ with fail-prone system $\mathcal{F}_1$, and $\mathcal{Q}_2$ a BQS defined on processes $\mathcal{P}_2$ with fail-prone system $\mathcal{F}_2$, where $\mathcal{P}_1 \cap \mathcal{P}_2 = \emptyset$. Then the *Cartesian composition* of $\mathcal{Q}_1$ and $\mathcal{Q}_2$ is defined on processes $\mathcal{P}_3 = \mathcal{P}_1 \cup \mathcal{P}_2$ and tolerates the failure of any combination of fail-prone sets of the original BQS. Formally,

$$\mathcal{F}_3 = \{F_i \cup F_j \mid F_i \in \mathcal{F}_1, F_j \in \mathcal{F}_2\}.$$

Lemma 15. *If $Q^3(\mathcal{F}_1)$ and $Q^3(\mathcal{F}_2)$, then for the fail-prone system $\mathcal{F}_3$ according to Construction 14, $Q^3(\mathcal{F}_3)$.*

Proof. Any $B \in \mathcal{F}_3$ satisfies $B|_{\mathcal{P}_1} \in \mathcal{F}_1$ and $B|_{\mathcal{P}_2} \in \mathcal{F}_2$, since $\mathcal{P}_1 \cap \mathcal{P}_2 = \emptyset$. Hence, the composition in Definition 14 satisfies the Properties 1 and 2, and, by Lemma 11, also $Q^3(\mathcal{F}_3)$. □

The previous lemma implies the existence of a BQS $\mathcal{Q}_3$ with fail-prone system $\mathcal{F}_3$. Such a $\mathcal{Q}_3$ can be obtained, as earlier, by

$$\mathcal{Q}_3 = \{Q_i \cup Q_j \mid Q_i \in \mathcal{Q}_1, Q_j \in \mathcal{Q}_2\}.$$

It is easy to show, in a similar way as in Lemma 13, that this $\mathcal{Q}_3$ satisfies consistency and availability properties. Moreover, if $\mathcal{Q}_1$ and $\mathcal{Q}_2$ are canonical, $\mathcal{Q}_3$ will be the canonical BQS for $\mathcal{F}_3$.

Example 16. Let us consider the threshold case. Suppose $\mathcal{Q}_1$ and $\mathcal{Q}_2$ be two BQS, defined on $\mathcal{P}_1$ and $\mathcal{P}_2$, where $\mathcal{P}_1 \cap \mathcal{P}_2 = \emptyset$, containing 7 and 10 processes, and tolerating the failure of any 2 and 3 processes, respectively. This means that the first fail-prone system contains $\binom{7}{2} = 21$ sets of processes and the second fail-prone system contains $\binom{10}{3} = 120$ sets. Because in this work we join systems with already existing failure assumptions, we refrain from changing these assumptions for the composite system. Nevertheless, according to Lemma 15, the Cartesian product of the fail-prone systems leads to a fail-prone system where the Q^3-condition holds, assuming that the starting systems both satisfy the Q^3-condition and are disjoint.

We apply Construction 14 here, observing that the Q^3-condition is the generalization of the condition $n > 3f$ for the threshold case. As a result we obtain an assumption on 17 processes, which tolerates the failure of 5 processes, where 2 processes are from $\mathcal{P}_1$ and 3 from $\mathcal{P}_2$. More formally, the failure of a set F is tolerated in the composite system if and only if $|F \cap P_1| \leq 2 \wedge |F \cap P_2| \leq 3$.

This gives a total of 2520 possible tolerated subsets. Observe that $\mathcal{Q}_3$ is not a threshold BQS any more, and this was intended. A threshold BQS made of 17 processes would tolerate the failure of any 5 processes, which would lead to a total of $\binom{17}{5} = 6188$ fail-prone sets.

Example 17. We now show how Construction 14 fails to create a BQS $\mathcal{Q}_3$ if $\mathcal{P}_1$ and $\mathcal{P}_2$ intersect, because the Q^3-condition may not hold in the composite system. Let $\mathcal{Q}_1$ defined on $\mathcal{P}_1 = \{a, b, c, d, e\}$ with fail-prone system $\mathcal{F}_1 = \{\{a\}, \{b, c\}, \{d\}, \{c, e\}\}$ and $\mathcal{Q}_2$ defined on $\mathcal{P}_2 = \{d, e, f, g, h\}$ with fail-prone system $\mathcal{F}_2 = \{\{d\}, \{e\}, \{f, g\}, \{h\}\}$.

It is easy to verify that the Q^3-condition is satisfied in $\mathcal{Q}_1$ and $\mathcal{Q}_2$. We also see that, according to Construction 14, $\mathcal{Q}_3$ with processes $\mathcal{P}_3 = \mathcal{P}_1 \cup \mathcal{P}_2$ contains, among others, the fail-prone sets $\{a, f, g\}, \{b, c, h\}, \{c, e, d\}$, which cover $\mathcal{P}_3$. Consequently, $\mathcal{Q}_3$ is not a BQS.

Example 17 shows that the Cartesian composition among fail-prone systems does not lead to a fail-prone system where the Q^3-condition holds, if the two systems have common processes. To overcome this issue, we introduce a third construction.

Definition 18. Let $\mathcal{A} = \{A_1, \ldots, A_m\}$ and $\mathcal{B} = \{B_1, \ldots, B_n\}$ be two sets of subsets of $\mathcal{P}_1$ and $\mathcal{P}_2$, respectively. We define $\mathcal{A} \otimes \mathcal{B}$ as the set that contains the union of all sets $A_i \in \mathcal{A}^*$ and $B_j \in \mathcal{B}^*$, under the restriction that either both A_i and B_j contain exactly the same subset of the processes common to $\mathcal{P}_1$ and $\mathcal{P}_2$ or they do not have anything in common. Formally,

$$\mathcal{A} \otimes \mathcal{B} = \big\{A_i \cup B_j \mid A_i \in \mathcal{A}^* \wedge B_j \in \mathcal{B}^* \wedge (\forall C \subseteq \mathcal{P}_1 \cap \mathcal{P}_2 : C \subseteq A_i \Leftrightarrow C \subseteq B_j)\big\}.$$

