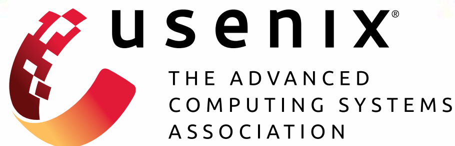

WaterBear: Practical Asynchronous BFT Matching Security Guarantees of Partially Synchronous BFT

Haibin Zhang, *Beijing Institute of Technology;* Sisi Duan, *Tsinghua University, Zhongguancun Laboratory;* Boxin Zhao, *Zhongguancun Laboratory;* Liehuang Zhu, *Beijing Institute of Technology*

https://www.usenix.org/conference/usenixsecurity23/presentation/zhang-haibin

This paper is included in the Proceedings of the
32nd USENIX Security Symposium.

August 9–11, 2023 • Anaheim, CA, USA

978-1-939133-37-3

Open access to the Proceedings of the
32nd USENIX Security Symposium
is sponsored by USENIX.

WaterBear: Practical Asynchronous BFT Matching Security Guarantees of Partially Synchronous BFT

Haibin Zhang
bchainzhang@aliyun.com
Beijing Institute of Technology

Sisi Duan*
duansisi@tsinghua.edu.cn
Tsinghua University, Zhongguancun Laboratory

Boxin Zhao*
zhaobx@zgclab.edu.cn
Zhongguancun Laboratory

Liehuang Zhu*
liehuangz@bit.edu.cn
Beijing Institute of Technology

Abstract

Asynchronous Byzantine fault-tolerant (BFT) protocols assuming no timing assumptions are inherently more robust than their partially synchronous counterparts, but typically have much weaker security guarantees.

We design and implement WaterBear, a family of new and efficient asynchronous BFT protocols matching all security guarantees of partially synchronous protocols. To achieve the goal, we have developed the local coin (flipping a coin locally and independently at each replica) based BFT approach—one long deemed as being inefficient—and designed more efficient asynchronous binary agreement (ABA) protocols and their reproposable ABA (RABA) versions from local coins.

We implemented in total five BFT protocols in a new golang library, including four WaterBear protocols and BEAT. Via extensive evaluation, we show that our protocols are efficient under both failure-free and failure scenarios, achieving at least comparable or superior performance to BEAT with much weaker security guarantees. Specifically, the most efficient WaterBear protocol consistently outperforms BEAT in terms of all metrics. For instance, when the number of replicas is 16, the latency of our protocol is about 1/8 of that of BEAT and the throughput of our protocol is 1.23x that of BEAT.

Our work pushes the boundaries of asynchronous BFT, showing the strongest security levels that we know of and high performance can co-exist for asynchronous BFT.

1 Introduction

Byzantine fault-tolerant state machine replication (BFT), a technique traditionally used to build mission-critical systems, has nowadays been the standard model for permissioned blockchains [6, 15, 32, 33, 53, 55, 56] and is used in various ways in hybrid blockchains. Due to their inherent robustness against performance and DoS attacks, asynchronous BFT protocols—relying on no timing assumptions—have been receiving significant attention [8]. While one line of works focuses on performance [25, 39, 40, 47], some other works aim at improving their "security." For instance, BEAT [30], PACE [57], and FIN [31] eliminated the less-established pairing assumption in these protocols; EPIC and HALE aimed at providing adaptive security [44, 60]; DAG-Rider strived to achieve quantum safety (though not quantum liveness) [42]; recent works studied how to avoid trusted setup [2, 27, 43, 58].

*Corresponding author

Table 1 summarizes the security levels that can be achieved for asynchronous BFT protocols implemented. The situation is in sharp contrast to their partially synchronous BFT counterparts (relying on timing assumptions): for example, the classic PBFT protocol [19]—based on authenticated channels only—easily achieves all the properties listed in the table. It is thus our goal to design and implement *practical* asynchronous BFT protocols achieving all these properties in Table 1—the same security guarantees as in partially synchronous BFT.

1.1 Background on Security Guarantees

Authenticated channels only vs. no PKC vs. quantum security. When designing practical fault-tolerant and cryptographic protocols, it is vital to use weak assumptions to reduce the attack surface and thus achieve strong guarantees. Arguably the "minimal" (and most frequently used) assumption is the point-to-point authenticated channel. A slightly stronger assumption is to use symmetric cryptography only—no public-key cryptography (PKC). Symmetric cryptography primitives, such as message authentication codes (MACs) and hash functions—with appropriately chosen parameters—are believed to defend against quantum adversaries that may leverage quantum mechanical phenomena to solve the problem intractable for conventional computers (see NIST reports [20]).

In the partial synchrony setting, the classic PBFT protocol [19] is the first one relying on no PKC—authenticated channels and hash functions only. (The protocol in Castro's PhD thesis [18] and Cachin's formulation for PBFT [12] assume authenticated channels and hash functions too.) They can be modified to assume authenticated channels only, as commented in PBFT "It is also possible to modify the algorithm (PBFT) not to use a cryptographic hash function by replacing the hash of a message by the value of the message." Recently, Stern and Abraham proposed a variant of

	authenticated channel only	no pkc	quantum secure	no trusted setup	adaptive security	high WAN throughput
SINTRA [13]						
RITAS [48]		√	√	√	√	
HoneyBadger [47]; BEAT [30]						√
Dumbo family [39,40]						√
EPIC [44]					√	√
Tusk [25]; Bullshark [37]						√
PACE [57]						√
SodsBC [28]		√	√			√
WaterBear-QS (this work)		√	√	√	√	√
WaterBear (this work)	√	√	√	√	√	√

Table 1: Comparison of efficient asynchronous BFT systems.

HotStuff relying on authenticated channels only and achieving improved message complexity [54]. All these partially synchronous BFT protocols achieve quantum security.

In contrast, we do not know how to build asynchronous BFT with quantum security, asynchronous BFT without PKC, or asynchronous BFT with authenticated channels only.

Adaptive security vs. static security. Depending on how the adversary decides to corrupt replicas, there are two types of corruptions: static corruptions and adaptive corruptions. A static adversary is restricted to choose its set of corrupted replicas at the start of the protocol. An adaptive adversary can choose its set of corrupted replicas at any moment during the execution of the protocol, based on the information it has accumulated. There is a strong separation result that statically secure protocols are not necessarily adaptively secure [16, 24]. Classic partially synchronous protocols such as PBFT naturally achieve adaptive security, while all asynchronous BFT protocols implemented (except RITAS [48] and EPIC [44]) achieve static security and even for RITAS and EPIC, they achieve adaptive security with significant performance penalty.

1.2 Why Matching Security Guarantees of Partially Synchronous BFT Hard?

Unlike partially synchronous BFT protocols, asynchronous BFT protocols must be randomized to achieve liveness (due to the celebrated FLP impossibility result [36]). Existing asynchronous BFT protocols rely critically on cryptographic common coin protocols (a distributed object that generates the same random coins to all replicas) which are inefficient if no trusted setup is assumed (see Sec. 2 for detailed discussion). Even if relaxing to quantum security, we currently lack efficient common coin instantiations from quantum secure primitives (e.g., lattices). Note while SodsBC is a quantum-secure BFT protocol, it directly relies on trusted setup to generate the common coins and thus bypasses the core problem of generating common coins efficiently [28].

Meanwhile, achieving stronger security in asynchronous BFT typically comes with a higher cost. For instance, adding adaptive security in [44] makes the original BFT system much slower. We must guarantee that ensuring these properties altogether would not incur too much overhead.

1.3 Our Approach

Reducing the problem to local coin RABA then to ABA. Instead of using common coins, we revisit the local coin based BFT approach that has been long viewed as being inefficient. In the local coin based approach, replicas need to independently and locally flip coins and existing protocols terminate in exponential expected rounds and fail to scale [22, 48].

We thus take a detour and develop the PACE BFT framework [57] using authenticated channels only and achieving quantum security. (Recall existing instantiations in PACE use trusted setup and threshold cryptography, achieving static security only.) PACE devises a variant of asynchronous binary Byzantine agreement (ABA) called reproposable ABA (i.e., RABA). Compared to ABA, RABA has additional properties, allowing all RABA instances to run in parallel and hence improving system throughput. It is also shown that some ABA protocols can be efficiently converted to RABA protocols. Crucially, RABA in PACE enables a fast path for consensus. We observe that while PACE was designed with common coin based RABA, the protocol, even with local coin based RABA, can—on average—terminate within a single RABA round with high probability. As most RABA instances will terminate in one round, the exponential rounds in local coin based ABA is no longer a major efficiency obstacle.[1] Now the per-round complexity of RABA protocols becomes critical. Our strategy is to *reduce asynchronous BFT to RABA with local coins and then to ABA with local coins* and then *improve the per-round complexity of them*.

Improving the per-round complexity of (R)ABA. As reported in almost all asynchronous BFT systems [30, 40, 57], ABA is their major performance bottleneck. It is shown that the concrete steps *per round* of ABA protocols are vital to the performance of asynchronous BFT: even a single step improvement in ABA, the resulting BFT protocol could be improved by, say, 2x [57].

To our knowledge, only two local coin based ABA proto-

[1] Jumping ahead, we experimentally demonstrate through latency-breakdown and robustness tests to validate the claim.

protocol	reference implementation	RBC	RABA	building blocks
WaterBear	WaterBear-C	Bracha's RBC [11]	Cubic-RABA (this paper)	MAC and hash
	WaterBear-Q	Bracha's RBC [11]	Quadratic-RABA (this paper)	MAC and hash
WaterBear-QS	WaterBear-QS-C	CT RBC [14]	Cubic-RABA (this paper)	MAC
	WaterBear-QS-Q	CT RBC [14]	Quadratic-RABA (this paper)	MAC

Table 2: WaterBear-QS and WaterBear instantiations. As in PACE, both WaterBear-QS and WaterBear have a fast path allowing them to terminate in $O(\log n)$ time. As shown in PACE, the probability of triggering fast paths is high. WaterBear-C and WaterBear-QS-C have $O(n^4)$ messages on average due to the usage of Cubic-RABA, while WaterBear-Q and WaterBear-QS-Q have $O(n^3)$ messages on average due to the usage of Quadratic-RABA—matching those of HoneyBadger, BEAT, and PACE.

ABA (local coins)	messages/round	steps/round
Bracha's ABA [11]	n^3	9 to 12
Cubic-ABA (this work)	n^3	5 to 7
Quadratic-ABA (this work)	n^2	4 or 5

Table 3: Local coin based ABA protocols with optimal resilience. We consider the messages and steps in each round. Messages/round and steps/round denote number of messages and steps among all replicas per round.

cols have been proposed: Ben-Or's ABA [8] assuming $n > 5f$, and Bracha's ABA [10] with $n > 3f$ (the most efficient protocol for nearly three decades). Bracha's ABA, unfortunately, has a large number of steps (12 steps) and $O(n^3)$ messages per round [10]. (The situation is in sharp contrast to ABA assuming common coins which has 3 steps and $O(n^2)$ messages per round.) Our main technical contributions are indeed efficient local coin based ABA and RABA protocols with improved message complexity and reduced number of steps.

