

Write Markers for Probabilistic Quorum Systems

Michael G. Merideth and Michael K. Reiter

November 2007. Revised November 2008.
CMU-ISRI-07-118R

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Supersedes Technical Report CMU-ISRI-07-118.
Also appears as Computer Science Technical Report CMU-CS-07-165R.

Abstract

Probabilistic quorum systems can tolerate a larger fraction of faults than can traditional (strict) quorum systems, while guaranteeing consistency with an arbitrarily high probability for a system with enough replicas. However, the masking and opaque types of probabilistic quorum systems are hampered in that their optimal load—a best-case measure of the work done by the busiest replica, and an indicator of scalability—is little better than that of strict quorum systems. In this paper we present a variant of probabilistic quorum systems that uses *write markers* in order to limit the extent to which Byzantine-faulty servers act together. Our masking and opaque probabilistic quorum systems have asymptotically better load than the bounds proven for previous masking and opaque quorum systems. Moreover, the new masking and opaque probabilistic quorum systems can tolerate an additional 24% and 17% of faulty replicas, respectively, compared with probabilistic quorum systems without write markers.

This work was partially supported by NSF grant CCF-0424422.

Keywords: Distributed systems, Byzantine fault tolerance, probabilistic quorum systems

1 Introduction

Given a universe U of servers, a *quorum system* over U is a collection $\mathcal{Q} = \{Q_1, \dots, Q_m\}$ such that each $Q_i \subseteq U$ and

$$|Q \cap Q'| > 0 \quad (1)$$

for all $Q, Q' \in \mathcal{Q}$. Each Q_i is called a *quorum*. The intersection property (1) makes quorums a useful primitive for coordinating actions in a distributed system. For example, if clients perform writes at a quorum of servers, then a client who reads from a quorum will observe the last written value. Because of their utility in such applications, quorums have a long history in distributed computing.

In systems that may suffer Byzantine faults [10], the intersection property (1) is typically not adequate as a mechanism to enable consistent data access. Because (1) requires only that the intersection of quorums be non-empty, it could be that two quorums intersect only in a single server, for example. In a system in which up to $b > 0$ servers might suffer Byzantine faults, this single server might be faulty and consequently, could fail to convey the last written value to a reader, for example.

For this reason, Malkhi and Reiter [15] proposed various ways of strengthening the intersection property (1) so as to enable quorums to be used in Byzantine environments. For example, an alternative to (1) is

$$|Q \cap Q' \setminus B| > |Q' \cap B| \quad (2)$$

for all $Q, Q' \in \mathcal{Q}$, where B is the (unknown) set of all (up to b) servers that are faulty. In other words, the intersection of any two quorums contains more non-faulty servers than the faulty ones in either quorum. As such, the responses from these non-faulty servers will outnumber those from faulty ones. These quorum systems are called *masking* systems.

Opaque quorum systems, have an even more stringent requirement as an alternative to (1):

$$|Q \cap Q' \setminus B| > |(Q' \cap B) \cup (Q' \setminus Q)| \quad (3)$$

for all $Q, Q' \in \mathcal{Q}$. In other words, the number of correct servers in the intersection of Q and Q' (i.e., $|Q \cap Q' \setminus B|$) exceeds the number of faulty servers in Q' (i.e., $|Q' \cap B|$) together with the number of servers in Q' but not Q . The rationale for this property can be seen by considering the servers in Q' but not Q as “outdated”, in the sense that if Q was used to perform an update to the system, then those servers in $Q' \setminus Q$ are unaware of the update. As such, if the faulty servers in Q' behave as the outdated ones do, their behavior (i.e., their responses) will dominate that from the correct servers in the intersection ($Q \cap Q' \setminus B$) unless (3) holds.

The increasingly stringent properties of Byzantine quorum systems come with costs in terms of the smallest system sizes that can be supported while tolerating a number b of faults [15]. This implies that a system with a fixed number of servers can tolerate fewer faults when the property is more stringent as seen in Table 1, which refers to the quorums just discussed as *strict*. Table 1 also shows the negative impact on the ability of the system to disperse load amongst the replicas, as discussed next.

Naor and Wool [24] introduced the notion of an *access strategy* by which clients select quorums to access. An access strategy $p : \mathcal{Q} \rightarrow [0, 1]$ is simply a probability distribution on quorums, i.e.,

$\sum_{Q \in \mathcal{Q}} p(Q) = 1$. Intuitively, when a client accesses the system, it does so at a quorum selected randomly according to the distribution p .

The formalization of an access strategy is useful as a tool for discussing the load-dispersing properties of quorums. The *load* [24] of a quorum system, $\mathcal{L}(\mathcal{Q})$, is the probability with which the busiest server is accessed in a client access, under the best possible access strategy p . As listed in Table 1, tight lower bounds have been proven for the load of each type of strict Byzantine quorum system. The load for opaque quorum systems is particularly unfortunate—systems that utilize opaque quorum systems cannot effectively disperse processing load across more servers (i.e., by increasing n) because the load is at least a constant. Byzantine quorum systems are used by many modern Byzantine-fault-tolerant protocols, e.g., [1, 4, 6, 9, 16, 19] in order to tolerate the arbitrary failure of a subset of their replicas. As such, circumventing the bounds is an important topic.

One way to circumvent these bounds is with *probabilistic quorum systems*. Probabilistic quorum systems relax the quorum intersection properties, asking them to hold only with high probability. More specifically, they relax (2) or (3), for example, to hold only with probability $1 - \epsilon$ (for ϵ , a small constant), where probabilities are taken with respect to the selection of quorums according to an access strategy p [18, 21]. This technique yields masking quorum constructions tolerating $b < 2.62/n$ and opaque quorum constructions tolerating $b < 3.15/n$ as seen in Table 1. These bounds hold in the sense that for any $\epsilon > 0$ there is an n_0 such that for all $n > n_0$, the required intersection property ((2) or (3) for masking and opaque quorum systems, respectively) holds with probability at least $1 - \epsilon$. Unfortunately, probabilistic quorum systems alone do not materially improve the load of Byzantine quorum systems.

In this paper, we present an additional modification, *write markers*, that improves on the bounds further. Intuitively, in each update access to a quorum of servers, a write marker is placed at the accessed servers in order to evidence the quorum used in that access. This write marker identifies the quorum used; as such, faulty servers not in this quorum cannot respond to subsequent quorum accesses as though they were.

As seen in Table 1, by using this method to constrain how faulty servers can collaborate, we show that probabilistic masking quorum systems with load $O(1/\sqrt{n})$ can be achieved, allowing the systems to disperse load independently of the value of b . Further, probabilistic opaque quorum systems with load $O(b/n)$ can be achieved, breaking the constant lower bound on load for opaque systems. Moreover, the resilience of probabilistic masking quorums can be improved an additional 24% to $b < n/2$, and the resilience of probabilistic opaque quorum systems can be improved an additional 17% to $b < n/2.62$.

The probability of error in probabilistic quorums requires mechanisms to ensure that accesses are performed according to the required access strategy p if the clients cannot be trusted to do so. Therefore, we adapt one such mechanism, the access-restriction protocol of probabilistic opaque quorum systems [21], to accommodate write markers. Thus, as a side benefit, our implementation forces faulty clients to follow the access strategy. With this, we provide a protocol to implement write markers that tolerates Byzantine clients.

Our primary contributions are (i) the identification and analysis of the benefits of write markers; and (ii) a proposed implementation of write markers that handles the complexities of tolerating Byzantine clients. Our analysis yields the following results:

Masking Quorums: We show that the use of write markers allows probabilistic masking quorum systems to tolerate up to $b < n/2$ faults when quorums are of size $\Omega(\sqrt{n})$. Setting all quorums to size $\rho\sqrt{n}$ for some constant ρ , we achieve a load that is asymptotically optimal for any quorum system, i.e., $\rho\sqrt{n}/n = O(1/\sqrt{n})$ [24].

This represents an improvement in load and the number of faults that can be tolerated. Probabilistic masking quorums without write markers can tolerate up to $b < n/2.62$ faults [21] and achieve load no better than $\Omega(b/n)$ [18]. In addition, the maximum number of faults that can be tolerated is tied to the size of quorums [18]. Thus, without write markers, achieving optimal load requires tolerating fewer faults. Strict masking quorum systems can tolerate (only) up to $b < n/4$ faults [15] and can achieve load $\Omega(\sqrt{b/n})$ [17].

Opaque Quorums: We show that the use of write markers allows probabilistic opaque quorum systems to tolerate up to $b < n/2.62$ faults. We present a construction with load $O(b/n)$ when $b = \Omega(\sqrt{n})$, thereby breaking the constant lower bound of $1/2$ on the load of strict opaque quorum systems [15]. Moreover, if $b = O(\sqrt{n})$, we can set all quorums to size $\rho\sqrt{n}$ for some constant ρ , in order to achieve a load that is asymptotically optimal for any quorum system, i.e., $\rho\sqrt{n}/n = O(1/\sqrt{n})$ [24].

This represents an improvement in load and the number of faults that can be tolerated. Probabilistic opaque quorum systems without write markers can tolerate (only) up to $b < n/3.15$ faults [21]. Strict opaque quorum systems can tolerate (only) up to $b < n/5$ faults [15]; these quorum systems can do no better than constant load even if $b = 0$ [15].