Construction 19 (Cartesian composition). Let $\mathcal{Q}_1$ be a BQS defined on processes $\mathcal{P}_1$ with fail-prone system $\mathcal{F}_1$ and $\mathcal{Q}_2$ a BQS defined on processes $\mathcal{P}_2$ with fail-prone system $\mathcal{F}_2$, where $\mathcal{P}_1$ and $\mathcal{P}_2$ might contain common processes. Then the *composition* of $\mathcal{Q}_1$ and $\mathcal{Q}_2$ is defined on $\mathcal{P}_3 = \mathcal{P}_1 \cup \mathcal{P}_2$ and tolerates the failure of any combination of any fail-prone set (or subset of it) of the first system and any fail-prone set (or subset) of the second system, such that both contain exactly the same subset of the common processes. Formally,

$$\mathcal{F}_3 = \mathcal{F}_1 \otimes \mathcal{F}_2 = \big\{F_i \cup F_j \mid F_i \in \mathcal{F}_1^* \wedge F_j \in \mathcal{F}_2^* \wedge (\forall C \subseteq \mathcal{P}_1 \cap \mathcal{P}_2 : C \subseteq F_i \Leftrightarrow C \subseteq F_j)\big\}.$$

The rule of Construction 19 states that any fail-prone set in $\mathcal{F}_3$ is of the form $F_i \cup F_j$, where F_i and F_j are fail-prone sets (or subsets of fail-prone sets) that either do not have any processes in common or, if they do, both contain exactly the same subset of $\mathcal{P}_1 \cup \mathcal{P}_2$. We demand $F_i \in \mathcal{F}_1^*$ and $F_j \in \mathcal{F}_2^*$, instead of $F_i \in \mathcal{F}_1$ and $F_j \in \mathcal{F}_2$, in order to construct a maximal $\mathcal{F}_3$, in the sense that it contains the maximal fail-prone sets that satisfy Properties 1 and 2.

Lemma 20. *If $Q^3(\mathcal{F}_1)$ and $Q^3(\mathcal{F}_2)$, then $Q^3(\mathcal{F}_3)$, with $\mathcal{F}_3$ as in Construction 19.*

Proof. Any $B \in \mathcal{F}_3$ either does not contain a set of common processes C among $\mathcal{P}_1$ and $\mathcal{P}_2$ or it does. In the former case, it is immediate to see that $B|_{\mathcal{P}_1} \in \mathcal{F}_1^*$ and $B|_{\mathcal{P}_2} \in \mathcal{F}_2^*$. In the latter case, B has been created as the union between $F_i \in \mathcal{F}_1^*$ and $F_j \in \mathcal{F}_2^*$, both containing the same subset of $P_1 \cap P_2$, according to Construction 19. It is thus not possible that a new element of $\mathcal{P}_1$ appears in $B|_{\mathcal{P}_1}$ that was not already in F_i, and similarly that a new element of $\mathcal{P}_2$ appears in $B|_{\mathcal{P}_2}$ that was not already in F_j. This implies that $B|_{\mathcal{P}_1} \in \mathcal{F}_1^*$ and $B|_{\mathcal{P}_2} \in \mathcal{F}_2^*$, and from Lemma 11 we get $Q^3(\mathcal{F}_3)$. □

Lemma 21. *Given $\mathcal{F}_3$ as in Construction 19, a BQS $\mathcal{Q}_3$ is*

$$\mathcal{Q}_3 = \{Q_i \cup Q_j \mid Q_i \in \mathcal{Q}_1, Q_j \in \mathcal{Q}_2\},$$

with $\mathcal{Q}_1$ and $\mathcal{Q}_2$ BQS.

Proof. Consistency and availability properties of $\mathcal{Q}_3$ can be proved in a similar way as Lemma 13, assuming $\mathcal{Q}_1$ and $\mathcal{Q}_2$ to be BQS. In fact, as in Lemma 13, we have that for every $Q_1, Q_2 \in \mathcal{Q}_3$, such that $Q_1 = Q_i \cup Q_j$ and $Q_2 = Q_i' \cup Q_j'$, with $Q_i, Q_i' \in \mathcal{Q}_1$ and $Q_j, Q_j' \in \mathcal{Q}_2$, we have $Q_1 \cap Q_2 = (Q_i \cup Q_j) \cap (Q_i' \cup Q_j')$, which results in $(Q_i \cap Q_i') \cup (Q_i \cap Q_j') \cup (Q_j \cap Q_i') \cup (Q_j \cap Q_j')$. If $\mathcal{P}_1 \cap \mathcal{P}_2 = \emptyset$, it is trivial to prove the result. Otherwise, given $F \in \mathcal{F}_3$ with $F = F_i \cup F_j$, where $F_i \in \mathcal{F}_1^* \wedge F_j \in \mathcal{F}_2^* \wedge \forall C \subseteq \mathcal{P}_1 \cap \mathcal{P}_2 : C \subseteq F_i \Leftrightarrow C \subseteq F_j$, we have two cases. If there are no common processes between F_i and F_j, then observe that F_i is contained in $\mathcal{F}_i^*$ and it is then a subset of a fail-prone set $\overline{F}_i$ in $\mathcal{F}_i$. The same happens for F_j. By assumptions, $\mathcal{Q}_1$ (respectively, $\mathcal{Q}_2$) are BQS. It follow that, $(Q_i \cap Q_i')$ (respectively, $(Q_j \cap Q_j')$) is not a proper subset of $\overline{F}_i$ and consequently of F_i (respectively of F_j). The result follows. The same reasoning can be applied if F_i and F_j contain a common subset $C \subseteq \mathcal{P}_1 \cap \mathcal{P}_2$. □

Example 22. Let us consider again the threshold case, where $\mathcal{Q}_1$ is defined on participants $\mathcal{P}_1 = \{a, b, c, d, e, f, g\}$ and $\mathcal{Q}_2$ on $\mathcal{P}_2 = \{g, h, i, j, k, l, m, n, o, p\}$. According to Construction 19, any two processes in $\mathcal{P}_1$ together with any three processes in $\mathcal{P}_2$ are tolerated to fail, because these failures would be tolerated in the original systems. However, if g together with any other process in $\mathcal{P}_1$ fails, then only two more failures in $\mathcal{P}_2$ are tolerated, because $g \in \mathcal{P}_2$ has already failed in the first system.