1.4 Our Contributions

1.4.1 Technical Contributions

Efficient local coin based ABA. Table 3 shows two novel local coin based ABA protocols that we introduce in the paper: Cubic-ABA and Quadratic-ABA. Cubic-ABA is easy to understand and implement, and can be viewed as an optimized version of Bracha's ABA. Cubic-ABA has 7 steps per round in the worst case, while Bracha's ABA uses 12 steps, almost doubling the number of steps of Cubic-ABA. In contrast, Quadratic-ABA adopts a novel design, having 4 or 5 steps per round only and being the first local coin based ABA with $O(n^2)$ messages per round. In particular, Quadratic-ABA admits a fast (coin-free) allowing the protocol to terminate in a single step in the optimistic mode.

Tackling a subtle liveness issue for RABA. We go on to design Cubic-RABA and Quadratic-RABA based on Cubic-ABA and Quadratic-ABA, respectively. Unlike prior transformations following a generic approach in [57], we identify and tackle a subtle liveness problem when transforming Quadratic-ABA to Quadratic-RABA. The issue that we identify demonstrates the subtlety of transforming ABA to RABA, and once again underlines the importance of a full proof when designing Byzantine-resilient protocols.

ABA from perfect common coins. The techniques we introduce for Quadratic-ABA are of independent interests, allowing us to obtain an ABA protocol that works for perfect common coins. The protocol can be used to improve various BFT protocols such as PACE and Dumbo [40], and the recent distributed key generation protocol requiring the good-case-coin-free property [27].

ABA (common coins)	steps/round	rounds	good-case-coin-free
MMR15 [49, 2nd alg]	9 to 13	3	yes
Cobalt [46]	3 or 4	4	no
Crain [23, 1st alg]	5 to 7	3	yes
Crain [23, 2nd alg]	2 or 3†	4	no
Pillar [57]	2 or 3	4	no
Our ABA (this work)	4 or 5	3	yes

Table 4: Our ABA protocol using perfect common coins. †The second algorithm of Crain relies on high threshold common coins and is less efficient than Pillar. Compared to Pillar, our ABA protocol using perfect common coins has the good-case-coin-free property that is vital for applications such as asynchronous distributed key generation protocols [27].

1.4.2 Practical Contributions

The WaterBear family of BFT protocols. Table 2 summarizes the characteristics of WaterBear protocols. We use Cubic-RABA to build WaterBear-C and Quadratic-RABA to build WaterBear-Q. WaterBear has *all* desirable properties a BFT protocol one could think of, being optimally resilient, using authenticated channels only, achieving quantum security and adaptive security, and not relying on trusted setup—matching the security guarantees of the classic PBFT protocol. We comment that WaterBear does not attain information-theoretic (IT) security, as we use HMAC that is not IT-secure.

We also build WaterBear-QS using authenticated channels and hash functions and achieving quantum security for both safety and liveness properties. Similar to WaterBear, WaterBear-QS family also consists of two protocols: WaterBear-QS-C and WaterBear-QS-Q, quantum secure versions of WaterBear-C and WaterBear-Q, respectively.

A new BFT platform. Starting from HoneyBadger, existing asynchronous BFT protocols, including BEAT, Dumbo, and EPIC, use the HoneyBadger programming framework using Python 2.7 (end of life and end of support on January 1, 2020). We instead build a new platform using Golang that is more modular and developer-friendly than existing ones. Our platform currently supports WaterBear-C, WaterBear-Q,

WaterBear-QS-C, WaterBear-QS-Q, and BEAT (one of the most efficient open-source asynchronous BFT) [30]. Due to its modularity, the library only contains about 11,000 LOC.

Large-scale experiments and robustness evaluation. With a 61-instance deployment on Amazon EC2, we show our protocols offer comparable performance as the state-of-the-art asynchronous BFT protocols, while achieving much stronger security. We also design and evaluate various failure and attack scenarios, showing all our protocols are highly robust during failures and attacks. Specifically, one of our protocols, WaterBear-QS-Q, consistently outperforms BEAT (with much weaker security guarantees); for instance, when $n = 16$, the latency of WaterBear-QS-Q is about 1/8 that of BEAT and the throughput of WaterBear-QS-Q is 1.23x that of BEAT. The peak throughput of WaterBear-QS-Q (as n grows larger) is about 1.47x that of BEAT.

1.5 Paper Organization

In what follows, we first discuss related work (Sec. 2) and describe the system model and definitions (Sec. 3). Then we present Cubic-ABA and Quadratic-ABA from local coins (Sec. 4), their RABA counterparts—Cubic-RABA and Quadratic-RABA (Sec. 5), and WaterBear BFT protocols (Sec. 6). Last, we present our evaluation results (Sec. 7) before concluding the paper (Sec. 8). The proof of our protocols can be found in our full paper [59].

2 Related Work

ADKG. A line of recent works studied how to eliminate trusted setup by using asynchronous distributed key generation (ADKG) [2, 27, 43, 58]. Even if using ADKG in existing asynchronous BFT protocols, they would neither achieve quantum security nor security with no PKC.

Adaptive vs. static security for BFT. Most asynchronous BFT protocols implemented, including SINTRA, HoneyBadgerBFT, BEAT, and Dumbo, defend against static adversary only. These protocols rely critically on efficient but statically secure threshold cryptography. EPIC is an asynchronous BFT that uses adaptively secure threshold pseudorandom function (PRF) to achieve adaptive security but is not as efficient as its statically secure counterparts. RITAS [48] contains an adaptively secure BFT protocol, but due to inefficient local coin based ABA, it is less efficient than other protocols in large-size networks. The situation for asynchronous environments is in sharp contrast to that of partially synchronous protocols, most of which attain adaptive security [5, 19, 21, 29, 38, 41].

Quantum safety (but no quantum liveness). A BFT protocol is quantum secure if its safety is quantum resistant (quantum safety) and its liveness is quantum resistant (quantum liveness) [42]. DAG-Rider [42] achieves quantum safety. Moreover, the BKR protocol and their descendants (e.g., HoneyBadger [47], MiB [45], PACE [57]) achieve quantum safety if using techniques from EPIC [44]. All these protocols, however, do not achieve quantum liveness. Tusk [25] and Bullshark [37] are variants of DAG-Rider; they extensively use signatures and hashes and achieve neither quantum safety nor quantum liveness.

Byzantine agreement and common coins. Byzantine agreement (BA) is a central tool for both distributed computing and cryptography. The condition $n \geq 3f + 1$ is both necessary and sufficient for both synchronous and asynchronous BA protocols [51]. The celebrated impossibility result of Fischer, Lynch, and Paterson [36] implies that a randomized BA protocol must have non-terminating executions. ABA protocol may thus be either $(1-\varepsilon)$-terminating, where correct replicas terminate the protocol with an overwhelming probability, or almost-surely terminating, where replicas terminate with probability 1. For both types, we review ABA protocols assuming authenticated channels only. Note there is no need to consider ABA with authenticated channels and hash functions, as the input to ABA is a binary value and we do not need to use hash functions to compress the input.

For almost-surely ABA terminating with probability 1, Ben-Or's ABA requires $n \geq 5f + 1$ [8], while Bracha's ABA [10] achieves optimal resilience. The two protocols use local coins and require an exponential expected running time. Feldman and Micali propose a BA protocol having a constant expected running time in synchronous environments and extend it to build a polynomial-time ABA protocol requiring $n \geq 4f + 1$ [35]. Abraham, Dolev, and Halpern [1] provide the first almost-surely ABA with polynomial efficiency (concretely, expected $O(n^2)$ time) and optimal resilience. Bangalore, Choudhury, and Patra [7] improve the expected running time of [1] by a factor of n.

For $(1-\varepsilon)$-terminating ABA, Canetti and Rabin [17] build an expected constant-round ABA protocol with optimal resilience. Patra, Choudhury, and Rangan [50] build a more efficient construction in terms of communication complexity.

Both types of ABA protocols follow the classic framework of Feldman and Micali [35] that reduces ABA to asynchronous verifiable secret sharing (AVSS). The framework uses AVSS to build common coins. (The original idea of using common coin for ABA is due to Rabin [52].) Unfortunately, the framework of using AVSS for common coins is prohibitively expensive. For instance, to build AVSS, the approach of Canetti and Rabin [17] needs to begin with an information checking protocol, then asynchronous recoverable sharing, then asynchronous weak secret sharing, and finally AVSS. The improved approach of Patra, Choudhury, and Rangan [50] remains complex, following the route of information-checking protocol, then asynchronous weak commitment, and then AVSS. Moreover, the transformation from AVSS to ABA is equally expensive, requiring running n^2 AVSS instances to generate a *single* (weak) coin. Patra, Choudhury, and Rangan [50] also propose an approach for sharing multiple secrets simultaneously. While such an approach is useful for building more efficient multi-valued BA (MBA), it is unknown

if it would yield more efficient ABA protocols. While, for instance, the CNV framework [22] does use MBA, it may run $O(n)$ consecutive MBA instances (which is inefficient).

3 System Model and Definitions

3.1 System and Threat Model

This section describes the system model for distributed protocols in the paper, where f out of n replicas may fail arbitrarily (Byzantine failures). We assume point-to-point authenticated channels between each pair of replicas; some of our protocols additionally assume hash functions. The WaterBear BFT protocols (and subprotocols we use or we invent) have the following properties:

- **Optimal resilience**: The protocols in this work assume $f \leq \lfloor \frac{n-1}{3} \rfloor$, which is optimal. A (Byzantine) *quorum* is a set of $\lceil \frac{n+f+1}{2} \rceil$ replicas. For simplicity, we may assume $n = 3f+1$ and a quorum size of $2f+1$.
- **Asynchronous network**: We consider completely asynchronous systems making no timing assumptions on message processing or transmission delays. In contrast, partially synchronous systems assume that there exist an upper bound on message processing and transmission delays but the bound may be unknown to anyone [34].

  Designing asynchronous systems is challenging, because in asynchronous environments it is impossible to distinguish Byzantine faulty replicas from "slow" replicas. In particular, one cannot use timers or timeout to assist in the design of asynchronous systems.
- **No dealer/trusted setup**: We do not assume the existence of a trusted dealer or trusted setup. Neither do we assume there exists an interactive protocol for any public keys, reference strings, or public parameters.
- **Adaptive corruptions**: We consider adaptive adversary that can choose its set of corrupted replicas at any moment during the execution of the protocol, based on the information it has accumulated thus far (i.e., the messages observed and the states of previously corrupted replicas).

We may associate a protocol instance with a unique identifier *id*, tagging each message in the instance with *id*. If no ambiguity arises, we may omit the identifiers.

3.2 Definitions and Preliminaries

BFT. In a BFT protocol, a replica *a-delivers* (atomically deliver) *transactions*, each *submitted* by some client. The client computes a final response to its submitted transaction from the responses it receives from replicas. We consider the following properties:

- **Agreement**: If any correct replica *a-delivers* a transaction *tx*, then every correct replica *a-delivers tx*.
- **Total order**: If a correct replica *a-delivers* a transaction *tx* before *a-delivering tx'*, then no correct replica *a-delivers* a transaction *tx'* without first *a-delivering tx*.
- **Liveness**: If a transaction *tx* is *submitted* to all correct replicas, then all correct replicas eventually *a-deliver tx*.

Below, we first introduce ABA and then its variant—RABA. Then we review the PACE framework using RBC and RABA.

Asynchronous binary Byzantine agreement (ABA). An ABA protocol is specified by *propose* and *decide*. Each replica proposes an initial binary value (called *vote*) for consensus and replicas will decide on some value. ABA should satisfy the following properties:

- **Validity**: If all correct replicas *propose* v, then any correct replica that terminates *decides* v.
- **Agreement**: If a correct replica *decides* v, then any correct replica that terminates *decides* v.
- **Termination**: Every correct replica eventually *decides* some value.
- **Integrity**: No correct replica *decides* twice.