Table 1: Improvements due to write markers. (**Bold** entries are properties of particular constructions; others are lower bounds.)

| <i>Non-Byzantine:</i> | load | | faults | |
|-----------------------|-----------------------------------|--------|---------------------|--------|
| strict | $\Omega(1/\sqrt{n})$ | [24] | $< n$ | |
| <i>Masking:</i> | load | | faults | |
| strict | $\Omega(\sqrt{b/n})$ | [15] | $< n/4.00$ | [17] |
| probabilistic | $\Omega(b/n)$ | [18] | $< \mathbf{n/2.62}$ | [21] |
| write markers | $O(1/\sqrt{n})$ | [here] | $< \mathbf{n/2.00}$ | [here] |
| <i>Opaque:</i> | load | | faults | |
| strict | $\geq 1/2$ | [15] | $< n/5.00$ | [15] |
| probabilistic | unproven | | $< \mathbf{n/3.15}$ | [21] |
| write markers | $O(b/n)$ | [here] | $< \mathbf{n/2.62}$ | [here] |

2 Definitions and System Model

We assume a system with a set U of servers, $|U| = n$, and an arbitrary but bounded number of clients. Clients and servers can fail arbitrarily (i.e., Byzantine faults [10]). We assume that up to b servers can fail, and denote the set of faulty servers by B , where $B \subseteq U$. Any number of clients can fail. Failures are permanent. Clients and servers that do not fail are said to be *non-faulty*. We allow that faulty clients and servers may collude, and so we assume that faulty clients and servers all know the membership of B (although non-faulty clients and servers do not). However, for our implementation of write markers, as is typical for many Byzantine-fault-tolerant protocols

(c.f., [1, 4, 6, 19]), we assume that faulty clients and servers are computationally bounded such that they cannot subvert standard cryptographic primitives such as digital signatures.

Communication. Write markers require no communication assumptions beyond those of the probabilistic quorums for which they are used. For completeness, we summarize the model of [21], which is common to prior works in probabilistic [18] and signed [25] quorum systems: we assume that each non-faulty client can successfully communicate with each non-faulty server with high probability, and hence with all non-faulty servers with roughly equal probability. This assumption is in place to ensure that the network does not significantly bias a non-faulty client’s interactions with servers either toward faulty servers or toward different non-faulty servers than those with which another non-faulty client can interact. Put another way, we treat a server that can be reliably reached by none or only some non-faulty clients as a member of B .

Access set; access strategy; operation. We abstractly describe client operations as either *writes* that alter the state of the service or *reads* that do not. Informally, a non-faulty client performs a write to update the state of the service such that its value (or a later one) will be observed with high probability by any subsequent operation; a write thus successfully performed is called “established” (we define established more precisely below). A non-faulty client performs a read to obtain the value of the latest established write, where “latest” refers to the value of the most recent write preceding this read in a linearization [7] of the execution.

In the introduction, we discussed *access strategies* as probability distributions on quorums used for operations. For the remainder of the paper, we follow [21] in strictly generalizing the notion of access strategy to apply instead to *access sets* from which quorums are chosen. An access set is a set of servers from which the client selects a quorum. If the client is non-faulty, we assume that this selection is done uniformly at random. We adopt the access strategy that all access sets are chosen uniformly at random (even by faulty clients). In Section 4, we adapt a protocol to support write markers from one in [21] that approximately ensures this access strategy. Our analysis allows that access sets may be larger than quorums, though if access sets and quorums are of the same size, then our protocol effectively forces even faulty clients to select quorums uniformly at random as discussed in the introduction. In our analysis, all access sets used for reads and writes are of constant size a_{rd} and a_{wt} , respectively. All quorums used for reads and writes are of constant size q_{rd} and q_{wt} respectively.

Candidate; conflicting; error probability; established; participant; qualified; vote. Each write yields a corresponding *candidate* at some number of servers. A candidate is an abstraction used in part to ensure that two distinct write operations are distinguishable from each other, even if the corresponding data values are the same. A candidate is *established* once it is accepted by all of the non-faulty servers in some write quorum of size q_{wt} within the write access set of size a_{wt} . In opaque quorum systems, property (3) anticipates that different non-faulty servers each may hold a different candidate due to concurrent writes. A candidate that is characterized by the property that a non-faulty server would accept either it or a given established candidate, but not both, is called a *conflicting* candidate. Two candidates may conflict because, e.g., they both bear the same timestamp. In either masking or opaque quorum systems, a faulty server may try to forge a conflicting candidate. No non-faulty server accepts two candidates that conflict with each other.

A server can try to *vote* for some candidate (e.g., by responding to a read operation) if the server is a *participant* in voting (i.e., if the server is a member of the client’s read access set). However, a server becomes *qualified* to vote for a particular candidate only if the server is a member of the client’s write access set selected for the write operation for which it votes. Non-faulty clients wait for responses from a read quorum of size q_{rd} contained in the read access set of size a_{rd} . An *error* is said to occur in a read operation when a non-faulty client fails to observe the latest value or a faulty client obtains sufficiently many votes for a conflicting value.¹ The *error probability* is the probability of this occurring.

Behavior of faulty clients. We assume that faulty clients seek to maximize the error probability by following specific strategies [21]. This is a conservative assumption; a client cannot increase—but may decrease—the probability of error by failing to follow these strategies. At a high level, the strategies are as follows: a faulty client, which may be completely restricted in its choices: (i) when establishing a candidate, writes the candidate to as few non-faulty servers as possible to minimize the probability that it is observed by a non-faulty client; and (ii) writes a conflicting candidate to as many servers as will accept it (i.e., faulty servers plus, in the case of an opaque quorum system, any non-faulty server that has not accepted the established candidate) in order to maximize the probability that it is observed. However, it is important to note that a faulty client performing a write does not have knowledge of the read access set used by a non-faulty client [21].

3 Analysis of Write Markers

Intuitively, when a client submits a write, the candidate is associated with a write marker. We require that the following three properties are guaranteed by an implementation of write markers:

- W1. Every candidate has a write marker that identifies the access set chosen for the write;
- W2. A verifiable write marker implies that the access set was selected uniformly at random (i.e., according to the access strategy);
- W3. Every non-faulty client and server can verify a write marker.

When considering a candidate, non-faulty clients and servers verify the candidate’s write marker. Because of this verification, no non-faulty node will accept a vote for a candidate unless the issuing server is qualified to vote for the candidate.

Since each write access set is chosen uniformly at random (W2), the faulty servers that can vote for a candidate, i.e., the faulty qualified servers, are therefore a random subset of the faulty servers.

Thus, write markers remove the advantage enjoyed by faulty servers in strict and traditional-probabilistic masking and opaque quorum systems, where any faulty participant can vote for

Table 2: Ability of a server to vote for a given candidate: ● (traditional quorums); ★ (write markers).

| Type of server | Vote |
|--------------------------------------|------|
| Non-faulty qualified participant | ● ★ |
| Faulty qualified participant | ● ★ |
| Non-faulty non-qualified participant | |
| Faulty non-qualified participant | ● |

¹Faulty clients may be able to affect the system with such votes in some protocols [21].

any candidate—and therefore can collude to have a conflicting, potentially fabricated candidate chosen instead of an established candidate. This aspect of write markers is summarized in Table 2, which shows the impact of write markers in terms of the abilities of faulty and non-faulty servers to vote for a given candidate.

3.1 Consistency Constraints

Probabilistic quorum systems must satisfy constraints similar to those of strict quorum systems (e.g., (2), (3)), but only with probability $1 - \epsilon$. As with strict quorum systems, the purpose of these constraints is to guarantee that operations can be observed consistently in subsequent operations by receiving enough votes.

First, the constraints must ensure in expectation that a non-faulty client can observe the latest established candidate if such a candidate exists. Let Q_{rd} represent a read quorum chosen uniformly at random, i.e., a random variable, from a read access set itself chosen uniformly at random. (Think of this quorum as one used by a non-faulty client.) Let Q_{wt} represent a write quorum chosen by a potentially faulty client; Q_{wt} must be chosen from A_{wt} , an access set chosen uniformly at random. (Think of Q_{wt} as a quorum used for an established candidate.) Then the threshold r number of votes necessary to observe a value must be less than the expected number of non-faulty qualified participants, which is

$$\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|]. \quad (4)$$

The use of write markers has no impact here on (4) because $(Q_{rd} \cap Q_{wt}) \setminus B$ contains no faulty servers. However, write markers do enable us to set r smaller, as the following shows.

Second, the constraints must ensure that a conflicting candidate (which is in conflict with an established candidate as described in Section 2) is, in expectation, not observed by any client (non-faulty or faulty). In general, it is important for all clients to observe only established candidates so as to enable higher-level protocols (e.g., [1]) that employ repair phases that may affect the state of the system within a read [21]. Let A'_{rd} and A'_{wt} represent read and write access sets, respectively, chosen uniformly at random. (Think of A'_{wt} as the access set used by a faulty client for a conflicting candidate, and of A'_{rd} as the access set used by a faulty client for a read operation. How faulty clients can be forced to choose uniformly at random is described in Section 4.) We consider the cases for masking and opaque quorums separately:

Probabilistic Masking Quorums. In a masking quorum system, (2) dictates that only faulty servers may vote for a conflicting candidate. Using write markers, we require that the faulty qualified participants alone cannot produce sufficient votes for a candidate to be observed in expectation. Taking (4) into consideration, we require:

$$\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|] > \mathbb{E} [|(A'_{rd} \cap A'_{wt}) \cap B|]. \quad (5)$$

Contrast this with (2) and with the consistency requirement for traditional probabilistic masking quorum systems [18] (adapted to consider access sets), which requires that the faulty participants

(qualified or not) cannot produce sufficient votes for a candidate to be observed in expectation:

$$\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|] > \mathbb{E} [|A'_{rd} \cap B|]. \quad (6)$$

Intuitively, the intersection between access sets can be smaller with write markers because the right-hand side of (5) is less than the right-hand side of (6) if $a_{wt} < n$.