Example 23. Let $\mathcal{Q}_1$ be defined on processes $\mathcal{P}_1 = \{a, b, c, d, e\}$ and with fail-prone system $\mathcal{F}_1 = \{\{a\}, \{b, c\}, \{d\}, \{c, e\}\}$ and $\mathcal{Q}_2$ be defined on processes $\mathcal{P}_2 = \{d, e, f, g, h\}$ with fail-prone system $\mathcal{F}_2 = \{\{d\}, \{e\}, \{f, g\}, \{h\}\}$. Then, according to Construction 19 processes in $\mathcal{P}_3 = \{a, b, c, d, e, f, g, h\}$ have fail-prone system

$$\mathcal{F}_3 = \{\{a, f, g\}, \{a, h\}, \{b, c, f, g\}, \{b, c, h\}, \{d\}, \{c, e\}\}.$$

It is easy to verify that $Q^3(\mathcal{F}_3)$.

5 Composition of asymmetric BQS

We now explore the composition of two asymmetric Byzantine quorum systems. Given two ABQS, $\mathbb{Q}_1$ defined on processes $\mathcal{P}_1$ with fail-prone system $\mathbb{F}_1$, and $\mathbb{Q}_2$ defined on processes $\mathcal{P}_2$ with fail-prone system $\mathbb{F}_2$, we want to provide a *composition* rule that allows the processes $\mathcal{P}_3 = \mathcal{P}_1 \cup \mathcal{P}_2$ to form an ABQS $\mathcal{Q}_3$ with fail-prone system $\mathbb{F}_3$.

5.1 Approaches to asymmetric trust

In the context of blockchains, different models of asymmetric trust have been proposed, united by a shared principle regarding the subjectivity of truth among the participants, but differentiated by fundamental properties that determine their limitations and strengths.

In this section, we compare our model of asymmetric trust [4] with the model of personal Byzantine quorum system (PBQS) introduced by Losa *et al.* [12]. PBQS extend and improve the formalization used by Stellar [17, 11] and give a new interpretation of a quorum system for subjective trust.

In a PBQS, each participant has its own notion of a quorum, with the requirement that if $Q_i \subseteq \mathcal{P}$ is a quorum for a process p_i and $p_j \in Q_i$, then it exists Q_j for p_j such that $Q_j \subseteq Q_i$. In other words, a quorum Q_i for p_i should contain at least one quorum for each $p_j \in Q_i$. Given this definition, a PBQS consists of a set of participants $\mathcal{P}$, a set of faulty processes $F \subseteq \mathcal{P}$, a set of correct processes $\mathcal{W} = \mathcal{P} \setminus F$, and a function mapping a participant p_i to its non-empty set of quorums. In other words, Losa *et al.* construct a PBQS starting from an arbitrary set F. This notion of a quorum system differs also from the well-known formalization [14] because a quorum in a PBQS is a private notion. In the traditional model, all quorums are public and known to every participant. From this, it follows that a global intersection property is absent from PBQS.

An asymmetric Byzantine quorum system (ABQS, cf. Section 3.2) is defined from an asymmetric fail-prone system, which contains the fail-prone systems of every participant, and requires a global intersection property for consistency. ABQS extend the traditional notion of Byzantine quorum systems [14]. With an ABQS, the correct processes can be grouped into naïve and wise ones, depending on their trust assumptions. According to this distinction, one can guarantee most properties of a protocol only to wise processes.

An useful structure in an ABQS $\mathbb{Q}$ is a *kernel* [4, 5] of each quorum system $\mathcal{Q}_i$ for p_i. This is a set $K_i \subseteq \mathcal{P}$ with the property that for every $Q \in \mathcal{Q}_i : K \cap Q \neq \emptyset$. In other words, a kernel is a set of processes that intersects every quorum in a quorum system $\mathcal{Q}_i$ for a process p_i; it generalizes sets of size $f+1$ in the traditional symmetric threshold model. Losa *et al.* define a similar structure called a *blocking* set. In particular, given $\mathcal{R}$ a set of participants, a process p_i is *blocked* by $\mathcal{R}$ when every quorum of p_i intersects $\mathcal{R}$. Moreover, they show that if a process p_i is blocked by the set of faulty processes F, then it is impossible to guarantee liveness for p_i. With ABQS, this cannot happen: by the availability property of an ABQS, for every set of faulty processes, it always exists a quorum $Q_i \in \mathcal{Q}_i$ for p_i that consists only of correct processes. It follows that, eventually, a process p_i will hear from a quorum of processes for itself, even if all the malicious processes remain silent.

Finally, an ABQS execution gives rise to a *guild*, a set of wise processes that contains at least one quorum for each of its members. The existence of a guild is essential for protocols with ABQS and it also plays a fundamental role for the composition of ABQS. Guilds cannot be disjoint and the union of two guilds is again a guild [5]. The analogue of a guild in a PBQS is a *consensus cluster*, which is a subset $\mathcal{S} \subseteq \mathcal{W}$ such that for every two quorums Q_i and Q_j of some members of $\mathcal{S}$, it holds $Q_i \cap Q_j \cap \mathcal{W} \neq \emptyset$, and for every $p_i \in \mathcal{S}$, it exists a quorum Q_i for p_i such that $Q_i \subseteq \mathcal{S}$. However, despite these similarities to a guild, two consensus clusters can be disjoint (due to the missing intersection requirement in a PBQS). This implies that for consensus with PBQS, agreement may hold only locally, and achieving consensus across disjoint clusters may not be possible.

5.2 The tolerated system of an ABQS

For defining composition with ABQS, we first introduce the central notion of the *tolerated system* of an ABQS. Recall that symmetric BQS start from a common understanding of the world; the participants agree on the possible failures, that is, on which participants might crash or collaborate to break security. In an asymmetric BQS, no such common understanding exists, either because there is not enough knowledge to make such an assumption on the system, or because the participants simply do not agree with each other. In this model, every participant expresses its own beliefs and expectations, and no global notion of "correct" belief exists. In every execution, however, there will be a ground truth, manifested

by a set of actually faulty participants, and not all members of the system will have correctly anticipated this ground truth. Again, since there is no global understanding of the world, this is expected to happen. However, the participants might still be able to make progress (where progress is defined by the protocol they are running), exactly in those executions when a guild exists. Recent works on consensus with ABQS have conditioned safety and liveness properties on the existence of such a set. In the context of Byzantine consensus, Cachin and Zanolini [5] show that a guild is required to solve asynchronous consensus and that consensus properties are guaranteed in all executions with a guild.