RABA. Reproposable ABA (RABA) is a new distributed computing primitive introduced in PACE [57]. In contrast to conventional ABA protocols, where replicas can vote once only, RABA allows replicas to change their votes. Formally, a RABA protocol tagged with a unique identifier *id* is specified by $propose(id,\cdot)$, $repropose(id,\cdot)$, and $decide(id,\cdot)$, with the input domain being $\{0,1\}$. For our purpose, RABA is "biased towards 1." Each replica can propose a vote v at the beginning of the protocol. Each replica can propose a vote only once. A correct replica that proposed 0 is allowed to change its mind and repropose 1. A replica that proposed 1 is not allowed to repropose 0. If a replica reproposes 1, it does so at most once. A replica terminates the protocol identified by *id* by generating a decide message. RABA (biased towards 1) satisfies the following properties:

- **Validity**: If all correct replicas *propose* v and never *repropose* $\bar{v}$, then any correct replica that terminates *decides* v.
- **Unanimous termination**: If all correct replicas *propose* v and never *repropose* $\bar{v}$, then all correct replicas eventually terminate.
- **Agreement**: If a correct replica *decides* v, then any correct replica that terminates *decides* v.
- **Biased validity**: If $f+1$ correct replicas *propose* 1, then any correct replica that terminates *decides* 1.
- **Biased termination**: Let Q be the set of correct replicas. Let Q_1 be the set of correct replicas that propose 1 and never repropose 0. Let Q_2 be correct replicas that propose 0 and later repropose 1. If $Q_2 \neq \emptyset$ and $Q = Q_1 \cup Q_2$, then each correct replica eventually terminates.
- **Integrity**: No correct replica decides twice.

Validity is slightly different from those for ABA. They are modified to accommodate the RABA syntax. Integrity is defined to ensure RABA decides once and once only.

Unanimous termination and biased termination are carefully introduced to help achieve RABA termination in certain scenarios. External operations would have to force the protocol to meet these termination conditions.

Biased validity in RABA requires that if $f+1$ correct repli-

cas, not simply all correct replicas, propose 1, then a correct replica that terminates decides 1. The property guarantees the PACE framework has sufficient transactions delivered.

This paper introduces new RABA protocols from local coins.

RBC. In a Byzantine reliable broadcast (RBC) protocol [3, 4, 11, 14, 26], a replica p first starts the protocol by executing *r-broadcast* with messages m, and all replicas terminate the protocol by executing *r-deliver* with message m. We consider the following properties:

- **Validity**: If a correct replica p *r-broadcasts* a message m, then p eventually *r-delivers* m.
- **Agreement**: If some correct replica *r-delivers* a message m, then every correct replica eventually *r-delivers* m.
- **Integrity**: For any message m, every correct replica *r-delivers* m at most once. Moreover, if the sender is correct, then m was previously *r-broadcast* by the sender.

This paper uses Bracha's broadcast [10] that assumes authenticated channels only and has a bandwidth of $O(n^2|m|)$, and uses CT RBC due to Cachin and Tessaro [14] that additionally uses hash functions (with output length λ) to reduce the bandwidth to $O(n|m| + \lambda n^2 \log n)$.

PACE framework. PACE uses RBC and RABA in a black-box manner to construct efficient asynchronous BFT. The framework allows all RABA instances to run in parallel, removing a well-known bottleneck in the original framework of Ben-Or, Kelmer, and Rabin [9]. Concretely, PACE has a RBC phase and a RABA phase; correct replicas can run the RABA phase in parallel once $n - f$ RBC instances have completed. PACE also provides a fast path for consensus, allowing the protocol to terminate using a single RABA round.

Steps, phases, and rounds. In asynchronous environments, the network delay is unbounded. To measure the latency of asynchronous protocols, we use the standard notion of *asynchronous steps* [17], where a protocol runs in x asynchronous steps if its running time is at most x times the maximum message delay between honest replicas during the execution.

We also use the notion of *phases* for ease of description, where a phase in a protocol consists of a fixed number of steps. When describing some of our protocols, we may divide a protocol into several phases, each of which has several steps.

In this paper, the notion of *rounds* is restricted to ABA protocols: an ABA protocol proceeds in rounds, where an ABA round consists of a fixed number of steps. For instance, local coin ABA protocols terminate in expected exponential rounds, while ABA assuming perfect common coins (including our ABA protocol that we introduce in this paper) terminates in expected constant rounds. An ABA round may consist of several phases and each phase consists of several steps. In asynchronous ABA systems, replicas proceed in rounds and they might not be always in the same round, but each correct replica eventually terminates every round that it has participated in.

4 ABA from Local Coins

Summary of our results. The state-of-the-art local coin based ABA protocol, Bracha's ABA [10], has $O(n^3)$ messages and 12 steps in each round. We design two new ABA protocols from local coins, Cubic-ABA and Quadratic-ABA, with two goals in mind—being more efficient than Bracha's ABA and being compatible with RABA.

We begin with the simpler one, Cubic-ABA, that achieves the same message complexity as Bracha's ABA but has only 7 steps in each round. Cubic-ABA admits a clean and intuitive proof of correctness. We go on to suggest Quadratic-ABA on top of Cubic-ABA. Compared to Cubic-ABA, Quadratic-ABA reduces the messages from $O(n^3)$ to $O(n^2)$ and reduces the number of steps to 5 in each round. *The improvement is significant, allowing WaterBear to attain the same average message complexity in normal cases as PACE—$O(n^3)$.* Equally important, both ABA protocols can be modified for efficient RABA protocols.

As an important by-product, extending the idea of Quadratic-ABA and assuming the existence of perfect common coins, we can present ABA protocols that have expected constant rounds and outperform the state-of-the-art protocols, as shown in Table 4 in the introduction.

4.1 Cubic-ABA

Overview. Our motivation for Cubic-ABA is to reduce the number of parallel RBCs in Bracha's ABA from three to two. In particular, in each round, Bracha's ABA has three phases, where in each phase, replicas run n parallel RBCs. In total, Bracha's ABA has 12 steps and $O(n^3)$ messages per round. (We recall Bracha's ABA in Appendix A.) In Cubic-ABA, we replace the first two RBC phases with one or two steps of all-to-all broadcast so Cubic-ABA only has 5 to 7 steps.

The protocol. Figure 1 describes the pseudocode of Cubic-ABA and Figure 2 illustrates the workflow. Cubic-ABA uses the *broadcast* primitive (multicasting messages to all replicas) and the *r-broadcast* and *r-deliver* primitives of RBC. The protocol proceeds in rounds, beginning with round 0. Each round r consists of three phases. In the first phase, a replica p_i broadcasts a $\textsf{pre-vote}_r(iv_r)$ message, where pre-vote is the message type, r is the round when the message was sent, and $iv_r \in \{0,1\}$ is the input value of p_i for round r (ln 07). At ln 08-09, if p_i receives $f+1$ $\textsf{pre-vote}_r(v)$ for some $v \in \{0,1\}$ and has not previously broadcast $\textsf{pre-vote}_r(v)$, it also broadcasts $\textsf{pre-vote}_r(v)$.

At ln 10-14, p_i enters the second phase. If p_i receives $2f+1$ $\textsf{pre-vote}_r(v)$, it adds v to its $bset_r$, a set consisting only 0 and 1 (ln 10-11). Letting v be the *first* value added to $bset_r$ for p_i, p_i broadcasts a $\textsf{main-vote}_r(v)$ message (ln 12-14).

In the third phase, a correct replica p_i accepts a $\textsf{main-vote}_r(v)$ message only if v has already been added locally to $bset_r$ (ln 15). If p_i has received $n-f$ $\textsf{main-vote}_r(v)$, p_i *r-broadcasts* a $\textsf{final-vote}_r(v)$ message (ln 16-17). Oth-

```
01 initialization
02   r ← 0                                                    {round}
03 func propose(v_input)
04   iv_0 ← v_input                              {set input for round 0}
05   start round 0
06 round r
07   broadcast pre-vote_r(iv_r)                              {▷ phase 1}
08   upon receiving pre-vote_r(v) from f+1 replicas
09     if pre-vote_r(v) has not been sent, broadcast pre-vote_r(v)
10   upon receiving pre-vote_r(v) from 2f+1 replicas         {▷ phase 2}
11     bset_r ← bset_r ∪ {v}
12   wait until bset_r ≠ ∅
13     if main-vote_r() has not been sent
14       broadcast main-vote_r(v) where v ∈ bset_r
15   upon receiving n−f main-vote_r() such that for each received
main-vote_r(b), b ∈ bset_r                                   {▷ phase 3}
16     if there are n−f main-vote_r(v)
17       r-broadcast final-vote_r(v)
18     else r-broadcast final-vote_r(∗)
19   upon r-delivering n−f final-vote_r() such that for each
final-vote_r(v), v ∈ bset_r; for each final-vote_r(∗), bset_r = {0,1}
20     if there are n−f final-vote_r(v)
21       iv_{r+1} ← v, decide v
22     else if there are f+1 final-vote_r(v)
23       iv_{r+1} ← v
24     else
25       iv_{r+1} ← Random()                           {obtain local coin}
26   r ← r+1
```

Figure 1: Cubic-ABA. The code for p_i. $v \in \{0,1\}$.

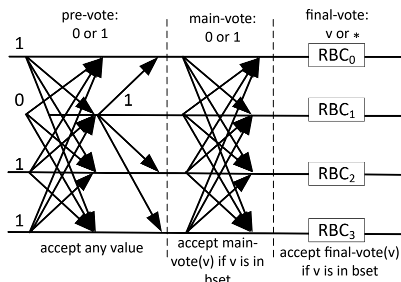

Figure 2: The workflow of Cubic-ABA.

erwise, p_i *r-broadcasts* final-vote$_r(*)$, where $*$ is a distinguished symbol that is neither 0 nor 1 (ln 18).

A correct p_i accepts a final-vote$_r()$ message, if one of the following two conditions holds (ln 23):

- For a final-vote$_r(v)$ message with $v \in \{0,1\}$, v has been added to $bset_r$ for p_i.
- For a final-vote$_r(*)$ message, $bset_r$ contains both 0 and 1.

Upon *r-delivering* $n-f$ valid final-vote$_r()$ messages, we distinguish three cases:

- Ln 20-21: If p_i *r-delivers* $n-f$ valid final-vote$_r(v)$ for the same $v \in \{0,1\}$, p_i decides v and uses v as iv_{r+1} to enter the next round. Each correct replica that decides in round r continues for one more round (up to the final-vote$_r()$ step) and terminates the protocol.
- Ln 22-23: If p_i *r-delivers* at least $f+1$ valid final-vote$_r(v)$ for some $v \in \{0,1\}$, p_i uses v as input for the next round.
- Ln 24-25: Otherwise, a replica generates a local random coin and uses it as input for the next round.

Analysis. In our approach, the first two phases of Cubic-ABA resemble those of common coin based ABA protocols [23, 46, 49, 57], where we ask replicas to broadcast their values. In particular, the first phase ensures that all correct replicas eventually acknowledge the same set of values $bset_r$; the second phase ensures that no two correct replicas will vote for opposite values in the third phase, though one correct replica may vote for $b \in \{0,1\}$ and one may vote for $*$ (a distinguished vote). Accordingly, we do not have to rely on RBC for the first two phases, as our first two phases already guarantee that correct replicas will not vote for conflicting values for the third phase.