Probabilistic Opaque Quorums. With write markers, we have the benefit, described above for probabilistic masking quorums, in terms of the number of faulty participants that can vote for a candidate in expectation. However, as shown in (3), opaque quorum systems must additionally consider the maximum number of non-faulty qualified participants that vote for the same conflicting candidate in expectation. As such, instead of (5), we have:

$$\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|] > \mathbb{E} [(A'_{rd} \cap A'_{wt}) \cap B] + \mathbb{E} [|(A'_{rd} \cap A'_{wt}) \setminus B| \setminus Q_{wt}]. \quad (7)$$

Contrast this with the consistency requirement for traditional probabilistic opaque quorums [21]:

$$\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|] > \mathbb{E} [|A'_{rd} \cap B|] + \mathbb{E} [|(A'_{rd} \cap A'_{wt}) \setminus B| \setminus Q_{wt}]. \quad (8)$$

Again, intuitively, the intersection between access sets can be smaller with write markers because the right-hand side of (7) is less than the right-hand side of (8) if $a_{wt} < n$.

3.2 Implied Bounds

In this subsection, we are concerned with quorum systems for which we can achieve error probability (as defined in Section 2) no greater than a given ϵ for any n sufficiently large. For such quorum systems, there is an upper bound on b in terms of n , akin to the bound for strict quorum systems.

Intuitively, the maximum value of b is limited by the relevant constraint (i.e., either (5) or (7)). Of primary interest are Theorem 3.17 and its corollaries, which demonstrate the benefits of write markers for probabilistic masking quorum systems, and Theorem 3.23 and its corollaries, which demonstrate the benefits of write markers for probabilistic opaque quorum systems. They utilize Lemmas 3.5 and 3.6, which together present basic requirements for the types of quorum systems with which we are concerned.

The following theorem is a restatement of the Molloy and Reed statement [23, p. 172] of the McDiarmid Inequality that can be used to show that a random variable computed on a series of independent permutations is concentrated about its expectation.

Theorem 3.1 ([23]). *Let $Z = z(\Pi_1, \dots, \Pi_l)$ be a random variable that is a non-negative function of a series Π_1, \dots, Π_l of independent random variables, where each Π_i takes on a random permutation (bijection) $\pi : \{1, \dots, |P|\} \rightarrow P$ of a finite non-empty set P . Also, for some positive constants δ and μ , let the following conditions hold (where if $\Pi_j = \pi_j$ then the mapping $\langle i, j, m \rangle$ indicates that $\pi_j(i) = m$):*

M1. Swapping the mappings of any two elements in a single permutation π_j (i.e., changing $\{\langle i, j, m \rangle, \langle i', j, m' \rangle\}$ to $\{\langle i', j, m \rangle, \langle i, j, m' \rangle\}$, where $i \neq i'$ and $m \neq m'$) changes the value of Z by at most δ .

M2. If $Z = z(\pi_1, \dots, \pi_l) = x$, then there exists a set of at most μx distinct mappings $\{\langle i_1, j_1, m_1 \rangle, \dots, \langle i_{\mu x}, j_{\mu x}, m_{\mu x} \rangle\}$ such that $z(\pi'_1, \dots, \pi'_l) \geq x$ for any π'_1, \dots, π'_l sharing the same set of mappings.

If $0 \leq \Delta \leq \mathbb{E}[Z]$, then:

$$\Pr(|Z - \mathbb{E}[Z]| \geq \Delta + 60\delta\sqrt{\mu\mathbb{E}[Z]} + 1) \leq 4/e^{(\Delta^2/8\delta^2\mu\mathbb{E}[Z])}.$$

We simplify Theorem 3.1 to create Corollary 3.2 that deals with asymptotic bounds.

Corollary 3.2. *Let $Z = z(\Pi_1, \dots, \Pi_l)$ be a random variable that is a non-negative function of a series Π_1, \dots, Π_l of independent random variables, where each Π_i takes on a random permutation (bijection) $\pi : \{1, \dots, |P|\} \rightarrow P$ of a finite non-empty set P . Also, for some positive constants δ and μ , let the following conditions hold (where if $\Pi_j = \pi_j$ then the mapping $\langle i, j, m \rangle$ indicates that $\pi_j(i) = m$):*

M1. Swapping the mappings of any two elements in a single permutation π_j (i.e., changing $\{\langle i, j, m \rangle, \langle i', j, m' \rangle\}$ to $\{\langle i', j, m \rangle, \langle i, j, m' \rangle\}$, where $i \neq i'$ and $m \neq m'$) changes the value of Z by at most δ .

M2. If $Z = z(\pi_1, \dots, \pi_l) = x$, then there exists a set of at most μx distinct mappings $\{\langle i_1, j_1, m_1 \rangle, \dots, \langle i_{\mu x}, j_{\mu x}, m_{\mu x} \rangle\}$ such that $z(\pi'_1, \dots, \pi'_l) \geq x$ for any π'_1, \dots, π'_l sharing the same set of mappings.

If $\Delta = \omega(\sqrt{\mathbb{E}[Z]})$, then,²

$$\Pr(|Z - \mathbb{E}[Z]| \geq \Delta) = 2/e^{\omega(1)} \quad \text{as } \Delta \rightarrow \infty.$$

Proof of Corollary 3.2. Since $\Delta = \omega(\sqrt{\mathbb{E}[Z]})$, the $60\delta\sqrt{\mu\mathbb{E}[Z]} + 1$ term is negligible, and, for any constant $\beta < 1/8\delta^2\mu$ and large enough value of $\mathbb{E}[Z]$, we have (c.f., [23, p. 81]),

$$\Pr(|Z - \mathbb{E}[Z]| \geq \Delta) \leq 2/e^{(\beta\Delta^2/\mathbb{E}[Z])}.$$

In other words, if $\Delta = \omega(\sqrt{\mathbb{E}[Z]})$, then,

$$\Pr(|Z - \mathbb{E}[Z]| \geq \Delta) = 2/e^{\omega(1)} \quad \text{as } \Delta \rightarrow \infty. \quad \square$$

So that we can apply Corollary 3.2 to bound the error probability, we present a method for defining Q_{wt} , A'_{rd} , A_{wt} , A'_{wt} , and Q_{rd} (where Q_{rd} represents a read quorum selected by a non-faulty client) in terms of independent random variables Π_1, Π_2, Π_3 , and Π_4 , each taking on a random permutation $\{1, \dots, |U|\} \rightarrow U$, where U is the set of all n servers. Fix any set of b servers to constitute B . Then consider the following definitions:

² ω is the little-oh analog of Ω , i.e., $f(n) = \omega(g(n))$ if $f(n)/g(n) \rightarrow \infty$ as $n \rightarrow \infty$.

- Define $A_{\text{wt}} = \{\Pi_1(1), \dots, \Pi_1(a_{\text{wt}})\}$.
- Define $A'_{\text{wt}} = \{\Pi_2(1), \dots, \Pi_2(a_{\text{wt}})\}$.
- Define $A'_{\text{rd}} = \{\Pi_3(1), \dots, \Pi_3(a_{\text{rd}})\}$.
- Define $Q_{\text{rd}} = \{\Pi_4(1), \dots, \Pi_4(q_{\text{rd}})\}$.

Because each permutation is randomly selected (independently of B), so too are A_{wt} , A'_{wt} , A'_{rd} , and Q_{rd} . We define Q_{wt} in accordance with Section 2. Specifically, first we choose each $\Pi_1(i)$ such that $\Pi_1(i) \in A_{\text{wt}} \cap B$. Next, for each $j = 1..a_{\text{wt}}$, we choose $\Pi_1(j)$ if we have not yet chosen q_{wt} servers and $\Pi_1(j) \in A_{\text{wt}} \setminus (A'_{\text{wt}} \cup B)$. Finally, for each $k = 1..a_{\text{wt}}$, we choose $\Pi_1(k)$ if we have not yet chosen q_{wt} servers and $\Pi_1(k) \in A_{\text{wt}} \cap (A'_{\text{wt}} \setminus B)$.

Let r be the threshold, discussed in Section 3.1, for the number of votes necessary to observe a candidate. Define MinCorrect to be a random variable for the number of non-faulty servers with the established candidate, i.e., $\text{MinCorrect} = |(Q_{\text{rd}} \cap Q_{\text{wt}}) \setminus B|$ as indicated in (4).

Lemma 3.3. *Let $Z = \text{MinCorrect}$. Let $\mathbb{E}[Z] > r$ and $\mathbb{E}[Z] - r = \omega(\sqrt{\mathbb{E}[Z]})$. Then,*

$$\Pr(Z \leq r) = 2/e^{\omega(1)} \quad \text{as } \mathbb{E}[Z] - r \rightarrow \infty.$$

Proof of Lemma 3.3. For a fixed B , MinCorrect can be treated as a function of independent permutations (i.e., using Q_{rd} and Q_{wt} , as defined above). Consider this in relation to Corollary 3.2 as follows. Swapping any two elements in either Q_{rd} or Q_{wt} can change the value of MinCorrect by at most 1; therefore, $\delta = 1$ in Condition M1. Additionally, if $\text{MinCorrect} = x$, then the mappings

$$\bigcup_{u \in (Q_{\text{rd}} \cap Q_{\text{wt}}) \setminus B} \{ \langle \Pi_1^{-1}(u), 1, u \rangle, \langle \Pi_4^{-1}(u), 4, u \rangle \}$$

suffice to satisfy Condition M2; therefore, $\mu = 2$.

Let $\Delta = \mathbb{E}[Z] - r$; then, by assumption, $\Delta = \omega(\sqrt{\mathbb{E}[Z]})$. We apply Corollary 3.2, yielding,

$$\begin{aligned} & \Pr(Z \leq r) \\ &= \Pr(Z \leq \mathbb{E}[Z] - \Delta) \\ &= \Pr(\mathbb{E}[Z] - Z \geq \Delta) \\ &\leq \Pr(|Z - \mathbb{E}[Z]| \geq \Delta) \\ &= 2/e^{\omega(1)} \quad \text{as } \Delta \rightarrow \infty \\ &= 2/e^{\omega(1)} \quad \text{as } \mathbb{E}[Z] - r \rightarrow \infty. \end{aligned} \quad \square$$

Define MaxConflicting to be a random variable for the maximum number of servers that vote for a conflicting candidate. For example: due to (5), in masking quorums with write markers, $\text{MaxConflicting} = |(A'_{\text{rd}} \cap A'_{\text{wt}}) \cap B|$; and due to (7), in opaque quorums with write markers, $\text{MaxConflicting} = |(A'_{\text{rd}} \cap A'_{\text{wt}}) \cap B| + |((A'_{\text{rd}} \cap A'_{\text{wt}}) \setminus B) \setminus Q_{\text{wt}}|$.