An external party examining an ABQS without any prior knowledge or beliefs about the participants cannot assess the trust assumptions of any individual participant. However, the third party can evaluate the ABQS based on its ability to make progress through a guild.

The central concept for composing two ABQS is the *tolerated system* of an ABQS. Recall that in an execution where all processes in $B \subset \mathcal{P}$ actually fail, there may also be naïve processes, wise processes that form a guild $\mathcal{G}$, and wise processes outside the guild ([5, Example 1]). For a specific guild $\mathcal{G} \neq \emptyset$, the union of all those processes outside $\mathcal{G}$ is called a *tolerated set* because the guild is autonomous without any of them. Hence, the tolerated set consists of the faulty, the naïve, and the wise processes outside the guild. The tolerated system contains all the tolerated sets. Formally, we have the following definition.

Definition 24 (Tolerated system). The *tolerated system* $\mathcal{T}$ of an ABQS $\mathbb{Q}$ defined on processes $\mathcal{P}$ is

$$\mathcal{T} = \{\mathcal{P} \setminus G, \text{for any possible guild } G \text{ of } \mathbb{Q}\}.$$

Intuitively, the tolerated system of an ABQS reflects the resilience of the ABQS: even without the processes in a tolerated set, there still exists a guild. Therefore, the tolerated system characterizes the executions in which some of the participants in the asymmetric system will be able to operate correctly and make progress. In that sense, the tolerated system of an ABQS is the counterpart of the fail-prone system for a BQS.

Notice that the tolerated system is a global notion emerging from the subjective trust choices of the participating processes; any party that knows the fail-prone and quorum systems of all processes can calculate it. We show later that the tolerated systems of two ABQS play a crucial role for composing them; the processes in the first system will use the tolerated sets of the second system as their trust assumptions, and vice versa. Consequently, the processes in the first system only need to know the tolerated system of the second system.

The following lemma shows that the tolerated system of a canonical ABQS naturally corresponds to a BQS.

Lemma 25. *Let $\mathbb{Q}$ be an ABQS on processes $\mathcal{P}$ with asymmetric fail-prone system $\mathbb{F} = \overline{\mathbb{Q}}$, i.e., such that $\mathbb{Q}$ is a canonical ABQS. Then the tolerated system $\mathcal{T}$ of $\mathbb{Q}$ is a BQS. In particular, if $B^3(\mathbb{F})$, then $Q^3(\mathcal{T})$.*

Proof. Towards a contradiction, let us assume that $\mathcal{T}$ does not satisfy the Q^3-condition. This means that there exist $T_1, T_2, T_3 \in \mathcal{T}$ such that $T_1 \cup T_2 \cup T_3 = \mathcal{P}$. Also, let $\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3$ be the corresponding guilds, i.e., $\mathcal{G}_1 = \mathcal{P} \setminus \mathcal{T}_1, \mathcal{G}_2 = \mathcal{P} \setminus \mathcal{T}_2$ and $\mathcal{G}_3 = \mathcal{P} \setminus \mathcal{T}_3$. Without loss of generality every guild contains at least a process, and at least a quorum for this process is fully contained in the guild. By the consistency property of an ABQS, these quorums must intersect pairwise, hence the guilds also intersect pairwise. This means that there exist processes $p_i \in \mathcal{G}_1 \cap \mathcal{G}_2$ and $p_j \in \mathcal{G}_2 \cap \mathcal{G}_3$. Now, because p_i is a member of $\mathcal{G}_1$, we can make the following reasoning: p_i has a quorum $Q_i \in \mathcal{Q}_i$ such that $Q_i \subseteq \mathcal{G}_1$, the BQS is canonical, so p_i has a fail-prone set $F_i = \mathcal{P} \setminus Q_i \in \mathcal{F}_i$, thus we get $T_1 \subseteq F_i$, i.e., $T_1 \in \mathcal{F}_i$. With similar reasoning, we get $T_2 \in \mathcal{F}_i$ (because $p_i \in \mathcal{G}_2$), $T_2 \in \mathcal{F}_j$ (because $p_j \in \mathcal{G}_2$), and $T_3 \in \mathcal{F}_j$ (because $p_j \in \mathcal{G}_3$). But this is a contradiction, because p_i and p_j with fail-prone sets T_1, T_2, and T_3 violate the B^3-condition in $\mathbb{Q}$. □

As has been known before, by Lemma 4, if $\mathcal{T}$ satisfies the Q^3-condition, then there exists also a symmetric BQS for the fail-prone system $\mathcal{T}$; for instance, this may be the canonical BQS $\overline{\mathcal{T}}$.

Lemma 25 confirms the intuition that the tolerated set of an ABQS is the counterpart of a fail-prone set in a BQS.

5.3 How clients interact with an ABQS

Many practical replication protocols separate clients from replicas; in state-machine replication, clients submit commands, replicas totally order and execute them, and then send back responses to the clients. When the expected failures among replicas are modeled as a BQS, that is, with a symmetric trust assumption, the clients wait for responses from a quorum of replicas. However, if the trust assumption among the replicas is asymmetric, it is unclear which sets of participants are capable to convince a client to accept a response. The subjective quorums of the replicas only express their personal beliefs, which the clients may not share.

One way to resolve this could be to let each client express trust in the replicas through its own quorum system. But if clients do not have sufficient knowledge to make such assumptions, they need a global property of the quorum system to decide on its responses, and this can be the tolerated system. Note that every guild formed by replicas corresponds to the complement of a tolerated set. This indicates that (at least some) replicas did agree on their trustworthiness, and this may convince the client. Indeed, we will use this idea in the composition procedure for ABQS. Specifically, the participants of each system may operate as clients of the other and could send a composition-request message, waiting for responses from a guild of participants.