In the third phase, we need to ensure that if a correct replica receives $n-f$ final-vote$_r(v)$ and decide, any correct replica will propose v and eventually decide v. Note for the case where $f+1$ correct replicas vote for v and f correct replicas vote for $*$, we need to guarantee that if a correct replica receives $n-f$ final-vote$_r(v)$, any correct replica will receive at least $f+1$ final-vote$_r(v)$ and therefore vote for v in the following round. Thus, we rely on RBC, ensuring that all correct replicas eventually receive consistent values, even in the presence of Byzantine replicas. As we show in the proof, this is crucial for the agreement property.

Note that our protocol presented is asynchronous and replicas are not always in the same round. If a replica p_i is in round r and receives a message from p_j tagged by $r' > r$ (e.g., pre-vote$_{r'}(v)$), p_i can simply store the message and will process it after p_i enters round r'. As shown in the termination property [59], each correct replica will proceed to the next round before it decides a value and all replicas will eventually decide in the same round.

As the number of n parallel RBC instances is 1 instead of 3, the number of steps is reduced from 12 to 7.

4.2 Quadratic-ABA

Overview. In Quadratic-ABA, we replace the only parallel RBC phase used in Cubic-ABA using a novel two-step all-to-all broadcast. The goal is to ensure that at the end of each round, if a correct replica receives $n-f$ matching votes for a value v, any correct replica will receive either $n-f$ votes for v or at least $f+1$ matching v. This will guarantee that all correct replicas will vote for v in the following round. By eliminating parallel RBC instances causing $O(n^3)$ messages and $O(n^3)$ communication, Quadratic-ABA now achieves $O(n^2)$ messages and $O(n^2)$ communication per round.

The protocol. The pseudocode of Quadratic-ABA is shown in Figure 3. The Quadratic-ABA protocol is round-based, starting from round 0. In each round, there are four phases—pre-vote$_r()$, vote$_r()$, main-vote$_r()$, and final-vote$_r()$, as shown in Figure 4. The pre-vote$_r()$ and

01 **initialization**
02 $r \leftarrow 0$ {round}
03 **func** propose(v_{input})
04 $iv_0 \leftarrow v_{input}$ {set input for round 0}
05 start round 0
06 **round** r
07 broadcast pre-vote$_r(iv_r)$ {$\triangleright$ phase 1}
08 **upon receiving** pre-vote$_r(v)$ from $f+1$ replicas
09 **if** pre-vote$_r(v)$ has not been sent, broadcast pre-vote$_r(v)$
10 **upon receiving** pre-vote$_r(v)$ from $2f+1$ replicas {$\triangleright$ phase 2}
11 $bset_r \leftarrow bset_r \cup \{v\}$
12 **wait until** $bset_r \neq \emptyset$
13 **if** vote$_r()$ has not been sent
14 broadcast vote$_r(v)$ where $v \in bset_r$
15 **upon receiving** $n-f$ vote$_r()$ such that for each vote$_r(v)$, $v \in bset_r$ {$\triangleright$ phase 3}
16 **if** there are $n-f$ vote$_r(v)$
17 broadcast main-vote$_r(v)$
18 **else** broadcast main-vote$_r(*)$
19 **upon** receiving $n-f$ main-vote$_r()$ such that for each main-vote$_r(v)$, at least $f+1$ vote$_r(v)$ have been received and for each main-vote$_r(*)$, $bset_r = \{0,1\}$ {$\triangleright$ phase 4}
20 **if** there are $n-f$ main-vote$_r(v)$
21 broadcast final-vote$_r(v)$
22 **else** broadcast final-vote$_r(*)$
23 **upon** receiving $n-f$ final-vote$_r()$ such that for each final-vote$_r(v)$, at least $f+1$ main-vote$_r(v)$ have been received and for each final-vote$_r(*)$, $bset_r = \{0,1\}$
24 **if** there are $n-f$ final-vote$_r(v)$
25 $iv_{r+1} \leftarrow v$, **decide** v
26 **else if** there are only final-vote$_r(v)$ and final-vote$_r(*)$
27 $iv_{r+1} \leftarrow v$
28 **else**
29 $iv_{r+1} \leftarrow Random()$ {obtain local coin}
30 $r \leftarrow r+1$

Figure 3: The Quadratic-ABA protocol. The code for p_i.

vote$_r()$ phases (ln 07-14) are similar to the pre-vote$_r()$ and main-vote$_r()$ phases in Cubic-ABA. In the first phase, every replica p_i broadcasts a pre-vote$_r(iv_r)$ message, where iv_r is the value p_i votes for in round r (ln 07). After receiving $f+1$ pre-vote$_r(v)$ and p_i has not previously broadcast pre-vote$_r(v)$, p_i also broadcasts pre-vote$_r(v)$ (ln 08-09). At ln 10-11, upon receiving $n-f$ pre-vote$_r(v)$, p_i adds v to $bset_r$. For the first value v added to $bset_r$, p_i broadcasts a vote$_r(v)$ message (ln 12-14).

For each vote$_r(v)$ message, p_i accepts it only if v has been added to $bset_r$. Upon receiving $n-f$ vote$_r()$ messages, one of the following two conditions holds.

- Ln 16-17: If p_i receives $n-f$ vote$_r(v)$ messages, it broadcasts a main-vote$_r(v)$ message.
- Ln 18: Otherwise, p_i broadcasts a main-vote$_r(*)$ message.

Every correct replica p_i accepts a main-vote$_r(v)$ message only if p_i has received $f+1$ vote$_r(v)$ messages. Every correct replica accepts a main-vote$_r(*)$ message only if

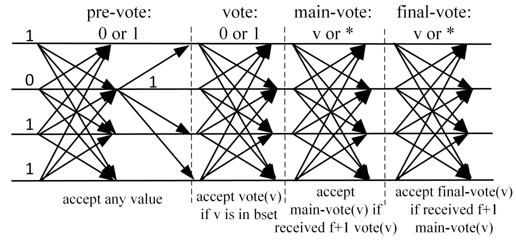

Figure 4: The workflow of Quadratic-ABA.

$bset_r = \{0,1\}$. Upon receiving $n-f$ main-vote$_r()$ messages, one of the following two conditions holds.

- Ln 20-21: If p_i receives $n-f$ main-vote$_r(v)$ messages, it broadcasts a final-vote$_r(v)$ message.
- Ln 22: Otherwise, p_i broadcasts final-vote$_r(*)$ message.

Every correct replica accepts a final-vote$_r(v)$ message only if it has received $f+1$ main-vote$_r(v)$ messages. Every correct replica accepts a final-vote$_r(*)$ message only if $bset_r = \{0,1\}$. Upon receiving $n-f$ final-vote$_r()$ messages, there are three cases:

- Ln 24-25: If p_i receives $n-f$ final-vote$_r(v)$, it decides v and also sets iv_{r+1} as v. It participates in the protocol for one more round and terminates the protocol.
- Ln 26-27: If p_i receives $n-f$ final-vote$_r()$ messages that carry only value v and $*$, it uses v for round $r+1$.
- Ln 28-29: Otherwise, p_i generates a local random coin and uses it as input for the next round.

Analysis. Quadratic-ABA guarantees that if a correct replica receives $n-f$ final-vote$_r(v)$, any correct replica will set iv_{r+1} as v. First, if a correct replica sends main-vote$_r(v)$, no correct replica will send main-vote$_r(\bar{v})$. In particular, if a correct replica sends main-vote$_r(v)$, it must have received $n-f$ vote$_r(v)$. If another correct replica sends main-vote$_r(\bar{v})$, at least $n-f$ replicas have sent vote$_r(v)$. Therefore, at least one correct replica has sent both vote$_r(v)$ and vote$_r(\bar{v})$, contradicting the fact that each correct replica only sends a single vote$_r()$ message in each round. Furthermore, if a correct replica sends final-vote$_r(v)$, no correct replica will send final-vote$_r(\bar{v})$ or even accept final-vote$_r(\bar{v})$ from other replicas. This is because if a correct replica accepts final-vote$_r(\bar{v})$, at least one correct replica has sent main-vote$_r(\bar{v})$. Meanwhile, if a correct replica accepts final-vote$_r(v)$, at least one correct replica has sent main-vote$_r(v)$, contradicting. Hence, at the end of each round, if a correct replica receives only final-vote$_r(v_1)$ and final-vote$_r(*)$, another correct replica receives only final-vote$_r(v_2)$ and final-vote$_r(*)$, it holds that $v_1 = v_2$. This result is crucial for agreement and termination.

Furthermore, if a correct replica decides v in round r, it must have received $n-f$ final-vote$_r(v)$. Among them, at least $f+1$ correct replicas have sent final-vote$_r(v)$. With $3f+1$ replicas in total, there are at most $2f$ final-vote$_r(\bar{v})$ or final-vote$_r(*)$. Hence, every correct replica receives at least one final-vote$_r(v)$. As no correct replica will accept final-vote$_r(\bar{v})$, every correct replica will only have final-vote$_r(v)$ and final-vote$_r(*)$. Thus, every correct replica

either decides in round r, or enters round $r+1$ and sets iv_{r+1} to v. Doing so ensures agreement.

5 RABA from Local Coins

As shown in PACE [57], the PACE framework with RABA significantly outperforms the conventional BKR diagram and enables a fast path for termination. Our goal here is to use local coin based ABA to design RABA without trusted setup. We use Cubic-ABA and Quadratic-ABA to build Cubic-RABA and Quadratic-RABA, respectively. Here, we focus on Quadratic-RABA and present Cubic-RABA in our full paper [59].

Subtlety of building Quadratic-RABA. PACE introduced a general approach to converting ABA to RABA [57]. Following their approach, we present Quadratic-RABA in Figure 5. Quadratic-RABA is identical to Quadratic-ABA except for the first round (round 0), where we make the following changes. First, we use a propose() event and a repropose() event (ln 03-07). Upon propose(v), a replica p_i starts round 0 and executes the *broadcast-vote*(v) function. Upon repropose(1) event, regardless of which round a replica is in, p_i still executes the *broadcast-vote*(v) function. The propose() and repropose() events are crucial for *biased termination*. So if a quorum of correct replicas either propose 1 or repropose 1, the protocol will eventually terminate.

Second, in the *broadcast-vote*(v) function, replica p_i broadcasts a pre-vote$_0(v)$ message (ln 09). At ln 10-14, if $v=1$, p_i immediately adds 1 to $bset_0$, and broadcasts vote$_0(1)$, main-vote$_0(1)$, and final-vote$_0(1)$.

Third, the coin value in round 0 is set to 1 (ln 38). The second and the third modifications guarantee both *biased validity* property and a *fast path allowing terminating the protocol in one step only*. Namely, if $f+1$ correct replicas propose 1, no correct replica will receive $2f+1$ final-vote$_0(0)$. As we will show in the proof, every correct replica will either directly decide 1 or set iv_1 as 1, so all correct replicas decide within two rounds. Furthermore, our protocol has a fast path: if all correct replicas propose 1, they will directly send vote$_0(1)$, main-vote$_0(1)$, final-vote$_0(1)$, allowing correct replicas to decide in one step.