Lemma 3.4. *Let $Z' = \text{MaxConflicting}$. Let $r > \mathbb{E}[Z']$ and $r - \mathbb{E}[Z'] = \omega(\sqrt{\mathbb{E}[Z']})$. Then,*

$$\Pr(Z' \geq r) = 2/e^{\omega(1)} \quad \text{as } r - \mathbb{E}[Z'] \rightarrow \infty.$$

Proof of Lemma 3.4. As seen in Sections 3.1 and 3.2, depending on the type of quorum system, MaxConflicting has a different definition. However, for each type of quorum system, for a fixed B , MaxConflicting can be treated as a function of independent permutations from the set $\{A'_{rd}, A'_{wt}, Q_{wt}\}$. Consider this in relation to Corollary 3.2 as follows. Swapping any two elements in one permutation can change the value of MaxConflicting by at most 1 because an additional server added to A_{rd} cannot be both faulty and non-faulty; therefore, $\delta = 1$ in Condition M1. Additionally, for any value of MaxConflicting considered in Section 3.1, if MaxConflicting = $|C| = x$, then the mappings

$$\bigcup_{u \in C} \{ \langle \Pi_1^{-1}(u), 1, u \rangle, \langle \Pi_2^{-1}(u), 2, u \rangle, \langle \Pi_3^{-1}(u), 3, u \rangle \}$$

suffice to satisfy Condition M2; therefore, $\mu \leq 3$.

Let $\Delta = r - \mathbb{E}[Z']$; then, by assumption, $\Delta = \omega(\sqrt{\mathbb{E}[Z']})$. We apply Corollary 3.2, yielding,

$$\begin{aligned} \Pr(Z' \geq r) &= \Pr(Z' \geq \Delta + \mathbb{E}[Z']) \\ &= \Pr(Z' - \mathbb{E}[Z'] \geq \Delta) \\ &\leq \Pr(|Z' - \mathbb{E}[Z']| \geq \Delta) \\ &= 2/e^{\omega(1)} \quad \text{as } \Delta \rightarrow \infty \\ &= 2/e^{\omega(1)} \quad \text{as } r - \mathbb{E}[Z'] \rightarrow \infty. \end{aligned} \quad \square$$

Lemma 3.5. *Let $n - b = \Omega(n)$. For all $c > 0$ there is a constant $d > 1$ such that for all q_{rd}, q_{wt} where $q_{rd}q_{wt} > dn$ and $q_{rd}q_{wt} - n = \Omega(1)$, it is the case that $\mathbb{E}[\text{MinCorrect}] > c$ for all n sufficiently large.*

Proof of Lemma 3.5. Note that,

$$\begin{aligned} \mathbb{E}[\text{MinCorrect}] &= q_{rd}q_{wt} \left(\frac{n-b}{n^2} \right) \\ &> dn \left(\frac{n-b}{n^2} \right) \\ &= d \left(\frac{n-b}{n} \right) \\ &= d \left(\frac{\Omega(n)}{n} \right) \\ &= d(\Omega(1)). \end{aligned} \quad \square$$

Lemma 3.6. *Let the following hold,*

$$\begin{aligned}\mathbb{E}[\text{MinCorrect}] - \mathbb{E}[\text{MaxConflicting}] &> 0, \\ \mathbb{E}[\text{MinCorrect}] - \mathbb{E}[\text{MaxConflicting}] &= \omega(\sqrt{\mathbb{E}[\text{MinCorrect}]}).\end{aligned}$$

Then it is possible to set r such that,

$$\text{error probability} \rightarrow 0 \quad \text{as } \mathbb{E}[\text{MinCorrect}] \rightarrow \infty.$$

Proof of Lemma 3.6. Set r as follows,

$$r = \frac{\mathbb{E}[\text{MinCorrect}] + \mathbb{E}[\text{MaxConflicting}]}{2}.$$

Then we can apply Lemma 3.3 to $\Pr(\text{MinCorrect} \leq r)$ because,

$$\begin{aligned}\mathbb{E}[\text{MinCorrect}] &> r, \text{ and} \\ \mathbb{E}[\text{MinCorrect}] - r &= \omega(\sqrt{\mathbb{E}[\text{MinCorrect}]}).\end{aligned}$$

Next, note that by assumption and our setting of r , $r - \mathbb{E}[\text{MaxConflicting}] = \omega(\sqrt{\mathbb{E}[\text{MinCorrect}]})$. But, since $\mathbb{E}[\text{MinCorrect}] - \mathbb{E}[\text{MaxConflicting}]$ grows when $\mathbb{E}[\text{MinCorrect}]$ grows, it must be that $\mathbb{E}[\text{MinCorrect}]$ grows faster than $\mathbb{E}[\text{MaxConflicting}]$. Therefore, $r - \mathbb{E}[\text{MaxConflicting}] = \omega(\sqrt{\mathbb{E}[\text{MaxConflicting}]})$. As such, we can apply Lemma 3.4 to $\Pr(\text{MaxConflicting} \geq r)$ because

$$\begin{aligned}r &> \mathbb{E}[\text{MaxConflicting}], \text{ and} \\ r - \mathbb{E}[\text{MaxConflicting}] &= \omega(\sqrt{\mathbb{E}[\text{MaxConflicting}]}).\end{aligned}$$

It is an error if $\text{MinCorrect} < r$ or $\text{MaxConflicting} \geq r$. Therefore, the error probability is bounded as follows:

$$\begin{aligned}\text{error probability} &= \Pr(\text{MaxConflicting} \geq r \vee \text{MinCorrect} < r) \\ &= \Pr(\text{MaxConflicting} \geq r) + \Pr(\text{MinCorrect} < r) - \\ &\quad \Pr(\text{MaxConflicting} \geq r \wedge \text{MinCorrect} < r) \\ &\leq \Pr(\text{MaxConflicting} \geq r) + \Pr(\text{MinCorrect} < r) \\ &\leq \Pr(\text{MaxConflicting} \geq r) + \Pr(\text{MinCorrect} \leq r) \\ &= 2/e^{\omega(1)} + 2/e^{\omega(1)} \\ &\quad \text{as } (\mathbb{E}[\text{MinCorrect}] - \mathbb{E}[\text{MaxConflicting}])/2 \rightarrow \infty \\ &= 2/e^{\omega(1)} + 2/e^{\omega(1)} \quad \text{as } \mathbb{E}[\text{MinCorrect}] \rightarrow \infty.\end{aligned}$$

Where the second-to-last line follows because,

$$(\mathbb{E}[\text{MinCorrect}] - \mathbb{E}[\text{MaxConflicting}])/2 = \mathbb{E}[\text{MinCorrect}] - r = r - \mathbb{E}[\text{MaxConflicting}].$$

And the final line follows because,

$$(\mathbb{E}[\text{MinCorrect}] - \mathbb{E}[\text{MaxConflicting}])/2 = \omega(\sqrt{\mathbb{E}[\text{MinCorrect}]}). \quad \square$$

Here and below, a suitable setting of r is one between $\mathbb{E}[\text{MinCorrect}]$ and $\mathbb{E}[\text{MaxConflicting}]$, inclusive. The remainder of the section is focused on determining, for each type of probabilistic quorum system, the upper bound on b and bounds on the load that Lemmas 3.5 and 3.6 imply. Important expected values, derived below, are as follows,

$$\mathbb{E}[|(Q_{rd} \cap Q_{wt}) \setminus B|] = \frac{q_{rd}(nq_{wt} - a_{wt}b)}{n^2}. \quad (9)$$

$$\mathbb{E}[|A'_{rd} \cap B|] = \frac{a_{rd}b}{n}. \quad (10)$$

$$\mathbb{E}[|(A'_{rd} \cap A'_{wt}) \cap B|] = \frac{a_{rd}a_{wt}b}{n^2}. \quad (11)$$

$$\mathbb{E}[|((A'_{rd} \cap A'_{wt}) \setminus B) \setminus Q_{wt}|] \leq \frac{a_{rd}}{n^3}(2a_{wt}n^2 - na_{wt}b - q_{wt}n^2 - a_{wt}^2n + a_{wt}^2b). \quad (12)$$

Lemma 3.7.

$$\mathbb{E}[|A'_{rd} \cap B|] = \frac{a_{rd}b}{n}.$$

Proof. A'_{rd} is selected independently of B . As such, $|A'_{rd} \cap B|$ is a hypergeometric random variable characterized by a_{rd} draws from a population of n elements containing b successes. Therefore, we use the formula for the expected value of a hypergeometric random variable. \square

Lemma 3.8.

$$\mathbb{E}[|(A'_{rd} \cap A'_{wt}) \cap B|] = \frac{a_{rd}a_{wt}b}{n^2}.$$

Proof. We calculate $\mathbb{E}[|(A'_{rd} \cap A'_{wt}) \cap B|]$ directly as follows. Consider an indicator random variable Ind_u , such that $\text{Ind}_u = 1$ if $u \in (A'_{rd} \cap A'_{wt}) \cap B$, and $\text{Ind}_u = 0$ otherwise. For each $u \in B$, we have $\Pr[\text{Ind}_u = 1] = \frac{a_{rd}a_{wt}}{n^2}$, since A_{rd} and A'_{wt} are chosen independently. By linearity of expectation:

$$\mathbb{E}[|(A'_{rd} \cap A'_{wt}) \cap B|] = \sum_{u \in B} \Pr(\text{Ind}_u = 1) = b \left(\frac{a_{rd}a_{wt}}{n^2} \right). \quad \square$$

In the proofs of the following lemmas, we use rules of conditional expectation (c.f., [22, Section 2.3]). In particular, the following.