5.4 Composition of ABQS

Based on the remarks above, the we claim that any form of composition between two ABQS must satisfy the following conditions. Regarding notation, we want to compose $\mathbb{Q}_1$ with $\mathbb{Q}_2$, resulting in $\mathbb{Q}_3$, with respective asymmetric fail-prone systems $\mathbb{F}_1$, $\mathbb{F}_2$, and $\mathbb{F}_3$. For $k = 1, 2$ and for any $p_i \in \mathcal{P}_k$, let $\mathcal{F}_i^{(k)}$ be the fail-prone system of p_i in $\mathbb{F}_k$, and $\mathcal{F}_i^{(3)}$ the fail-prone system of p_i in the resulting $\mathbb{F}_3$. Moreover, let $\mathcal{T}_k$ be the tolerated system of $\mathbb{Q}_k$.

1. If $p_i \in \mathcal{P}_1$ and $p_i \in \mathcal{P}_2$, then any $F_i \in \mathcal{F}_i^{(3)}$ must respect the trust assumptions of p_i in $\mathcal{P}_1$ and in $\mathcal{P}_2$, i.e., it must satisfy $F_i|_{\mathcal{P}_1} \in \mathcal{F}_i^{(1)^*}$ and $F_i|_{\mathcal{P}_2} \in \mathcal{F}_i^{(2)^*}$. If p_i is only in $\mathcal{P}_1$ (and the same holds for $\mathcal{P}_2$), then any $F_i \in \mathcal{F}_i^{(3)}$ must respect the assumptions of p_i in $\mathcal{P}_1$, i.e., $F_i|_{\mathcal{P}_1} \in \mathcal{F}_i^{(1)^*}$, and $F_i|_{\mathcal{P}_2}$ can only be one of the tolerated sets in $\mathcal{P}_2$, i.e., $F_i|_{\mathcal{P}_2} \in \mathcal{T}_2^*$, since p_i has no assumptions for $\mathcal{P}_2$. This generalizes Properties 1 and 2 of the symmetric composition.

2. If the B^3-condition holds for $\mathbb{F}_1$ and for $\mathbb{F}_2$, then it also holds for composite system, for $\mathbb{F}_3$. This is a generalization of Property 3 of the symmetric composition.

3. For any $p_i \in \mathcal{P}_3$ and any $F_i \in \mathcal{F}_i^{(3)}$, there exists a quorum $Q_i \in \mathcal{Q}_i^{(3)}$, such that $F_i \cap Q_i = \emptyset$.

Up to here, these three properties are generalizations of the corresponding properties of the symmetric composition. However, in the asymmetric case, we also want to achieve the following.

4. **Preserving wisdom.** In all executions, where there exists a guild $\mathcal{G}_1$ in ABQS $\mathbb{Q}_1$ and a guild $\mathcal{G}_2$ in $\mathbb{Q}_2$, the processes in $\mathcal{G}_1 \cup \mathcal{G}_2$ will form a guild in $\mathbb{Q}_3$. The intuition is that, given an execution with B as actual faulty set, if a process correctly foresees B (and thus enjoys the properties of a guild) in its own system, and if there is a guild in the other system, then this process should also enjoy the properties of a guild in the composite system.

5. **Reducibility to symmetric.** If all processes have the same trust assumptions (in which case $\mathbb{Q}_1$ and $\mathbb{Q}_2$ reduce to symmetric BQS), then the composite system $\mathbb{Q}_3$ is a symmetric BQS and satisfies the properties of symmetric composition.

Lemma 26. *Property 1 implies Property 5.*

Proof. This follows immediately by observing that when all processes in $\mathcal{P}_k$ have the same fail-prone system $\mathcal{F}_k$, for $k = 1, 2$, then the tolerated system $\mathcal{T}_k$ is $\mathcal{F}_k$ itself. Then, Property 1 implies that $\mathcal{F}_i^{(3)}$ is the same for every $p_i \in \mathcal{P}_3$, and that every $B \in \mathcal{F}_i^{(3)}$ satisfies $B|_{\mathcal{P}_1} \in {\mathcal{F}_1}^*$ and $B|_{\mathcal{P}_2} \in {\mathcal{F}_2}^*$, which is what Properties 1 and 2 of the symmetric composition require. □

Now let us consider two ABQS $\mathbb{Q}_1$ and $\mathbb{Q}_2$ on processes $\mathcal{P}_1$ and $\mathcal{P}_2$ with asymmetric fail-prone systems $\mathbb{F}_1$ and $\mathbb{F}_2$, respectively. All processes in $\mathcal{P}_1$ and $\mathcal{P}_2$ wish to jointly run a protocol, without making any extra assumption about the participants of the other group. Intuitively speaking, each group might have their own issues, their own agreements and disagreements, their own good and bad executions, but they still want to work together. As reasoned earlier, each participant in $\mathcal{P}_1$ is an external observer for $\mathcal{P}_2$. Hence, the best a participant in $\mathcal{P}_1$ can do, assuming they have no knowledge, beliefs, or assumptions for the participants of the second group, is to use the tolerated system of $\mathbb{Q}_2$. The same applies, of course, for participants in $\mathcal{P}_2$. This leads to the composition procedure we describe next.

Construction 27 (Purification). Let $\mathbb{Q}$ an ABQS on processes $\mathcal{P} = \{p_1, \ldots, p_n\}$, with asymmetric fail-prone system $\mathbb{F} = \{\mathcal{F}_1, \ldots, \mathcal{F}_n\}$, such that $B^3(\mathbb{F})$, and let $\mathcal{T}$ its tolerated system. Assume $Q^3(\mathcal{T})$. As we have seen, this is always the case for canonical ABQS. We want to *purify* $\mathbb{F}$ so that $B^3([\mathcal{F}_1, \ldots, \mathcal{F}_n, \mathcal{T}])$, i.e., $\forall F_i \in \mathcal{F}_i, \forall F_j \in \mathcal{T}, \forall F_{ij} \in \mathcal{F}_i^* \cap \mathcal{T}^*$ it holds that $\mathcal{P} \not\subseteq F_i \cup F_j \cup F_{ij}$. To do so, every process p_i evaluates the B^3-condition including $\mathcal{T}$ in the asymmetric fail-prone system $\mathbb{F}$. If it does not hold, then for any $F_i \in \mathcal{F}_i$ that violates the B^3-condition, p_i removes F_i from $\mathcal{F}_i$, and adds to $\mathcal{F}_i$ all those subsets of F_i that do not violate the B^3-condition. This results in a *purified* fail-prone system, which, by construction, satisfies the B^3-condition.