The above modifications largely follow the generic transformation. We find that these modifications are sufficient for a secure Cubic-ABA, just as all known ABA protocols that can be transformed into their secure RABA counterparts (shown in [57]). Unexpectedly, we find for Quadratic-ABA, however, there is a subtle liveness issue for round 0. Suppose f correct replicas propose 1 and $f+1$ correct replicas propose 0. The f replicas directly broadcast vote$_0(1)$, main-vote$_0(1)$, and final-vote$_0(1)$. Even if the $f+1$ correct replicas that proposed 0 may later repropose 1, they may have already sent vote$_0(0)$, main-vote$_0(0)$, and final-vote$_0(0)$. In this case, no correct replica will accept final-vote$_0(1)$, as they do not receive $f+1$ main-vote$_0(1)$. The issue is, in essence, caused by the fact that each correct replica accepts a main-vote$_r(v)$ message only if it has previously received $f+1$ vote$_r(v)$ messages, and each correct replica accepts a final-vote$_r(v)$ message only if it has previously received $f+1$ main-vote$_r(v)$ messages. We visualize the issue via an example below.

To resolve the above issue, we introduce another change to round 0 of the protocol. In particular, we relax the conditions for round 0: for each main-vote$_r(v)$ (ln 28) and final-vote$_r(v)$ (ln 32), a correct replica accepts it as long as $v \in bset_0$. With

```
01 initialization
02   r ← 0                                                    {round}
03 func propose(v)
04   broadcast-vote(v)
05   start round 0
06 func repropose(v)
07   broadcast-vote(v)
08 func broadcast-vote(v)
09   if pre-vote_0(v) has not been sent, broadcast pre-vote_0(v)
10   if v = 1
11     bset_0 ← bset_0 ∪ {1}
12     if vote_0() has not been sent, broadcast vote_0(1)
13     if main-vote_0() has not been sent, broadcast main-vote_0(1)
14     if final-vote_0() has not been sent, broadcast final-vote_0(1)
15 round r
16   if r > 0, broadcast pre-vote_r(iv_r)
17   upon receiving pre-vote_r(v) from f+1 replicas
18     if pre-vote_r(v) has not been sent, broadcast pre-vote_r(v)
19   upon receiving pre-vote_r(v) from 2f+1 replicas
20     bset_r ← bset_r ∪ {v}
21   wait until bset_r ≠ ∅
22     if vote_r() has not been sent
23       broadcast vote_r(v) where v ∈ bset_r
24   upon receiving n−f vote_r() such that for each received
     vote_r(b), b ∈ bset_r
25     if there are n−f vote_r(v)
26       broadcast main-vote_r(v)
27     else broadcast main-vote_r(∗)
28   upon receiving n−f main-vote_r() such that for each
     main-vote_r(v): 1) if r = 0, v ∈ bset_r, 2) if r > 0, at least f+1
     vote_r(v) have been received; for each main-vote_r(∗), bset_r = {0,1}
29     if there are n−f main-vote_r(v)
30       broadcast final-vote_r(v)
31     else broadcast final-vote_r(∗)
32   upon receiving n−f final-vote_r() such that for each
     final-vote_r(v), 1) if r = 0, v ∈ bset_r, 2) if r > 0, at least f+1
     main-vote_r(v) have been received; for each final-vote_r(∗), bset_r =
     {0,1}
33     if there are n−f final-vote_r(v)
34       iv_{r+1} ← v, decide v
35     else if there are only final-vote_r(v) and final-vote_r(∗)
36       iv_{r+1} ← v
37     else
38       if r = 0, iv_{r+1} ← 1              {coin in the first round is 1}
39       else iv_{r+1} ← Random()                    {obtain local coin}
40   r ← r+1
```

Figure 5: The Quadratic-RABA protocol. The code for p_i.

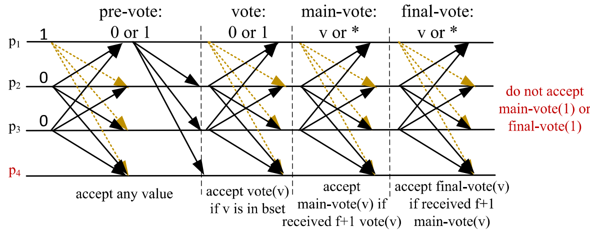

Figure 6: Subtle challenge when converting Quadratic-ABA to Quadratic-RABA. Votes for 1 are marked in dashed lines. We illustrate a scenario to show that if we only apply the three modifications mentioned in Sec. 5, the protocol may not terminate. In this example, we have four replicas—p_1 to p_4, and p_4 simply crashes. In round 0, the input of p_1 is 1, so it broadcasts the $\mathsf{pre\text{-}vote}_0(1)$, $\mathsf{vote}_0(1)$, $\mathsf{main\text{-}vote}_0(1)$, and $\mathsf{final\text{-}vote}_0(1)$ messages simultaneously. Meanwhile, p_2 and p_3 propose 0, so they send $\mathsf{pre\text{-}vote}_0(0)$ to all replicas. After p_1 receives $f+1$ $\mathsf{pre\text{-}vote}_0(0)$, it will also send a $\mathsf{pre\text{-}vote}_0(0)$ message to all replicas. However, it will not send $\mathsf{vote}_0(0)$, $\mathsf{main\text{-}vote}_0(0)$, or $\mathsf{final\text{-}vote}_0(0)$ messages, as it only sends any of these messages once. As p_2 and p_3 receive $n-f$ $\mathsf{pre\text{-}vote}_0(0)$ messages, they will send $\mathsf{vote}_0(0)$, $\mathsf{main\text{-}vote}_0(0)$, and $\mathsf{final\text{-}vote}_0(0)$. Hence, every replica receives two $\mathsf{final\text{-}vote}_0(0)$ messages and one $\mathsf{final\text{-}vote}_0(1)$. For p_2 or p_3 to proceed to round 1, it needs to *accept* $n-f$ $\mathsf{final\text{-}vote}_0()$ messages. According to our Quadratic-ABA specification, each replica accepts a $\mathsf{final\text{-}vote}_0(v)$ message only if it has received $f+1$ $\mathsf{main\text{-}vote}_0(v)$ and $bset_r = \{0,1\}$. Clearly, p_2 and p_3 accept $\mathsf{final\text{-}vote}_0(0)$ but will not accept $\mathsf{final\text{-}vote}_0(1)$, because they fail to receive $f+1$ $\mathsf{main\text{-}vote}_0(1)$ messages (they might have $bset_r = \{0,1\}$ if at least one of them reproposes). Therefore, the protocol may not terminate in this case.

this modification, the set of $f+1$ correct replicas that proposed 0 will repropose 1, so every correct replica will eventually add 1 in $bset_0$. Hence, every correct replica will eventually accept $\mathsf{main\text{-}vote}_0(1)$ and $\mathsf{final\text{-}vote}_0(1)$. Our result underlines the subtlety of constructing RABA from ABA and the importance of a full proof for a new protocol (proof in our full paper [59]).

6 The WaterBear Family

This section describes our asynchronous BFT protocols—WaterBear (WaterBear-C and WaterBear-Q), and WaterBear-QS (WaterBear-QS-C, and WaterBear-QS-Q). All the protocols are quantum secure, and WaterBear-C and WaterBear-Q rely on authenticated channels only.

6.1 The WaterBear Protocols

WaterBear follows the PACE paradigm but uses the trick in EPIC [44] to avoid the usage of threshold encryption (needed for achieving adaptive security). In particular, WaterBear uses *r-broadcast* and *r-deliver* primitives of Bracha's broadcast,

```
01 upon selecting m_i for p_i using the technique of EPIC
02    r-broadcast([e,i], m_i) for RBC_i
03 upon r-deliver([e,j], m_j) for RBC_j
04    if RABA_j has not been started
05       propose([e,j], 1) for RABA_j
06    else
07       repropose([e,j], 1) for RABA_j
08 upon delivery of n − f RBC instances
09    for RABA instances that have not been started
10       propose([e,j], 0)
11 upon decide([e,j], v) for any value v for all RABA instances
12    let S be set of indexes for RABA instances that decide 1
13    wait until r-deliver([e,j], m_j) for all RABA_j where j ∈ S
14       a-deliver(∪_{j∈S} {m_j})
```

Figure 7: The WaterBear family. The code for replica p_i in epoch e. WaterBear uses the EPIC technique to select transactions.

and *propose*, *repropose* and *decide* primitives of WaterBear RABA. Figure 7 depicts the pseudocode of WaterBear. Earlier asynchronous BFT protocols such as HoneyBadger and BEAT use threshold encryption to allow parallel, random transaction selection while achieving censorship resilience (liveness); however, no efficient adaptively secure threshold encryption was known. EPIC thus proposes a new crypto-free transaction selection strategy: replicas select random transactions *in plaintext* for most epochs, and to achieve liveness, they periodically switch to the first-in, first-out (FIFO) selection, where replicas maintain a log of transactions according to the order that transactions are received and select as input the first transaction group in the buffer. In WaterBear, we use this trick to avoid threshold encryption.

Following the PACE paradigm, for each epoch e, WaterBear consists of n parallel RBC instances and n parallel RABA instances. In the RBC phase, each replica p_i *r-broadcasts* a proposal m_i for RBC_i. If p_i *r-delivers* a proposal from RBC_j, it proposes 1 for RABA_j. Upon delivery of $n-f$ RBC instances, instead of waiting for $n-f$ RABA instances to terminate, p_i proposes 0 for all RABA instances that have not been started. If p_i later delivers a proposal from some RBC_j, it has proposed 0 for RABA_j, and has not terminated RABA_j, it reproposes 1 for RABA_j. We let S be the set of indexes where RABA_j decides 1. When all RABA instances terminate and all RBC_i ($i \in S$) instances are delivered, p_i *a-delivers* $\cup_{j\in S}\{m_j\}$. The security of WaterBear directly follows from that of the PACE paradigm. As we propose two RABA protocols Cubic-RABA and Quadratic-RABA. We use Cubic-RABA to build WaterBear-C and Quadratic-RABA to build WaterBear-Q.

6.2 The WaterBear-QS Protocols

We now describe WaterBear-QS, also consisting of two asynchronous BFT protocols. We use Cubic-RABA to build WaterBear-QS-C and Quadratic-RABA to build WaterBear-QS-Q. The difference between WaterBear and WaterBear-QS

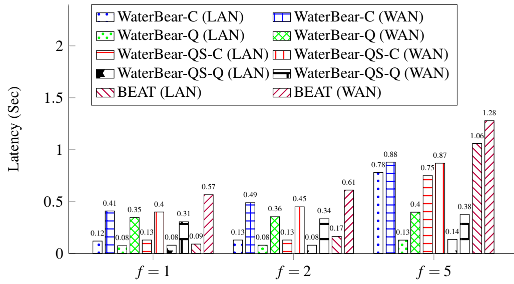

(a) Latency for $b = 1$.

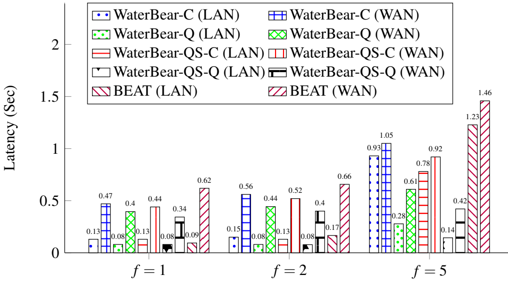

(b) Latency for $b = 100$.

Figure 8: Latency of the protocols.

is that WaterBear-QS additionally uses hash functions (used in CT RBC [14]) to reduce the communication complexity of RBC. Jumping ahead, we show the modification leads to a dramatic performance improvement.