Definition 3.9 ([22]). *The expression $\mathbb{E}[X \mid Y]$ is a random variable $f(Y)$ that takes on the value $\mathbb{E}[X \mid Y = y]$ when $Y = y$.*

Because $\mathbb{E}[X \mid Y]$ is a random variable, i.e., a function, it makes sense to consider its expectation.

Theorem 3.10 ([22, Theorem 2.7]).

$$\mathbb{E}[X] = \mathbb{E}[\mathbb{E}[X \mid Y]]. \quad (13)$$

Lemma 3.11.

$$\mathbb{E} [|\mathcal{Q}_{wt} \setminus B|] = \frac{q_{wt}n - a_{wt}b}{n}. \quad (14)$$

Proof. Let $\text{MalWrite} = |\mathcal{A}_{wt} \cap B|$. Since \mathcal{A}_{wt} is selected uniformly at random independently of B , MalWrite is a hypergeometric random variable, characterized by a_{wt} draws from a population of n elements containing b successes; therefore,

$$\mathbb{E} [\text{MalWrite}] = \frac{a_{wt}b}{n}.$$

Recall from Section 2, that a write is established once all of the non-faulty servers in any write quorum in \mathcal{A}_{wt} have accepted it. Therefore,

$$\mathbb{E} [|\mathcal{Q}_{wt} \setminus B| \mid \text{MalWrite} = m] = q_{wt} - m.$$

Applying Theorem 3.10 and linearity of expectation, we have that,

$$\begin{aligned} \mathbb{E} [|\mathcal{Q}_{wt} \setminus B|] &= \mathbb{E} [\mathbb{E} [|\mathcal{Q}_{wt} \setminus B| \mid \text{MalWrite}]] \\ &= \mathbb{E} [q_{wt} - \text{MalWrite}] \\ &= q_{wt} - \mathbb{E} [\text{MalWrite}] \\ &= q_{wt} - \frac{a_{wt}b}{n}. \end{aligned} \quad \square$$

Lemma 3.12.

$$\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|] = \frac{q_{rd}(nq_{wt} - a_{wt}b)}{n^2}.$$

Proof. Q_{rd} is independent of $Q_{wt} \setminus B$; therefore, $|(Q_{rd} \cap Q_{wt}) \setminus B| \mid |Q_{wt} \setminus B| = m$ is a conditional hypergeometric random variable characterized by q_{rd} draws from a population of n elements containing m successes, and,

$$\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B| \mid |Q_{wt} \setminus B| = m] = \frac{q_{rd}m}{n}.$$

Applying Theorem 3.10, by linearity of expectation and Lemma 3.11 we have that,

$$\begin{aligned} \mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|] &= \mathbb{E} [\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B| \mid |Q_{wt} \setminus B|]] \\ &= \mathbb{E} \left[\frac{q_{rd}}{n} (|Q_{wt} \setminus B|) \right] \\ &= \frac{q_{rd}}{n} \mathbb{E} [|\mathcal{Q}_{wt} \setminus B|] \\ &= \frac{q_{rd}(nq_{wt} - a_{wt}b)}{n^2}. \end{aligned} \quad \square$$

Lemma 3.13.

$$\mathbb{E}[(A'_{rd} \cap A'_{wt}) \setminus B] = a_{rd} \left(\frac{a_{wt}}{n} - \left(\frac{a_{wt}}{n} \right) \left(\frac{b}{n} \right) \right). \quad (15)$$

Proof. We calculate $\mathbb{E}[(A'_{rd} \cap A'_{wt}) \setminus B]$ directly as follows. Consider an indicator random variable Ind_u , such that $\text{Ind}_u = 1$ if $u \in (A'_{rd} \cap A'_{wt}) \setminus B$, and $\text{Ind}_u = 0$ otherwise. For each $u \in U \setminus B$, we have $\Pr[\text{Ind}_u = 1] = \frac{a_{rd}a_{wt}}{n^2}$, since A'_{rd} and A'_{wt} are chosen independently. By linearity of expectation:

$$\begin{aligned} \mathbb{E}[(A'_{rd} \cap A'_{wt}) \setminus B] &= \sum_{u \in U \setminus B} \Pr(\text{Ind}_u = 1) \\ &= (n - b) \left(\frac{a_{rd}a_{wt}}{n^2} \right) = a_{rd} \left(\frac{a_{wt}}{n} - \left(\frac{a_{wt}}{n} \right) \left(\frac{b}{n} \right) \right). \quad \square \end{aligned}$$

Lemma 3.14.

$$\mathbb{E}[(A_{rd} \cap A'_{wt}) \cap (Q_{wt} \setminus B)] \geq \frac{a_{rd}}{n} \left(q_{wt} - \frac{a_{wt}}{n} \left(b + \frac{(n - a_{wt})(n - b)}{n} \right) \right). \quad (16)$$

Proof. In calculating $\mathbb{E}[(A_{rd} \cap A'_{wt}) \cap (Q_{wt} \setminus B)]$, because a faulty client may perform *both* a write that becomes established at Q_{wt} and another write at A'_{wt} that conflicts with the first write, we cannot assume that Q_{wt} is selected independently of A'_{wt} . Let $\text{Cl} = A'_{wt} \cap (Q_{wt} \setminus B)$. Then, in particular, as described in Section 2, such a client seeks to maximize $\mathbb{E}[\text{MaxConflicting}]$, and therefore minimizes $\mathbb{E}[|\text{Cl}|]$ by choosing the servers for $Q_{wt} \setminus B$ from $A_{wt} \setminus (A'_{wt} \cup B)$ first. Thus,

$$\begin{aligned} \mathbb{E}[|\text{Cl}|] &= \mathbf{max}(0, \mathbb{E}[|Q_{wt} \setminus B|] - \mathbb{E}[|A_{wt} \setminus (A'_{wt} \cup B)|]) \\ &\geq \mathbb{E}[|Q_{wt} \setminus B|] - \mathbb{E}[|A_{wt} \setminus (A'_{wt} \cup B)|]. \end{aligned} \quad (17)$$

We calculate $\mathbb{E}[|A_{wt} \setminus (A'_{wt} \cup B)|]$ directly as follows. For clarity, note that $A_{wt} \setminus (A'_{wt} \cup B) = (A_{wt} \setminus A'_{wt}) \setminus B$. Consider an indicator random variable Ind_u , such that $\text{Ind}_u = 1$ if $u \in (A_{wt} \setminus A'_{wt}) \setminus B$, and $\text{Ind}_u = 0$ otherwise. For each $u \in U \setminus B$, we have $\Pr[\text{Ind}_u = 1] = \frac{a_{wt}(n - a_{wt})}{n^2}$, since A'_{wt} and A_{wt} are independent. By linearity of expectation:

$$\mathbb{E}[|A_{wt} \setminus (A'_{wt} \cup B)|] = \sum_{u \in U \setminus B} \Pr(\text{Ind}_u = 1) = (n - b) \left(\frac{a_{wt}(n - a_{wt})}{n^2} \right) = \frac{a_{wt}}{n^2} (n - a_{wt})(n - b). \quad (18)$$

By (17), (14), and (18), we have that,

$$\mathbb{E}[|\text{Cl}|] \geq \left(q_{wt} - \frac{a_{wt}b}{n} \right) - \frac{a_{wt}}{n^2} (n - a_{wt})(n - b) = q_{wt} - \frac{a_{wt}}{n} \left(b + \frac{(n - a_{wt})(n - b)}{n} \right). \quad (19)$$

Since A_{rd} is independent of Cl , we see that $|(A_{rd} \cap A'_{wt}) \cap (Q_{wt} \setminus B)| \mid |\text{Cl}| = c'$ is a hypergeometric random variable characterized by a_{rd} draws from a population of n elements containing c' successes, and,

$$\mathbb{E}[|(A_{rd} \cap A'_{wt}) \cap (Q_{wt} \setminus B)| \mid |\text{Cl}| = c'] = \frac{a_{rd}c'}{n}.$$

Applying Theorem 3.10, by linearity of expectation we have that,

$$\begin{aligned}
& \mathbb{E} [|(A_{rd} \cap A'_{wt}) \cap (Q_{wt} \setminus B)|] \\
&= \mathbb{E} [\mathbb{E} [|(A_{rd} \cap A'_{wt}) \cap (Q_{wt} \setminus B)| \mid |C|]] \\
&= \mathbb{E} \left[\frac{a_{rd}}{n} |C| \right] = \frac{a_{rd}}{n} \mathbb{E} [|C|].
\end{aligned} \tag{20}$$

Therefore, by (19) and (20) we have that,

$$\mathbb{E} [|(A_{rd} \cap A'_{wt}) \cap (Q_{wt} \setminus B)|] \geq \frac{a_{rd}}{n} \left(q_{wt} - \frac{a_{wt}}{n} \left(b + \frac{(n - a_{wt})(n - b)}{n} \right) \right). \quad \square$$

Lemma 3.15.

$$\mathbb{E} [|((A'_{rd} \cap A'_{wt}) \setminus B) \setminus Q_{wt}|] \leq \frac{a_{rd}}{n^3} (2a_{wt}n^2 - na_{wt}b - q_{wt}n^2 - a_{wt}^2n + a_{wt}^2b).$$