Intuitively, the purification procedure removes fail-prone systems that are "useless," in the sense that they do not influence the existence of a guild, as shown by the next lemma. Seen from a higher level, it is an expression of the fact that processes have their own beliefs, but also need to adapt to those of the others; a process p_i might expect a set F to fail during an execution and construct its fail-prone system $\mathcal{F}_i$ so as to be protected against F. However, if the beliefs of other processes are such that the failure of F does not lead to a guild, i.e., F is not tolerated, then p_i can not benefit from including F in $\mathcal{F}_i$.

Lemma 28. *For every possible execution with a guild $\mathcal{G}$, a process in $\mathcal{G}$ of the non-purified system is also contained in some guild of the purified system.*

Proof. Observe that the $F_i \in \mathcal{F}_i$ which p_i removes cannot be in $\mathcal{T}$, because otherwise it would be possible to cover all $\mathcal{P}$ with sets in $\mathcal{T}$; but this is not possible by the assumption $Q^3(\mathcal{T})$. This implies that the failure of F_i cannot lead to the existence of a guild, and can be removed from $\mathcal{F}_i$. On the other hand, subsets of F_i can possibly be in $\mathcal{T}$, and p_i keeps those subsets in $\mathcal{F}_i$. □

Observe that the purification procedure is deterministic and uses information that is available to every process in the system: evaluating the B^3-condition, for example, already assumes that every process in the system knows the asymmetric fail-prone systems of the others and that Byzantine processes do not lie about their assumptions.

Construction 29 (Composition of ABQS). Let $\mathcal{P}_1 = \{p_1, \ldots, p_{m+k}\}$ and $\mathcal{P}_2 = \{p_{m+1}, \ldots, p_n\}$ be two sets of processes, with processes $p_{m+1}, \ldots, p_{m+k}$ in common. Let $\mathbb{Q}_1$ be an ABQS on processes $\mathcal{P}_1$ with asymmetric fail-prone system $\mathbb{F}_1 = \{\mathcal{F}_1^{(1)}, \ldots, \mathcal{F}_{m+k}^{(1)}\}$, and $\mathbb{Q}_2$ an ABQS on processes $\mathcal{P}_2$ with asymmetric fail-prone system $\mathbb{F}_2 = \{\mathcal{F}_{m+1}^{(2)}, \ldots, \mathcal{F}_n^{(2)}\}$, where $\mathbb{F}_1$ and $\mathbb{F}_2$ are purified. Moreover, let $\mathcal{T}_1$ and $\mathcal{T}_2$ be the tolerated systems of the two ABQS, respectively. The *composite fail-prone system* $\mathbb{F}_3$ on processes $\mathcal{P}_3 = \mathcal{P}_1 \cup \mathcal{P}_2$ is

$$\mathbb{F}_3 = [\mathcal{F}_1^{(1)} \otimes \mathcal{T}_2, \ldots, \mathcal{F}_m^{(1)} \otimes \mathcal{T}_2, \mathcal{F}_{m+1}^{(1)} \otimes \mathcal{F}_{m+1}^{(2)}, \ldots, \mathcal{F}_{m+k}^{(1)} \otimes \mathcal{F}_{m+k}^{(2)}, \mathcal{F}_{m+k+1}^{(2)} \otimes \mathcal{T}_1, \ldots, \mathcal{F}_n^{(2)} \otimes \mathcal{T}_1].$$

and the *composite ABQS* $\mathbb{Q}_3$ is any asymmetric quorum system for $\mathbb{F}_3$.

Lemma 30. *The composed fail-prone system $\mathbb{F}_3$ resulting from Construction 29 satisfies the B^3-condition.*

Proof. Towards a contradiction, let us assume that the B^3-condition does not hold on $\mathbb{F}_3$. This means there exist processes p_i and p_j and fail-prone sets $F_i \in \mathcal{F}_i^{(3)}$, $F_j \in \mathcal{F}_j^{(3)}$, and $F_{ij} \in \mathcal{F}_i^{(3)^*} \cap \mathcal{F}_j^{(3)^*}$ such that $\mathcal{P}_3 = F_i \cup F_j \cup F_{ij}$. In the following we consider the restriction of F_i, F_j, and F_{ij} to $\mathcal{P}_1$, i.e., $F_i|_{\mathcal{P}_1}, F_j|_{\mathcal{P}_1}$, and $F_{ij}|_{\mathcal{P}_1}$, respectively. We distinguish two cases for p_i and p_j. First, consider the case where p_i and p_j belong to different sets of processes and let, w.l.o.g., $p_i \in \mathcal{P}_1 \setminus \mathcal{P}_2$ and $p_j \in \mathcal{P}_2 \setminus \mathcal{P}_1$. By the definition of the $\otimes$ operator, and with an argument similar to what we used in the proof of Lemma 20, we get that $F_i|_{\mathcal{P}_1} \in \mathcal{F}_i^{(1)^*}$, that $F_j|_{\mathcal{P}_1} \in \mathcal{T}_1{}^*$, that $F_{ij}|_{\mathcal{P}_1}$ is a common subset of $\mathcal{F}_i^{(1)*}$ and $\mathcal{T}_1^*$, and that their union covers $\mathcal{P}_1$. This is a contradiction because $\mathbb{F}_1$ is purified. Second, consider the case where at least one of p_i and p_j belongs to both $\mathcal{P}_1$ and $\mathcal{P}_2$, and let, w.l.o.g., $p_i \in \mathcal{P}_1, p_j \in \mathcal{P}_1 \cap \mathcal{P}_2$. (If $p_i \in \mathcal{P}_1 \cap \mathcal{P}_2$ the same reasoning can be applied by projecting in $\mathcal{P}_2$.) For this case, we observe that $F_i|_{\mathcal{P}_1} \in \mathcal{F}_i^{(1)^*}$, $F_j|_{\mathcal{P}_1} \in \mathcal{F}_j^{(1)^*}$, and that $F_{ij}|_{\mathcal{P}_1}$ is a common subset of a fail-prone set in $\mathcal{F}_i^{(1)*}$ and a fail-prone set in $\mathcal{F}_j^{(1)*}$. This contradicts the assumption that $B^3(\mathbb{F}_1)$. □