7 Implementation and Evaluation

Implementation. We implemented WaterBear-C, WaterBear-Q, WaterBear-QS-C, and WaterBear-QS-Q in a new Golang library. For comparison, we choose to implement BEAT[2] [30] in our library. We implemented a new version of BEAT, replacing MMR ABA with Cobalt-ABA, as Cobalt ABA addressed the liveness issue of MMR. Our implementation has been made publicly available[3] and involves more than 11,000 LOC for the protocol implementations and about 1,000 LOC for evaluation.

All the protocols use authenticated channels and WaterBear-QS additionally uses hash functions. We use HMAC to realize authenticated channels. We use SHA256 as the hash function. We use gRPC as the communication library.

All our implemented asynchronous protocols use RBC in their RBC phases, and WaterBear-C and WaterBear-QS-C additionally use RBC in the ABA phase. For WaterBear-C and WaterBear-Q, we use Bracha's broadcast in the RBC phase. For WaterBear-QS-C and WaterBear-QS-Q, we use CT RBC [14] (using erasure coding and hash functions) in the RBC phase. In ABA phases of WaterBear-C and WaterBear-QS-C, we directly use Bracha's broadcast because there is no bulk data (and no need to use erasure coding). To implement CT RBC, we use a Golang Reed-Solomon code library[4].

There are several reasons we chose BEAT as the baseline protocol. First, BEAT is one of the most efficient open-source asynchronous BFT implementations. As shown in PACE [57], BEAT is more efficient than Dumbo [40] for $n \leq 46$. Second, all WaterBear protocols achieve adaptive security, and EPIC is the only known adaptively secure asynchronous BFT protocol implemented. It is shown that BEAT significantly outperforms EPIC in both LAN and WAN settings [44]. Hence, if we show the performance difference between BEAT and our protocols, we can argue which is the most efficient adaptively secure asynchronous BFT protocol among EPIC and WaterBear protocols. We do not attempt to compare our protocols with other BFT protocols in Table 1, as those protocols neither achieve adaptive nor quantum security, relying on PKC and trusted setup. *Indeed, our goal is not to claim WaterBear protocols are the most efficient asynchronous BFT protocols, but we aim at refuting the conventional wisdom that asynchronous BFT protocols cannot match the security guarantees of partially synchronous protocols while preserving performance.*

Overview of evaluation. We evaluate the performance of our protocols on Amazon EC2 utilizing up to 61 virtual machines (VMs). We consider both LAN and WAN settings. In the LAN setting, the replicas are run in the same region of EC2—US Virginia. In the WAN setting, the replicas are evenly distributed in four different regions: us-west-2 (Oregon, US), us-east-2 (Ohio, US), ap-southeast-1 (Singapore), and eu-west-1 (Ireland). The lowest one-way latency (resp., the highest latency) is 24.5 ms for Ohio-Oregon (resp., 90 ms for Ireland-Singapore). We use both *t2.medium* and *m5.xlarge* instances for our evaluation. The *t2.medium* type has two virtual CPUs and 4GB memory and the *m5.xlarge* has four virtual CPUs and 16GB memory. Unless otherwise mentioned, we use *m5.xlarge* instances by default. We conduct the experiments under different network sizes and contention levels (batch size). We use f to denote the network size; in each experiment, we use $3f+1$ replicas in total. We let b denote the contention level. In particular, each replica proposes b transactions in each epoch. For each experiment, we vary the batch size b from 1 to 25,000. For each experiment, we run the tests 10 times and compute the average (for both throughput and latency). We evaluate the performance of the protocols with two different transaction sizes—100 bytes by default and also 250 bytes.

We assess the protocols under failure-free and failure scenarios. While our failure-case evaluation is not the first such evaluation for asynchronous BFT protocols, the testbed we built aims to be comprehensive, encompassing realistic failure and attack scenarios we can envision. We roughly summarize our main results in the following:

[2] https://github.com/fififish/beat

[3] https://github.com/fififish/waterbear

[4] https://github.com/klauspost/reedsolomon

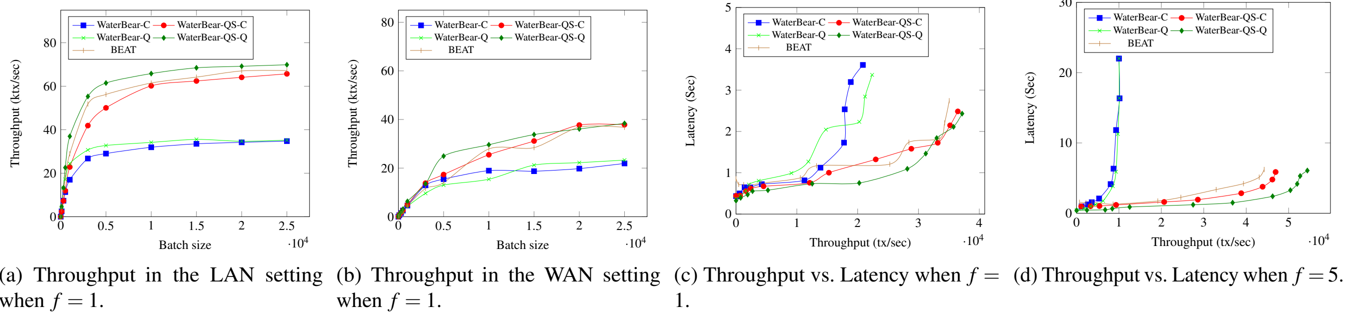

(a) Throughput in the LAN setting when $f = 1$.

(b) Throughput in the WAN setting when $f = 1$.

(c) Throughput vs. Latency when $f = 1$.

(d) Throughput vs. Latency when $f = 5$.

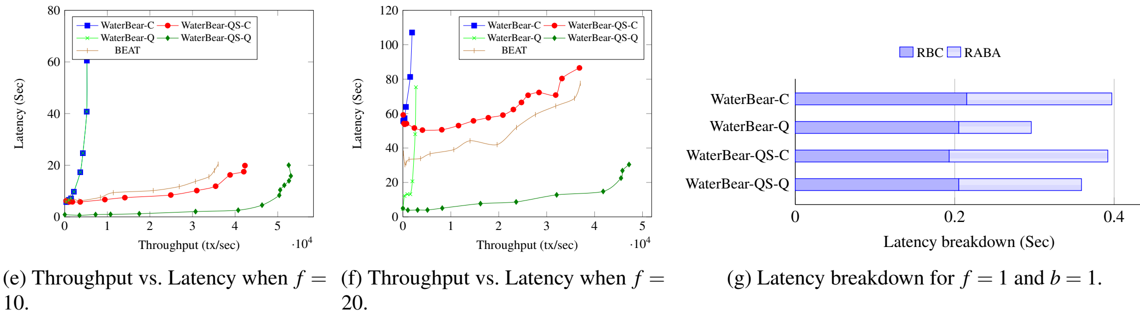

(e) Throughput vs. Latency when $f = 10$.

(f) Throughput vs. Latency when $f = 20$.

(g) Latency breakdown for $f = 1$ and $b = 1$.

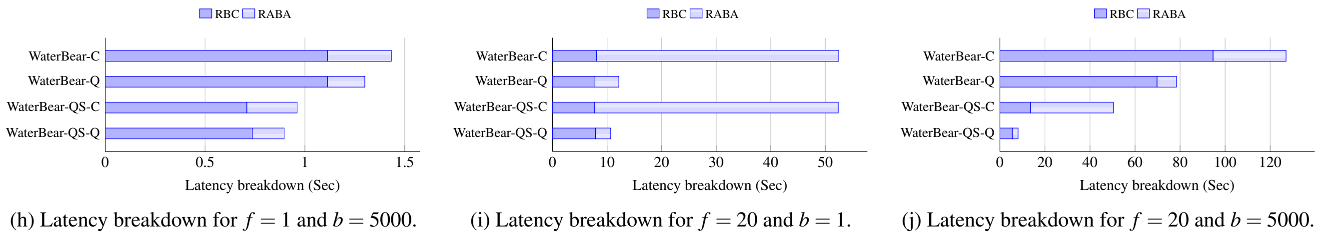

(h) Latency breakdown for $f = 1$ and $b = 5000$.

(i) Latency breakdown for $f = 20$ and $b = 1$.

(j) Latency breakdown for $f = 20$ and $b = 5000$.

Figure 9: Throughput vs. latency and latency breakdown on m5.xlarge instances for $f = 1$ to $f = 20$.

- The WaterBear protocols using Quadratic-RABA (WaterBear-QS-Q and WaterBear-Q) are much more efficient than the protocols using Cubic-RABA (WaterBear-QS-C and WaterBear-C), as Quadratic-RABA has $O(n^2)$ messages and much fewer steps than Cubic-RABA (with $O(n^3)$ messages). The result justifies the importance of designing Quadratic-RABA and Quadratic-ABA.
- The quantum secure WaterBear protocols (WaterBear-QS-C and WaterBear-QS-Q) drastically outperform their counterparts assuming authenticated channels only (WaterBear-C and WaterBear-Q), as the RBC used for WaterBear-QS-C and WaterBear-QS-Q is more bandwidth-efficient than that for WaterBear-C and WaterBear-Q. The finding highlights the cost of achieving security with authenticated channels only from quantum security for our protocols.
- Regarding latency, all WaterBear protocols have lower latency (under no contention) than BEAT. Regarding throughput, all our protocols, except WaterBear-QS-Q, share similar performance as BEAT.
- WaterBear-QS-Q consistently and significantly outpaces BEAT. For instance, when $n = 16$, WaterBear-QS-Q has about 1/8 the latency that of BEAT and 1.23x the throughput of BEAT. As n grows larger, the peak throughput of WaterBear-QS-Q is about 1.47x that of BEAT.
- All four protocols we propose are highly robust against various crash and Byzantine failures, just as BEAT.

7.1 Performance in Failure-Free Cases

Latency. We report the latency of the asynchronous protocols in both LAN and WAN settings for $f = 1, 2$, and 5 in Figure 8 with for $b = 1$ and 100. All WaterBear protocols consistently achieve lower latency than BEAT in both LAN and WAN environments, mainly because our protocols have a coin-free fast path. Among the protocols, WaterBear-QS-Q has consistently lower latency than all other protocols, as WaterBear-QS-Q has the lowest communication complexity among the WaterBear protocols. As f increases, the latency difference between WaterBear-QS-Q and other protocols becomes more visible. For instance, when $f = 5$ in the WAN setting, BEAT achieves 3.47x latency of that for WaterBear-QS-Q; in the LAN setting,

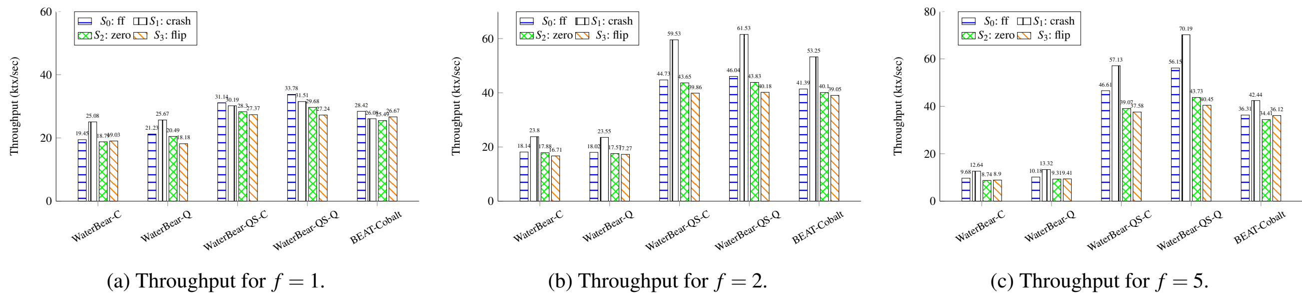

(a) Throughput for $f = 1$.