Proof. To calculate $\mathbb{E} [|((A'_{rd} \cap A'_{wt}) \setminus B) \setminus Q_{wt}|]$, first note that:

$$\begin{aligned}
& \mathbb{E} [|((A'_{rd} \cap A'_{wt}) \setminus B) \setminus Q_{wt}|] \\
&= \mathbb{E} [|(A'_{rd} \cap A'_{wt}) \setminus B|] - \mathbb{E} [|(A'_{rd} \cap A'_{wt}) \cap (Q_{wt} \setminus B)|].
\end{aligned} \tag{21}$$

Combining and simplifying equations (15), (16), and (21), we obtain (12). \square

Lemma 3.16. *For a probabilistic masking or opaque quorum system configuration with or without write markers,³*

$$\begin{aligned}
& \mathbb{E} [\text{MinCorrect}] > \mathbb{E} [\text{MaxConflicting}] \Rightarrow \\
& \mathbb{E} [\text{MinCorrect}] - \mathbb{E} [\text{MaxConflicting}] = \omega(\sqrt{\mathbb{E} [\text{MinCorrect}]}).
\end{aligned}$$

Proof of Lemma 3.16. Consider:

- **Masking without write markers:** $\mathbb{E} [\text{MinCorrect}] > \mathbb{E} [\text{MaxConflicting}]$ implies $\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|] - \mathbb{E} [|A'_{rd} \cap B|] \neq 0$, and so (9) and (10) show us that $\mathbb{E} [\text{MinCorrect}] - \mathbb{E} [\text{MaxConflicting}] = \omega(\sqrt{\mathbb{E} [\text{MinCorrect}]}).$
- **Masking with write markers:** $\mathbb{E} [\text{MinCorrect}] > \mathbb{E} [\text{MaxConflicting}]$ implies $\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|] - \mathbb{E} [|(A'_{rd} \cap A'_{wt}) \cap B|] \neq 0$, and so (9) and (11) show us that $\mathbb{E} [\text{MinCorrect}] - \mathbb{E} [\text{MaxConflicting}] = \omega(\sqrt{\mathbb{E} [\text{MinCorrect}]}).$
- **Opaque without write markers:** $\mathbb{E} [\text{MinCorrect}] > \mathbb{E} [\text{MaxConflicting}]$ implies $\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|] - (\mathbb{E} [|A'_{rd} \cap B|] + \mathbb{E} [|((A'_{rd} \cap A'_{wt}) \setminus B) \setminus Q_{wt}|]) \neq 0$, and so (9), (10), and (12) show us that $\mathbb{E} [\text{MinCorrect}] - \mathbb{E} [\text{MaxConflicting}] = \omega(\sqrt{\mathbb{E} [\text{MinCorrect}]}).$

³Though not shown here, Lemma 3.16 also applies trivially to dissemination quorum systems.

- Opaque with write markers: $\mathbb{E} [\text{MinCorrect}] > \mathbb{E} [\text{MaxConflicting}]$ implies $\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|] - (\mathbb{E} [|(A'_{rd} \cap A'_{wt}) \cap B|] + \mathbb{E} [|(A'_{rd} \cap A'_{wt}) \setminus B| \setminus Q_{wt}|]) \neq 0$, and so (9), (11), and (12) show us that $\mathbb{E} [\text{MinCorrect}] - \mathbb{E} [\text{MaxConflicting}] = \omega(\sqrt{\mathbb{E} [\text{MinCorrect}]})$. \square

Theorem 3.17. *For all ϵ there is a constant $d > 1$ such that for all q_{rd}, q_{wt} where $q_{rd}q_{wt} > dn$, $q_{rd}q_{wt} - n = \Omega(1)$, and*

$$b < \frac{q_{rd}q_{wt}n}{q_{rd}a_{wt} + a_{rd}a_{wt}},$$

any such probabilistic masking quorum system employing write markers achieves error probability no greater than ϵ given a suitable setting of r for all n sufficiently large.

Proof of Theorem 3.17. We show that Lemmas 3.5 and 3.6 apply. By (9) and (11),

$$\begin{aligned} b &< \frac{q_{rd}q_{wt}n}{q_{rd}a_{wt} + a_{rd}a_{wt}} \\ &\Leftrightarrow \mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|] > \mathbb{E} [|(A'_{rd} \cap A'_{wt}) \cap B|] \\ &\Leftrightarrow \mathbb{E} [\text{MinCorrect}] > \mathbb{E} [\text{MaxConflicting}]. \end{aligned}$$

Therefore, we can apply Lemma 3.16 and, thus, Lemma 3.6. By the restriction on b and the conditions of the Theorem, we can also apply Lemma 3.5. \square

Corollary 3.18. *Let $a_{rd} = q_{rd}$ and $a_{wt} = q_{wt}$. For all ϵ there is a constant $d > 1$ such that for all q_{rd}, q_{wt} where $q_{rd}q_{wt} > dn$, $q_{rd}q_{wt} - n = \Omega(1)$, and*

$$b < n/2,$$

any such probabilistic masking quorum system employing write markers achieves error probability no greater than ϵ given a suitable setting of r for all n sufficiently large.

In other words, with write markers, the size of quorums does not impact the maximum fraction of faults that can be tolerated when quorums are selected uniformly at random (i.e., when $a_{rd} = q_{rd}$ and $a_{wt} = q_{wt}$).

Corollary 3.19. *Let $a_{rd} = q_{rd}$, $a_{wt} = q_{wt}$, and $b < n/2$. For all ϵ there is a constant $\rho > 1$ such that if $q_{rd} = q_{wt} = \rho\sqrt{n}$, any such probabilistic masking quorum system employing write markers achieves error probability no greater than ϵ given a suitable setting of r for all n sufficiently large, and has load*

$$\rho\sqrt{n}/n = O(1/\sqrt{n}).$$

Theorem 3.20. *For all ϵ there is a constant $d > 1$ such that for all q_{rd}, q_{wt} where $q_{rd}q_{wt} > dn$, $q_{rd}q_{wt} - n = \Omega(1)$, and*

$$b < \frac{q_{rd}q_{wt}n}{q_{rd}a_{wt} + a_{rd}n},$$

any such probabilistic masking quorum system without write markers achieves error probability no greater than ϵ given a suitable setting of r for all n sufficiently large.

Proof of Theorem 3.20. We show that Lemmas 3.5 and 3.6 apply. By (9) and (10),

$$\begin{aligned} b &< \frac{q_{rd}q_{wt}n}{q_{rd}a_{wt} + a_{rd}n} \\ &\Leftrightarrow \mathbb{E} [|(\mathbf{Q}_{rd} \cap \mathbf{Q}_{wt}) \setminus B|] > \mathbb{E} [|A'_{rd} \cap B|] \\ &\Leftrightarrow \mathbb{E} [\text{MinCorrect}] > \mathbb{E} [\text{MaxConflicting}]. \end{aligned}$$

Therefore, we can apply Lemma 3.16 and, thus, Lemma 3.6. By the restriction on b and the conditions of the Theorem, we can also apply Lemma 3.5. \square

Corollary 3.21. *Let $a_{rd} = q_{rd}$ and $a_{wt} = q_{wt}$. For all ϵ there is a constant $d > 1$ such that for all q_{rd}, q_{wt} where $q_{rd}q_{wt} > dn$, $q_{rd}q_{wt} - n = \Omega(1)$, and*

$$b < \frac{q_{wt}n}{q_{wt} + n},$$

any such probabilistic masking quorum system without write markers achieves error probability no greater than ϵ given a suitable setting of r for all n sufficiently large.

Corollary 3.22. *Let $a_{rd} = q_{rd}$ and $a_{wt} = q_{wt} = n - b$. For all ϵ there is a constant $d > 1$ such that for all q_{rd}, q_{wt} where $q_{rd}q_{wt} > dn$, $q_{rd}q_{wt} - n = \Omega(1)$, and*

$$b < n/2.62,$$

any such probabilistic masking quorum system without write markers achieves error probability no greater than ϵ given a suitable setting of r for all n sufficiently large.

Theorem 3.23. *For all ϵ there is a constant $d > 1$ such that for all q_{rd}, q_{wt} where $q_{rd}q_{wt} > dn$, $q_{rd}q_{wt} - n = \Omega(1)$, and*

$$b < \frac{n(a_{rd}a_{wt}^2 + a_{rd}q_{wt}n + q_{rd}q_{wt}n - 2a_{rd}a_{wt}n)}{a_{wt}(a_{rd}a_{wt} + q_{rd}n)},$$

any such probabilistic opaque quorum system employing write markers achieves error probability no greater than ϵ given a suitable setting of r for all n sufficiently large.

Proof of Theorem 3.23. We show that Lemmas 3.5 and 3.6 apply. By (9), (11), and (12),

$$\begin{aligned} b &< \frac{n(a_{rd}a_{wt}^2 + a_{rd}q_{wt}n + q_{rd}q_{wt}n - 2a_{rd}a_{wt}n)}{a_{wt}(a_{rd}a_{wt} + q_{rd}n)} \\ &\Rightarrow \mathbb{E} [|(\mathbf{Q}_{rd} \cap \mathbf{Q}_{wt}) \setminus B|] > \mathbb{E} [|(A'_{rd} \cap A'_{wt}) \cap B|] + \mathbb{E} [|((A'_{rd} \cap A'_{wt}) \setminus B) \setminus \mathbf{Q}_{wt}|] \\ &\Leftrightarrow \mathbb{E} [\text{MinCorrect}] > \mathbb{E} [\text{MaxConflicting}]. \end{aligned}$$

Therefore, we can apply Lemma 3.16 and, thus, Lemma 3.6. By the restriction on b and the conditions of the Theorem, we can also apply Lemma 3.5. \square

Corollary 3.24. Let $a_{rd} = q_{rd}$ and $a_{wt} = q_{wt}$. For all ϵ there is a constant $d > 1$ such that for all q_{rd}, q_{wt} where $q_{rd}q_{wt} > dn$, $q_{rd}q_{wt} - n = \Omega(1)$, and

$$b < \frac{q_{wt}n}{q_{wt} + n},$$

any such probabilistic opaque quorum system employing write markers achieves error probability no greater than ϵ given a suitable setting of r for all n sufficiently large.