Remark 31. Given an ABQS $\mathbb{Q}_1$ for an asymmetric fail-prone system $\mathbb{F}_1$ on processes $\mathcal{P}_1$, and an ABQS $\mathbb{Q}_2$ for $\mathbb{F}_2$ on $\mathcal{P}_2$, and assuming that the processes of each BQS make no assumptions about processes in the other, a composition of the two systems is only possible if the corresponding tolerated systems $\mathcal{T}_1$ and $\mathcal{T}_2$ both satisfy the Q^3-condition. This is because the processes of the first ABQS (and vice versa) are only external observers for the second system, and therefore only assess it through its tolerated system. Processes in $\mathcal{P}_1$ want to make sure that whenever the second system is able to make progresses (that is, for every $T \in \mathcal{T}_2$ that leads to a guild $\mathcal{G}$), they will also be able to make progress. To achieve this, they must consider all the $T \in \mathcal{T}_2$ as a possible actual failed set. However, because the processes of the first system do not assume anything about the second system, the only way to achieve this is to include all the $T \in \mathcal{T}_2$ in their fail-prone sets. This leads to an ABQS if and only if the Q^3-condition holds in the second system (and vice versa).

Lemma 30 and Lemma 8 together imply the existence of an ABQS for $\mathbb{F}_3$ as defined in Construction 29. This is the asymmetric canonical quorum system $\mathbb{Q}_3 = \overline{\mathbb{F}}_3$.

For instance, let us consider two ABQS $\mathbb{Q}_1$ and $\mathbb{Q}_2$ on processes $\mathcal{P}_1 = \{p_1, \ldots, p_m\}$ and $\mathcal{P}_2 = \{p_{m+1}, \ldots, p_n\}$ with asymmetric fail-prone systems $\mathbb{F}_1$ and $\mathbb{F}_2$, respectively, such that $\mathcal{P}_1 \cap \mathcal{P}_2 = \emptyset$. Then, the asymmetric canonical quorum system for $\mathbb{F}_3$ is

$$\mathbb{Q}_3 = [\mathcal{Q}_1 \cup \overline{\mathcal{T}}_2, \ldots, \mathcal{Q}_m \cup \overline{\mathcal{T}}_2, \mathcal{Q}_{m+1} \cup \overline{\mathcal{T}}_1, \ldots, \mathcal{Q}_n \cup \overline{\mathcal{T}}_1],$$

where $\mathcal{Q}_i = \overline{\mathcal{F}}_i$, $\mathcal{Q}_i \cup \overline{\mathcal{T}}_j = \{Q_k \cup \mathcal{G}_l \mid Q_k \in \mathcal{Q}_i \wedge \mathcal{G}_l \in \overline{\mathcal{T}}_j\}$ and $\mathcal{G}_i$ is a guild for a tolerated set in $\mathcal{T}_j$. Notice that, by definition, $\overline{\mathcal{T}}$ contains all the guilds that can be obtained within an ABQS.

As a short proof of why $\mathbb{Q}_3$ is the canonical asymmetric quorum system of $\mathbb{F}_3$, we observe that, by assuming $\mathcal{P}_1 \cap \mathcal{P}_2 = \emptyset$, the asymmetric fail-prone system $\mathbb{F}_3$ in Construction 29 reduces to

$$\mathbb{F}_3 = [\mathcal{F}_1 \cup \mathcal{T}_2, \ldots, \mathcal{F}_m \cup \mathcal{T}_2, \mathcal{F}_{m+1} \cup \mathcal{T}_1, \ldots, \mathcal{F}_n \cup \mathcal{T}_1],$$

where $\mathcal{F}_i \cup \mathcal{T}_j = \{F_k \cup T_l \mid F_k \in \mathcal{F}_i \wedge T_l \in \mathcal{T}_j\}$. If we consider the bijective complement of $\mathcal{F}_i \cup \mathcal{T}_j$ this is made by all the sets of the form $\overline{F_k \cup T_l}$ in $\mathcal{P}_3 = \mathcal{P}_1 \cup \mathcal{P}_2$. Then, $\overline{F_k \cup T_l} = \overline{F}_k \cap \overline{T}_l = (Q_k \cup \mathcal{P}_2) \cap (\mathcal{G}_l \cup \mathcal{P}_1)$ where $Q_k = \overline{F}_k$ in $\mathcal{P}_1$. Finally, $(Q_k \cup \mathcal{P}_2) \cap (\mathcal{G}_l \cup \mathcal{P}_1) = (Q_k \cap \mathcal{G}_l) \cup (Q_k \cap \mathcal{P}_1) \cup (\mathcal{P}_2 \cap \mathcal{G}_l) \cup (\mathcal{P}_2 \cap \mathcal{P}_1)$. Observe that, by assumption on the sets of processes, it follows that $(\mathcal{P}_2 \cap \mathcal{P}_1) = \emptyset$ and $(Q_k \cap \mathcal{G}_l) = \emptyset$. So, $\overline{F_k \cup T_l} = (Q_k \cap \mathcal{P}_1) \cup (\mathcal{P}_2 \cap \mathcal{G}_l) = Q_k \cup \mathcal{G}_l$.

5.5 Composition in practice

We now sketch a protocol that can be used by two (possibly disjoint) sets of processes $\mathcal{P}_1$ and $\mathcal{P}_2$ that form two asymmetric Byzantine quorum systems $\mathbb{Q}_1$ and $\mathbb{Q}_2$ with asymmetric fail-prone systems $\mathbb{F}_1$ and $\mathbb{F}_2$, respectively. We assume that processes in $\mathcal{P}_1$ and $\mathcal{P}_2$ are running two different instances of the same Byzantine consensus protocol (i.e., providing total-order broadcast) and that $\mathbb{F}_1$ and $\mathbb{F}_2$ are publicly known.