(b) Throughput for $f = 2$.

(c) Throughput for $f = 5$.

Figure 10: Performance of the protocols in failure scenarios.

the latency for BEAT is 8.78x of that for WaterBear-QS-Q.

Let us explain the latency results in more detail. First, take WaterBear-QS-Q in WANs for an example. In the optimistic mode (with no failures and synchronous networks), it takes four steps (3 for RBC and 1 for RABA due to the fast path in Quadratic-RABA) to terminate. When $f = 1$, the latency reported for WaterBear-QS-Q is 310 ms, which is consistent with the ping latency mentioned above (24.5 ms$\sim$90 ms). In contrast, WaterBear-QS-C in its optimistic mode has 6 steps, which justifies its latency of 570 ms.

Throughput and scalability. We report throughput and throughput vs. latency of all our implemented protocols in Figure 9 by varying the network size f from 1 to 20.

Our results show that the throughput of WaterBear-C and WaterBear-Q are consistently lower than the other protocols. As WaterBear-C (resp. WaterBear-Q) and WaterBear-QS-C (resp. WaterBear-QS-Q) differ in RBC only, RBC is clearly one performance bottleneck. The result highlights the overhead of achieving security with authenticated channels only.

We assess the throughput of all the protocols for $f = 1$ in WAN as depicted in Figure 9b: the peak throughput of WaterBear-QS-Q is slightly higher in most experiments. We also conduct a separate experiment in LANs, as shown in Figure 9a. Unlike the results in WANs, the throughput of BEAT in LANs is marginally higher than WaterBear-QS-C, and the throughput of WaterBear-QS-Q is marginally higher than BEAT: the peak throughput of BEAT is 2.3% higher than WaterBear-QS-C, and the peak throughput of WaterBear-QS-Q is 3.9% higher than BEAT. The peak throughput of WaterBear-QS-C is 65.7 ktx/sec in LANs and 37.8 ktx/sec in WANs, and the peak throughput of WaterBear-QS-Q is 69.9 ktx/sec in LANs and 38.4 ktx/sec in WANs.

When f increases, in general, WaterBear-QS-Q and WaterBear-QS-C outpace all the other protocols. The peak throughput of WaterBear-QS-C is higher than BEAT when $f = 5$ and $f = 10$ but lower when $f = 20$ only. Meanwhile, WaterBear-QS-Q is consistently more efficient than all other asynchronous protocols (higher throughput and lower latency). For instance, when $f = 10$, the peak throughput of WaterBear-QS-Q is 47.4% higher than BEAT. The reason is WaterBear-QS-Q uses a more communication-efficient RABA protocol.

We report the latency breakdown of the WaterBear protocols for $b = 1$ and 5000 in Figure 9g-9j. These experiments justify the design of our most efficient protocol—WaterBear-QS-Q (with CT RBC that is communication-efficient and Quadratic-RABA with asymptotically reduced per-round message complexity and concretely reduced number of steps). Indeed, as f increases, RABA dominates the latency and Quadratic-RABA performs indeed much better than Cubic-RABA; when b increases, RBC becomes the latency bottleneck and CT RBC outperforms Bracha's RBC significantly.

7.2 Performance under Failures

To assess the protocol performance under failures and attacks, we carefully design various experiments as follows. We focus on the five asynchronous protocols.

- S_0**: (failure-free)** All replicas are correct. S_0 is the baseline scenario used to compare with failure scenarios.
- S_1**: (crash)** f replicas crash by not participating in the protocols.
- S_2**: (Byzantine; keep voting 0)** We control all f faulty replicas to keep voting 0 in each step of (R)ABA. For all protocols, doing so would intuitively make fewer (R)ABA instances to decide 1 and would likely decrease the throughput of the protocols. We aim to observe the throughput reduction compared to failure-free scenarios.
- S_3**: (Byzantine; flipping the (R)ABA input)** We let f replicas exhibit Byzantine behavior in the (R)ABA phase. The strategy is to vote for a flipped value in (R)ABA. In other words, in each (R)ABA step, each Byzantine replica inputs $\bar{b}$ when it should have input b. Doing so could potentially force each (R)ABA instance to experience more steps to terminate for all five protocols. For WaterBear and WaterBear-QS, the strategy would, at first glance, likely be more fruitful. For both protocols, a RABA instance may terminate in round 0, thanks to the biased validity property of RABA. The flipping strategy illustrated above may make them not decide in round 0 and force them to enter the second round of RABA, where the two protocols start to query the local coins.

We assess the performance for $f = 1$ (Figure 10a), $f = 2$ (Figure 10b), and $f = 5$ (Figure 10c).

Performance under crash failures (S_1). The throughput of all the five protocols implemented under crash failures is higher than that in the failure-free case, except for $f = 1$, where all protocols share similar performance between the two scenarios. Our result echoes those of previous works. The reason is that under crash failures, the network bandwidth consumption is much lower (about 33% lower) than in the failure-free case. Note that when $f = 1$, the network bandwidth consumption is not as dominating as in other cases; hence, the performance difference among the protocols for $f = 1$ is less visible.

Performance under Byzantine failures (S_2 and S_3). The performance of all the protocols under Byzantine failures is slightly lower than that in the failure-free scenario and the crash failure scenario. WaterBear-QS-C and WaterBear-QS-Q suffer from slightly higher performance degradation under Byzantine failures compared to BEAT. The higher performance degradation is due to the use of local coins. As replicas start to use local coins in round $r > 0$, the RABA protocol may decide in more rounds. In all cases, WaterBear-QS-C and WaterBear-QS-Q remain more efficient than BEAT.

The difference between S_2 and S_3 is that faulty replicas broadcast 0 in S_2 but broadcast the flipped value in S_3. For BEAT, the performance in S_3 is higher for $f = 1$ and $f = 5$ but lower for $f = 2$; the difference in all the cases is not significant though. In contrast, for WaterBear-QS-C and WaterBear-QS-Q, the performance in S_3 is consistently lower than S_2, showing that the flipping strategy in S_3 works slightly better than that in S_2.

8 Conclusion

This paper designs and implements a family of practical asynchronous BFT protocols matching the security guarantees of their partially synchronous counterparts. Our experiments demonstrate that our protocols are efficient in both failure and failure-free scenarios. In particular, one of our protocols, WaterBear-QC-Q, consistently outperforms the state-of-the-art asynchronous protocols with much weaker security guarantees. We also build in different settings more efficient ABA and RABA protocols that can be used to improve various high-level Byzantine-resilient protocols. Our work, for the first time, shows that the strongest security models and high performance can co-exist for asynchronous BFT.

Acknowledgment

This work was supported in part by the National Key R&D Program of China under 2022YFB2701500 and 2022YFB2701700, the National Natural Science Foundation of China under 62272043 and 92267203, Beijing Natural Science Foundation under M23015, Major Program of Shandong Provincial Natural Science Foundation for the Fundamental Research under ZR2022ZD03, and the Piloting Fundamental Research Program for the Integration of Scientific Research, Education and Industry of Qilu University of Technology (Shandong Academy of Sciences) 2022XD001.

References

[1] Ittai Abraham, Danny Dolev, and Joseph Y. Halpern. An almost-surely terminating polynomial protocol for asynchronous byzantine agreement with optimal resilience. PODC, 2008.

[2] Ittai Abraham, Philipp Jovanovic, Mary Maller, Sarah Meiklejohn, Gilad Stern, and Alin Tomescu. Reaching consensus for asynchronous distributed key generation. In *PODC*, 2021.

[3] Nicolas Alhaddad, Sourav Das, Sisi Duan, Ling Ren, Mayank Varia, Zhuolun Xiang, and Haibin Zhang. Balanced byzantine reliable broadcast with near-optimal communication and improved computation. In *PODC*, pages 399–417, 2022.

[4] Nicolas Alhaddad, Mayank Varia, and Haibin Zhang. High-threshold avss with optimal communication complexity. In *FC*, 2021.

[5] Yair Amir, Brian Coan, Jonathan Kirsch, and John Lane. Prime: Byzantine replication under attack. *TDSC*, 8(4):564–577, 2011.

[6] Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, Srinivasan Muralidharan, Chet Murthy, Binh Nguyen, Manish Sethi, Gari Singh, Keith Smith, Alessandro Sorniotti, Chrysoula Stathakopoulou, Marko Vukolić, Sharon Weed Cocco, and Jason Yellick. Hyperledger fabric: A distributed operating system for permissioned blockchains. 2018.

[7] Laasya Bangalore, Ashish Choudhury, and Arpita Patra. The power of shunning: Efficient asynchronous byzantine agreement revisited. *J. ACM*, 2020.

[8] Michael Ben-Or. Another advantage of free choice: Completely asynchronous agreement protocols (extended abstract). In *PODC*, pages 27–30, 1983.

[9] Michael Ben-Or, Boaz Kelmer, and Tal Rabin. Asynchronous secure computations with optimal resilience. In *PODC*, pages 183–192. ACM, 1994.

[10] Gabriel Bracha. An asynchronous [(n-1)/3]-resilient consensus protocol. In *PODC*, pages 154–162. ACM, 1984.

[11] Gabriel Bracha. Asynchronous byzantine agreement protocols. *Information and Computation*, 75(2):130–143, 1987.

[12] Christian Cachin. Yet another visit to paxos. IBM Research Report RZ 3754, 2010.

[13] Christian Cachin and Jonathan A Poritz. Secure intrusion-tolerant replication on the internet. In *DSN*, pages 167–176. IEEE, 2002.

[14] Christian Cachin and Stefano Tessaro. Asynchronous verifiable information dispersal. In *SRDS*, pages 191–201. IEEE, 2005.

[15] Christian Cachin and Marko Vukolic. Blockchain consensus protocols in the wild. In *DISC*, 2017.

[16] Ran Canetti, Uri Feige, Oded Goldreich, and Moni Naor. Adaptively secure multi-party computation. In *STOC*, 1996.

[17] Ran Canetti and Tal Rabin. Fast asynchronous byzantine agreement with optimal resilience. In *STOC*, volume 93, pages 42–51. Citeseer, 1993.

[18] Miguel Castro. Practical byzantine fault tolerance. PhD thesis, 2001.

[19] Miguel Castro and Barbara Liskov. Practical Byzantine fault tolerance and proactive recovery. *TOCS*, 20(4):398–461, 2002.

[20] Lily Chen, Lily Chen, Stephen Jordan, Yi-Kai Liu, Dustin Moody, Rene Peralta, Ray A Perlner, and Daniel Smith-Tone. *Report on post-quantum cryptography*, volume 12. US Department of Commerce, National Institute of Standards and Technology, 2016.

[21] Allen Clement, Edmund L Wong, Lorenzo Alvisi, Michael Dahlin, and Mirco Marchetti. Making byzantine fault tolerant systems tolerate byzantine faults. In *NSDI*, volume 9, pages 153–168, 2009.