Comparing Corollary 3.24 with Corollary 3.18, we see that in the opaque quorum case q_{wt} cannot be set independently of b .

Corollary 3.25. Let $a_{rd} = q_{rd}$, $a_{wt} = q_{wt}$, and $b < (q_{wt}n)/(q_{wt} + n)$. For all ϵ there is a constant $d > 1$ such that for all q_{rd}, q_{wt} where $q_{rd}q_{wt} > dn$ and $q_{rd}q_{wt} - n = \Omega(1)$, any such probabilistic opaque quorum system employing write markers achieves error probability no greater than ϵ given a suitable setting of r for all n sufficiently large, and has load

$$\Omega(b/n).$$

Corollary 3.26. Let $b = \Omega(\sqrt{n})$. For all ϵ there is a constant $d > 1$ such that for all $a_{rd}, a_{wt}, q_{rd}, q_{wt}$ where $a_{rd} = a_{wt} = q_{rd} = q_{wt} = lb$ for a value l such that $c' \geq l > n/(n - b)$ for some constant c' , $(lb)^2 > dn$ and $(lb)^2 - n = \Omega(1)$, any such probabilistic opaque quorum system employing write markers achieves error probability no greater than ϵ given a suitable setting of r for all n sufficiently large, and has load

$$O(b/n).$$

Corollary 3.27. Let $a_{rd} = q_{rd}$ and $a_{wt} = q_{wt} = n - b$. For all ϵ there is a constant $d > 1$ such that for all q_{rd}, q_{wt} where $q_{rd}q_{wt} > dn$, $q_{rd}q_{wt} - n = \Omega(1)$, and

$$b < n/2.62,$$

any such probabilistic opaque quorum system employing write markers achieves error probability no greater than ϵ given a suitable setting of r for all n sufficiently large.

Theorem 3.28. For all ϵ there is a constant $d > 1$ such that for all q_{rd}, q_{wt} where $q_{rd}q_{wt} > dn$, $q_{rd}q_{wt} - n = \Omega(1)$, and

$$b < \frac{n(a_{rd}q_{wt}n - 2a_{rd}a_{wt}n + a_{wt}^2a_{rd} + q_{rd}q_{wt}n)}{n^2a_{rd} - a_{rd}a_{wt}n + a_{wt}^2a_{rd} + q_{rd}a_{wt}n},$$

any such probabilistic opaque quorum system without write markers achieves error probability no greater than ϵ given a suitable setting of r for all n sufficiently large.

Proof of Theorem 3.28. We show that Lemmas 3.5 and 3.6 apply. By (9), (10), and (12),

$$\begin{aligned} b &< \frac{(a_{rd}q_{wt}n - 2a_{rd}a_{wt}n + a_{wt}^2a_{rd} + q_{rd}q_{wt}n)n}{n^2a_{rd} - a_{rd}a_{wt}n + a_{wt}^2a_{rd} + q_{rd}a_{wt}n} \\ &\Rightarrow \mathbb{E}[|(Q_{rd} \cap Q_{wt}) \setminus B|] > \mathbb{E}[|A'_{rd} \cap B|] + \mathbb{E}[|((A'_{rd} \cap A'_{wt}) \setminus B) \setminus Q_{wt}|] \\ &\Leftrightarrow \mathbb{E}[\text{MinCorrect}] > \mathbb{E}[\text{MaxConflicting}]. \end{aligned}$$

Therefore, we can apply Lemma 3.16 and, thus, Lemma 3.6. By the restriction on b and the conditions of the Theorem, we can also apply Lemma 3.5. \square

Corollary 3.29. *Let $a_{rd} = q_{rd}$ and $a_{wt} = q_{wt}$. For all ϵ there is a constant $d > 1$ such that for all q_{rd}, q_{wt} where $q_{rd}q_{wt} > dn$, $q_{rd}q_{wt} - n = \Omega(1)$, and*

$$b < \frac{q_{wt}^2 n}{q_{wt}^2 + n^2},$$

any such probabilistic opaque quorum system without write markers achieves error probability no greater than ϵ given a suitable setting of r for all n sufficiently large.

Corollary 3.30. *Let $a_{rd} = q_{rd}$ and $a_{wt} = q_{wt} = n - b$. For all ϵ there is a constant $d > 1$ such that for all q_{rd}, q_{wt} where $q_{rd}q_{wt} > dn$, $q_{rd}q_{wt} - n = \Omega(1)$, and*

$$b < n/3.15,$$

any such probabilistic opaque quorum system without write markers achieves error probability no greater than ϵ given a suitable setting of r for all n sufficiently large.

4 Implementation

In this section, we provide an implementation of write markers that achieves the behavior assumed in Section 3, even with Byzantine clients. Specifically, it ensures properties W1–W3. (Though, technically, it ensures W2 only approximately in the case of opaque quorum systems, in which, as we explain below, a faulty server might be able to create a conflicting candidate using a write marker for a stale, i.e., out-of-date, access set—but to no advantage.)

Because clients may be faulty, we cannot rely on, e.g., digital signatures issued by them to implement write markers. Instead, we adapt mechanisms of our access-restriction protocol for probabilistic opaque quorum systems [21]. The access-restriction protocol is designed to ensure that all clients follow the access strategy. It already enables non-faulty *servers* to verify this before accepting a write. And, since it is the only way of which we are aware for a probabilistic quorum system to tolerate Byzantine clients when write markers are of benefit (i.e., when the sizes of write access sets are restricted), its mechanisms are appropriate.

The relevant parts of the preexisting protocol work as follows [21]. From a pre-configured number of servers, a client obtains a *verifiable recent value* (VRV), the value of which is unpredictable to clients and b or fewer servers prior to its creation. This VRV is used to generate a pseudorandom sequence of access sets. Since a VRV can be verified using only public information, both it and the sequence of access sets it induces can be verified by clients and servers. Non-faulty clients choose the next unused access set for each operation. However, a faulty client is motivated to maximize the probability of error. If the use of the next access set in the sequence does not maximize the probability of error given the current state of the system (i.e., the candidates accepted by the servers), such a client may try to skip ahead some number of access sets. Alternatively, such a client might try to wait to use the next access set until the state of the system changes. If allowed to follow either strategy, such a client would circumvent the access strategy because its choice of access set would not be independent from the state of the system.

Three mechanisms are used together to coerce a faulty client to follow the access strategy. First, the client must perform exponentially increasing work in expectation in order to use later access sets. As such, a client requires exponentially increasing time in expectation in order to choose an access set. This is implemented by requiring that the client solve a client puzzle [8] of the appropriate difficulty. The solution to the puzzle is, in expectation, difficult to find but easy to verify. Second, the VRV and sequence of access sets become invalid as the non-faulty servers accept additional candidates, or as the system otherwise progresses (e.g., as time passes). Non-faulty servers verify that an access set is still valid, i.e., not stale, before accepting it. Thus, system progress forces the client to start its work anew, and, as such, makes the work solving the puzzle for any unused access set wasted. Finally, during the time that the client is working, the established candidate propagates in the background to the non-faulty servers that are non-qualified (c.f., [14]). This decreases the window of vulnerability in which a given access set in the sequence is useful for a conflicting write by making non-qualified servers aware that (i) there is an established candidate (so that they will not accept a conflicting candidate) and (ii) that the state of the system has progressed (so that they will invalidate the current VRV if appropriate).

The impact of these three mechanisms is that a non-faulty *server* can be confident that the choice of write access set adheres (at least approximately) to the access strategy upon having verified that the access set is valid, current, and is accompanied by an appropriate puzzle solution.

| Masking write | | | |
|---------------|--------------------------------------|---|----------------------|
| α | access set solution data value | β promise | γ certificate |
| | | | δ status |
| Opaque write | | | |
| a | access set solution data value | b status | |
| Read | | | |
| i | query | ii data value certificate (masking) access set, solution (opaque) | |

Figure 2: Message types. (Write marker emphasized with gray.)

date to the client during any read operation. If the server is non-faulty—a fact of which a non-faulty client cannot be certain—the access set used for the operation was indeed chosen according to the access strategy because the server performed verification before accepting the candidate. However, because the server may be faulty, the client performs verification as well; it verifies the write marker and that the server is a member of the access set. This allows us to guarantee points W1–




Figure 1: Read operation with write markers: messages and stages of verification of access set. (Changes in gray.)

For write markers, we extend the protocol so that, as seen in Figure 1, *clients* can also perform verification. This requires that information about the puzzle solution and access set (including the VRV used to generate it) be returned by the servers to clients. (As seen in Figure 2 and explained below, this information varies across masking and opaque quorum systems.) In the preexisting access-restriction protocol, this information is verified and discarded by each server. For write markers, this information is instead stored by each server in the verification stage as a write marker. It is sent along with the data value as part of the candidate to the client during any read operation.

W3. As such, faulty non-qualified servers are unable to vote for the candidates for which qualified servers can vote.

Figures 1, 2, 3, and 4 illustrate relevant pieces of the preexisting protocol and our modifications for write markers in the context of read and write operations in probabilistic masking and opaque quorum systems. The figures highlight that the additions to the protocol for write markers involve saving the write markers and returning them to clients so that clients can also verify them.

The differences in the structure of the write marker for probabilistic opaque and masking quorum systems mentioned above results in subtly different guarantees. The remainder of the section discusses these details.

4.1 Probabilistic Opaque Quorums

As seen in Figure 2 (message *ii*), a write marker for a probabilistic opaque quorum system consists of the write-access-set identifier (including the VRV) and the solution to the puzzle that unlocks the use of this access set. Unlike a non-faulty server that verifies the access set at the time of use, a non-faulty client cannot verify that an access set was not already stale when the access set was accepted by a faulty server. Initially, this may appear problematic because it is clear that, given sufficient time, a faulty client will eventually be able to solve the puzzle for its preferred access set to use for a conflicting write—this access set may contain all of the servers in B . In addition, the faulty client can delay the use of this access set because non-faulty clients will be unable to verify whether it was already stale when it was used.