The composition can be initiated by any process p_i in $\mathcal{P}_1$. To that end, process p_i, acting as a client for $\mathbb{Q}_1$, sends a *composition-request* message to every process in $\mathcal{P}_2$. Upon receiving this request, processes in $\mathcal{P}_2$ start a round of Byzantine consensus: if a sufficient number of processes votes for the composition, it will be agreed. Assume the protocol instance run by $\mathbb{Q}_2$ has a history of delivered messages $\mathcal{H}_2$ at this point. Then, upon deciding, processes in $\mathcal{P}_2$ send a *composition-response* message, which includes $\mathcal{H}_2$, back to $\mathcal{P}_1$.

The rest of the protocol is symmetric to the first part; any process in $\mathcal{P}_1$ that receives the same composition response from a guild of $\mathcal{P}_2$ participates in a round of Byzantine consensus, this time within $\mathcal{P}_1$. This results in $\mathcal{P}_1$ sending a *composition-acknowledgment* message to $\mathcal{P}_2$, which now includes $\mathcal{H}_1$, the history of delivered messages in the instance run by $\mathcal{P}_1$. The histories $\mathcal{H}_1$ and $\mathcal{H}_2$ can be used by the composed system to calculate the initial state of the new protocol instance, presumably using a generic *merge function*.

The composition-acknowledgment message signals the start a new protocol instance run by $\mathcal{P}_1 \cup \mathcal{P}_2$. From this point on, processes use the composed fail-prone and quorum systems. Since $\mathbb{F}_2$ is known, processes in $\mathcal{P}_1$ can calculate both the tolerated system $\mathcal{T}_2$ of $\mathbb{Q}_2$ (in the simplest case by trying all possible failures of $\mathcal{P}_2$) and the purified version of $\mathbb{F}_2$, and vice versa for processes in $\mathcal{P}_2$. Should the fail-prone systems not be public, the processes could send them in the composition messages; however, privacy aspects are beyond the focus of this work.

6 Conclusions

Our work shows how trust assumptions of (possibly disjoint) systems can composed deterministically, such that groups of strangers may join each other and collaborate under a composed trust assumption with appealing properties. We present composition rules that work in both symmetric and asymmetric-trust models. Moreover, we overcome existing impossibility results for consensus among disjoint personal Byzantine quorum systems systems from the literature [12]; given two systems that can reach consensus on their own, our composition method results in a system that achieves consensus. As intermediate results we define the tolerated system of an ABQS, which reflects the overall resilience of the ABQS, and present a purification procedure, which aligns the expectations of a process with the realistic capabilities of an ABQS. We expect these contributions to be of independent interest towards a deeper understanding and practical adoption of subjective decentralized trust.

Acknowledgments

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

References

[1] O. Alpos and C. Cachin, "Consensus beyond thresholds: Generalized byzantine quorums made live," in *SRDS*, pp. 21–30, IEEE, 2020.

[2] L. Alvisi, E. T. Pierce, D. Malkhi, M. K. Reiter, and R. N. Wright, "Dynamic byzantine quorum systems," in *DSN*, pp. 283–292, IEEE Computer Society, 2000.

[3] E. Buchman, J. Kwon, and Z. Milosevic, "The latest gossip on BFT consensus," *CoRR*, vol. abs/1807.04938, 2018.

[4] C. Cachin and B. Tackmann, "Asymmetric distributed trust," in *OPODIS*, vol. 153 of *LIPIcs*, pp. 7:1–7:16, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[5] C. Cachin and L. Zanolini, "Asymmetric byzantine consensus," *CoRR*, vol. abs/2005.08795, 2020.

[6] M. Castro and B. Liskov, "Practical byzantine fault tolerance and proactive recovery," *ACM Trans. Comput. Syst.*, vol. 20, no. 4, pp. 398–461, 2002.

[7] I. Damgård, Y. Desmedt, M. Fitzi, and J. B. Nielsen, "Secure protocols with asymmetric trust," in *ASIACRYPT*, vol. 4833 of *Lecture Notes in Computer Science*, pp. 357–375, Springer, 2007.

[8] M. Hirt and U. M. Maurer, "Player simulation and general adversary structures in perfect multiparty computation," *J. Cryptol.*, vol. 13, no. 1, pp. 31–60, 2000.

[9] H. Howard, A. Charapko, and R. Mortier, "Fast flexible paxos: Relaxing quorum intersection for fast paxos," in *ICDCN*, pp. 186–190, ACM, 2021.

[10] S. Liu, P. Viotti, C. Cachin, V. Quéma, and M. Vukolic, "XFT: practical fault tolerance beyond crashes," in *OSDI*, pp. 485–500, USENIX Association, 2016.

[11] 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 *SOSP*, pp. 80–96, ACM, 2019.

[12] G. Losa, E. Gafni, and D. Mazières, "Stellar consensus by instantiation," in *DISC*, vol. 146 of *LIPIcs*, pp. 27:1–27:15, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[13] D. Malkhi, K. Nayak, and L. Ren, "Flexible byzantine fault tolerance," in *CCS*, pp. 1041–1053, ACM, 2019.

[14] D. Malkhi and M. K. Reiter, "Byzantine quorum systems," *Distributed Comput.*, vol. 11, no. 4, pp. 203–213, 1998.

[15] D. Malkhi, M. K. Reiter, and A. Wool, "The load and availability of byzantine quorum systems," *SIAM J. Comput.*, vol. 29, no. 6, pp. 1889–1906, 2000.

[16] D. Malkhi, M. K. Reiter, A. Wool, and R. N. Wright, "Probabilistic quorum systems," *Inf. Comput.*, vol. 170, no. 2, pp. 184–206, 2001.

[17] 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.

[18] M. Naor and A. Wool, "The load, capacity, and availability of quorum systems," *SIAM J. Comput.*, vol. 27, no. 2, pp. 423–447, 1998.

[19] M. Yin, D. Malkhi, M. K. Reiter, G. Golan-Gueta, and I. Abraham, "Hotstuff: BFT consensus with linearity and responsiveness," in *PODC*, pp. 347–356, ACM, 2019.