Fortunately, because non-faulty servers will not accept a stale candidate (i.e., a candidate accompanied by a stale access set), the fact that a stale access set may be accepted by a faulty server does not impact the benefit of write markers for opaque quorum systems. In general, consistency requires (7), i.e.,

$$\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|] > \mathbb{E} [|(A'_{rd} \cap A'_{wt}) \cap B|] + \mathbb{E} [|((A'_{rd} \cap A'_{wt}) \setminus B) \setminus Q_{wt}|] .$$

However, only faulty servers will accept a stale candidate. Therefore, if the candidate was stale when written to A'_{wt} , no non-faulty server would have accepted it. Thus, in this case, the consistency constraint is equivalent to,

$$\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|] > \mathbb{E} [|(A'_{rd} \cap A'_{wt}) \cap B|] .$$

Even if the access set contains all of the faulty servers, i.e., $B \subset A'_{wt}$, then this becomes,

$$\mathbb{E} [|(Q_{rd} \cap Q_{wt}) \setminus B|] > \mathbb{E} [|A'_{rd} \cap B|] .$$




Figure 3: Write operation in opaque quorum systems: messages and stages of verification of write marker. (Changes in gray.)

However, this is (6), the constraint on probabilistic masking quorum systems without write markers. In effect, a faulty client must either: (i) use a recent access set that is therefore chosen approximately uniformly at random, and be limited by (7); or (ii), use a stale access set and be limited by (6). If quorums are the sizes of access sets, both inequalities have the same upper bound on b as seen in Corollary 3.21 and Corollary 3.24; otherwise, a faulty client is disadvantaged by using a stale access set because a system that satisfies (6) can tolerate more faults than one that satisfies (7), and is therefore less likely to result in error. (Compare the bounds in Theorem 3.20 and Theorem 3.23.)

4.2 Probabilistic Masking Quorums

Protocols for masking quorum systems involve an additional round of communication (an echo phase, c.f., [16] or broadcast phase, c.f., [20]) during write operations in order to tolerate Byzantine or concurrent clients. This round prevents non-faulty servers from accepting conflicting data values, as assumed by (2). In order to write a data value, a client must first obtain a *write certificate* (a quorum of replies that together attest that the non-faulty servers will accept no conflicting data value). In contrast to optimistic protocols that use opaque quorum systems, these protocols are pessimistic.

This additional round allows us to prevent clients from using stale access sets. Specifically, in the request to authorize a data value (message α in Figure 2 and Figure 4), the client sends the access set identifier (including the VRV), the solution to the puzzle enabling use of this access set, and the data value. We require that the certificate come from servers in the access set that is chosen for the write operation. Each server verifies the VRV and that the puzzle solution enables use of the indicated access set before returning authorization (message β in Figure 2 and Figure 4). The non-faulty servers that contribute to the certificate all implicitly agree that the access set is not stale, for otherwise they would not agree to the write. This certificate (sent to each server in message γ in Figure 2 and Figure 4) is stored along with the data value as a write marker. Thus, unlike in probabilistic opaque quorum systems, a verifiable write marker in a probabilistic masking quorum system implies that a stale access set was not used. The reading client verifies the certificate (returned in message *ii* in Figure 1 and Figure 2) before accepting a vote for a candidate. Because a writing client will be unable to obtain a certificate for a stale access set, votes for such a candidate will be rejected by reading clients. Therefore, the analysis in Section 3 applies without additional complications.




Figure 4: Write operation in masking quorum systems: messages and stages of verification of write marker. (Changes in gray.)

5 Additional Related Work

Probabilistic quorum systems were explored in the context of dynamic systems with non-uniform access strategies by Abraham and Malkhi [2]. Recently, probabilistic quorum systems have been used in the context of security for wireless sensor networks [5] as well as storage for mobile ad hoc networks [13]. Lee and Welch make use of probabilistic quorum systems in randomized algorithms for distributed read-write registers [11] and shared queue data structures [12].

Signed quorum systems presented by Yu [25] also weaken the requirements of strict quorum systems but use different techniques. However, signed quorum systems have not been analyzed in the context of Byzantine faults, and so they are not presently affected by write markers.

Another implementation of write markers was introduced by Alvisi et al. [3] for purposes different than ours. We achieve the goals of (i) improving the load, and (ii) increasing the maximum fraction of faults that the system can tolerate by using write markers to prevent some faulty servers from colluding. In contrast to this, Alvisi et al. use write markers in order to increase accuracy in estimating the number of faults present in Byzantine quorum systems, and for identifying faulty servers that consistently return incorrect results. Because the implementation of Alvisi et al. does not prevent faulty servers from lying about the write quorums of which they are members, it cannot be used directly for our purposes. In addition, our implementation is designed to tolerate Byzantine clients, unlike theirs.

6 Conclusion

We have presented write markers, a way to improve the load of masking and opaque quorum systems asymptotically. Moreover, our new masking and opaque probabilistic quorum systems with write markers can tolerate an additional 24% and 17% of faulty replicas, respectively, compared with the proven bounds of probabilistic quorum systems without write markers. Write markers achieve this by limiting the extent to which Byzantine-faulty servers may cooperate to provide incorrect values to clients. We have presented a proposed implementation of write markers that is designed to be effective even while tolerating Byzantine-faulty clients and servers.

References

- [1] M. Abd-El-Malek, G. R. Ganger, G. R. Goodson, M. K. Reiter, and J. J. Wylie. Fault-scalable Byzantine fault-tolerant services. In *Symposium on Operating Systems Principles*, October 2005.
- [2] I. Abraham and D. Malkhi. Probabilistic quorums for dynamic systems. *Distributed Computing*, 18(2):113 – 124, 2005.
- [3] L. Alvisi, D. Malkhi, E. Pierce, and M. K. Reiter. Fault detection for Byzantine quorum systems. *IEEE Transactions on Parallel and Distributed Systems*, 12(9):996–1007, 2001.

- [4] M. Castro and B. Liskov. Practical Byzantine fault tolerance. In *Symposium on Operating Systems Design and Implementation*, 1999.
- [5] W. Du, J. Deng, Y. S. Han, P. K. Varshney, J. Katz, and A. Khalili. A pairwise key predistribution scheme for wireless sensor networks. *ACM Transactions on Information and System Security*, 8(2):228–258, 2005.
- [6] G. R. Goodson, J. J. Wylie, G. R. Ganger, and M. K. Reiter. Efficient Byzantine-tolerant erasure-coded storage. In *International Conference on Dependable Systems and Networks*, June 2004.
- [7] M. Herlihy and J. Wing. Linearizability: A correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems*, 12(3):463–492, 1990.
- [8] A. Juels and J. Brainard. Client puzzles: A cryptographic countermeasure against connection depletion attacks. In *Network and Distributed Systems Security Symposium*, pages 151–165, 1999.
- [9] L. Kong, D. Manohar, A. Subbiah, M. Sun, M. Ahamad, and D. Blough. Agile store: Experience with quorum-based data replication techniques for adaptive Byzantine fault tolerance. In *IEEE Symposium on Reliable Distributed Systems*, pages 143–154, 2005.
- [10] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, July 1982.
- [11] H. Lee and J. L. Welch. Applications of probabilistic quorums to iterative algorithms. In *International Conference on Distributed Computing Systems*, pages 21–30, April 2001.
- [12] H. Lee and J. L. Welch. Randomized shared queues applied to distributed optimization algorithms. In *International Symposium on Algorithms and Computation*, December 2001.
- [13] J. Luo, J.-P. Hubaux, and P. T. Eugster. Pan: providing reliable storage in mobile ad hoc networks with probabilistic quorum systems. In *International symposium on mobile ad hoc networking and computing*, pages 1–12, 2003.
- [14] D. Malkhi, Y. Mansour, and M. K. Reiter. Diffusion without false rumors: On propagating updates in a Byzantine environment. *Theoretical Computer Science*, 299(1–3):289–306, 2003.
- [15] D. Malkhi and M. Reiter. Byzantine quorum systems. *Distributed Computing*, 11(4):203–213, 1998.
- [16] D. Malkhi and M. K. Reiter. An architecture for survivable coordination in large distributed systems. *IEEE Transactions on Knowledge and Data Engineering*, 12(2):187–202, 2000.
- [17] D. Malkhi, M. K. Reiter, and A. Wool. The load and availability of Byzantine quorum systems. *SIAM Journal of Computing*, 29(6):1889–1906, 2000.

- [18] D. Malkhi, M. K. Reiter, A. Wool, and R. N. Wright. Probabilistic quorum systems. *Information and Computation*, 170(2):184–206, 2001.
- [19] J.-P. Martin and L. Alvisi. Fast Byzantine consensus. *IEEE Transactions on Dependable and Secure Computing*, 3(3):202–215, 2006.
- [20] J.-P. Martin, L. Alvisi, and M. Dahlin. Minimal Byzantine storage. In *International Symposium on Distributed Computing*, 2002.
- [21] M. G. Merideth and M. K. Reiter. Probabilistic opaque quorum systems. In *International Symposium on Distributed Computing*, 2007.
- [22] M. Mitzenmacher and E. Upfal. *Probability and Computing*. Cambridge University Press, 2005.
- [23] M. Molloy and B. Reed. *Graph Colouring and the Probabilistic Method*. Springer, 2002.
- [24] M. Naor and A. Wool. The load, capacity, and availability of quorum systems. *SIAM Journal on Computing*, 27(2):423–447, 1998.
- [25] H. Yu. Signed quorum systems. *Distributed Computing*, 18(4):307–323, 2006.