[22] Miguel Correia, Nuno Ferreira Neves, and Paulo Veríssimo. From consensus to atomic broadcast: Time-free byzantine-resistant protocols without signatures. *Comput. J.*, 49(1):82–96, 2006.

[23] Tyler Crain. Two more algorithms for randomized signature-free asynchronous binary byzantine consensus with t<n/3 and o(n^2) messages and O(1) round expected termination. *CoRR*, abs/2002.08765, 2020.

[24] Ronald Cramer, Ivan Damgård, Stefan Dziembowski, Martin Hirt, and Tal Rabin. Efficient multiparty computations secure against an adaptive adversary. In *EUROCRYPT*, 1999.

[25] George Danezis, Lefteris Kokoris-Kogias, Alberto Sonnino, and Alexander Spiegelman. Narwhal and tusk: A dag-based mempool and efficient bft consensus. EuroSys '22.

[26] Sourav Das, Zhuolun Xiang, and Ling Ren. Asynchronous data dissemination and its applications. In *CCS*, pages 2705–2721, 2021.

[27] Sourav Das, Thomas Yurek, Zhuolun Xiang, Andrew Miller, Lefteris Kokoris-Kogias, and Ling Ren. Practical asynchronous distributed key generation. *IEEE Symposium on Security and Privacy*, 2022.

[28] Shlomi Dolev and Ziyu Wang. Sodsbc: Stream of distributed secrets for quantum-safe blockchain. In *2020 IEEE International Conference on Blockchain (Blockchain)*, pages 247–256. IEEE, 2020.

[29] Sisi Duan, Hein Meling, Sean Peisert, and Haibin Zhang. BChain: Byzantine replication with high throughput and embedded reconfiguration. In *OPODIS*, pages 91–106, 2014.

[30] Sisi Duan, Michael K Reiter, and Haibin Zhang. BEAT: Asynchronous BFT made practical. In *CCS*, pages 2028–2041. ACM, 2018.

[31] Sisi Duan, Xin Wang, and Haibin Zhang. Fin: Practical signature-free asynchronous common subset and bft in constant time. Cryptology ePrint Archive, Paper 2023/154, 2023. https://eprint.iacr.org/2023/154.

[32] Sisi Duan and Haibin Zhang. Foundations of dynamic bft. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1317–1334. IEEE, 2022.

[33] Sisi Duan and Haibin Zhang. Recent progress on bft in the era of blockchains. *National Science Review*, 9(10):nwac132, 2022.

[34] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *JACM*, 35(2):288–323, 1988.

[35] Paul Feldman and Silvio Micali. Optimal algorithms for byzantine agreement. In *STOC*, 1988.

[36] Michael J Fischer, Nancy A Lynch, and Michael S Paterson. Impossibility of distributed consensus with one faulty process. Technical report, Massachusetts Inst of Tech Cambridge lab for Computer Science, 1982.

[37] Neil Giridharan, Lefteris Kokoris-Kogias, Alberto Sonnino, and Alexander Spiegelman. Bullshark: Dag bft protocols made practical. *ACM CCS*, 2022.

[38] Rachid Guerraoui, Nikola Knežević, Vivien Quéma, and Marko Vukolić. The next 700 bft protocols. *ACM Transactions on Computer Systems*, 32(4):12:1–12:45, 2015.

[39] Bingyong Guo, Yuan Lu, Zhenliang Lu, Qiang Tang, Jing Xu, and Zhenfeng Zhang. Speeding dumbo: Pushing asynchronous bft closer to practice. *NDSS*, 2022.

[40] Bingyong Guo, Zhenliang Lu, Qiang Tang, Jing Xu, and Zhenfeng Zhang. Dumbo: Faster asynchronous bft protocols. In *CCS*, 2020.

[41] James Hendricks, Shafeeq Sinnamohideen, Gregory R Ganger, and Michael K Reiter. Zzyzx: Scalable fault tolerance through byzantine locking. In *DSN*, pages 363–372. IEEE, 2010.

[42] Idit Keidar, Eleftherios Kokoris-Kogias, Oded Naor, and Alexander Spiegelman. All you need is dag. In *PODC*, 2021.

[43] Eleftherios Kokoris Kogias, Dahlia Malkhi, and Alexander Spiegelman. Asynchronous distributed key generation for computationally-secure randomness, consensus, and threshold signatures. CCS, 2020.

[44] Chao Liu, Sisi Duan, and Haibin Zhang. Epic: Efficient asynchronous bft with adaptive security. In *DSN*, 2020.

[45] Chao Liu, Sisi Duan, and Haibin Zhang. MiB: Asynchronous BFT with more replicas. *CoRR*, abs/2108.04488, 2021.

[46] Ethan MacBrough. Cobalt: Bft governance in open networks. *arXiv preprint arXiv:1802.07240*, 2018.

[47] Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. The honey badger of bft protocols. In *CCS*, pages 31–42. ACM, 2016.

[48] Henrique Moniz, Nuno Ferreria Neves, Miguel Correia, and Paulo Verissimo. Ritas: Services for randomized intrusion tolerance. *TDSC*, 8(1):122–136, 2008.

[49] Achour Mostéfaoui, Hamouma Moumen, and Michel Raynal. Signature-free asynchronous binary byzantine consensus with t < n/3, o(n2) messages, and O(1) expected time. *J. ACM*, 62(4):31:1–31:21, 2015.

[50] A. Patra, A. Choudhury, and C.P. Rangan. Asynchronous byzantine agreement with optimal resilience. *Distrib. Comput.*, 27:111–146, 2014.

[51] M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *JACM*, 27(2):228–234, April 1980.

[52] Michael O Rabin. Randomized byzantine generals. In *SFCS*, pages 403–409. IEEE, 1983.

[53] Alan T Sherman, Farid Javani, Haibin Zhang, and Enis Golaszewski. On the origins and variations of blockchain technologies. *IEEE Security & Privacy*, 17(1):72–77, 2019.

[54] Gilad Stern and Ittai Abraham. Information theoretic hotstuff. In *OPODIS*, 2018.

[55] Marko Vukolić. The quest for scalable blockchain fabric: Proof-of-work vs. bft replication. In *International workshop on open problems in network security*, pages 112–125. Springer, 2015.

[56] Xin Wang, Sisi Duan, James Clavin, and Haibin Zhang. Bft in blockchains: From protocols to use cases. *ACM Computing Surveys (CSUR)*, 54(10s):1–37, 2022.

[57] Haibin Zhang and Sisi Duan. PACE: fully parallelizable bft from reproposable byzantine agreement. In *CCS 2022*.

[58] Haibin Zhang, Sisi Duan, Chao Liu, Boxin Zhao, Xuanji Meng, Shengli Liu, Yong Yu, Fangguo Zhang, and Liehuang Zhu. Practical asynchronous distributed key generation: Improved efficiency, weaker assumption, and standard model. *DSN*, 2023.

[59] Haibin Zhang, Sisi Duan, Boxin Zhao, and Liehuang Zhu. Waterbear: Practical asynchronous bft matching security guarantees of partially synchronous bft. Cryptology ePrint Archive, Paper 2022/021, 2022.

[60] Haibin Zhang, Chao Liu, and Sisi Duan. How to achieve adaptive security for asynchronous bft? *Journal of Parallel and Distributed Computing*, 169:252–268, 2022.

A Bracha's ABA

We present Bracha's ABA [10]. The pseudocode is shown in Figure 11. Bracha's ABA has three phases. In each phase, each replica broadcasts its value via a RBC instance, i.e., there are n parallel RBC instances in each of the three phases. As the underlying RBC has $O(n^2)$ messages and 4 steps, Bracha's ABA has $O(n^3)$ messages and 12 steps in each round.

In Bracha's ABA, every replica maintains a set *vset* containing *valid* values. In each phase, every replica only accepts messages that carry valid values. The valid values *vset* must be congruent with the values each replica receives from the previous phase (or the last phase of the previous round). In the first phase of round 0, both 0 and 1 are considered valid. In the second and third phases, a value is added to *vset* only if the replica receives the value from enough replicas.

In the first phase, every replica p_i *r-broadcasts* a $\textsf{pre-vote}_r(iv_r)$ message (ln 08), where iv_r is the input value of p_i for round r.

In the second phase, p_i waits for $n-f$ $\textsf{pre-vote}_r()$ messages such that for each $\textsf{pre-vote}_r(v)$, $v \in vset$. There are two cases:

- Ln 10-13: If p_i has received $n-f$ $\textsf{pre-vote}_r(v)$ for some $v \in \{0,1\}$, p_i decides v and sets both *vset* and iv_{r+1} as v. Replica p_i continues for one more round and terminates the protocol (up to either ln 10 or ln 25 before p_i decides some value again).

```
01 Initialization
02   r ← 0                                                    {round}
03 func propose(v)
04   iv_0 ← v
05   vset ← {0,1}        {valid binary values that will be accepted}
06   start round 0
07 round r
08   r-broadcast pre-vote_r(iv_r)                             {▷ phase 1}
09   upon r-delivering n − f pre-vote_r() such that for each
pre-vote_r(v), v ∈ vset                                       {▷ phase 2}
10     if there are n − f pre-vote_r(v)
11       decide v
12       iv_{r+1} ← v
13       vset ← {v}
14     else
15       v ← majority value in the set of pre-vote_r() messages
16     r-broadcast main-vote_r(v)
17   upon r-delivering n − f main-vote_r() such that for each
main-vote_r(v), v ∈ vset                                      {▷ phase 3}
18     if there are at least n/2 main-vote_r(v)
19       vset ← {v}
20     else
21       v ← {⊥}
22       vset ← {0,1}
23     r-broadcast final-vote_r(v)
24   upon r-delivering n − f final-vote_r() such that for each
final-vote_r(v), v ∈ vset; for each final-vote_r(∗), vset = {0,1}
25     if there are at least 2f + 1 final-vote_r(v)
26       decide v
27       iv_{r+1} ← v
28       vset ← {v}
29     else if there are f + 1 final-vote_r(v)
30       iv_{r+1} ← v
31       vset ← {0,1}
32     else
33       iv_{r+1} ← Random()                          {obtain local coin}
34       vset ← {0,1}
35   r ← r + 1
```

Figure 11: The Bracha's ABA protocol [10]. The code for p_i.

- Ln 14-15: Otherwise, p_i sets v as the majority value in the set of pre-vote$_r()$ messages it receives. The set *vset* is not changed, i.e., $vset = \{0,1\}$.

In both cases, p_i *r-broadcasts* a main-vote$_r(v)$ message (ln 16).

In the third phase, every replica p_i waits for $n-f$ valid main-vote$_r()$ messages (ln 17). There are two cases:

- Ln 18-19: If p_i receives at least $n/2$ main-vote$_r(v)$, it sets *vset* as $\{v\}$.
- Ln 20-22: Otherwise, p_i sets v as $*$ and *vset* as $\{0,1\}$.

In both cases, p_i *r-broadcasts* a final-vote$_r(v)$ message (ln 23). Then every replica waits for $n-f$ valid final-vote$_r()$ messages (ln 24). Note that final-vote$_r(*)$ is considered valid only if $vset = \{0,1\}$. There are three cases:

- Ln 25-28: If p_i receives at least $2f+1$ final-vote$_r(v)$, it decides v and sets iv_{r+1} as v. Replica p_i continues for one more round (up to either ln 10 or ln 25) and terminates the protocol.
- Ln 29-31: If p_i receives at least $f+1$ final-vote$_r(v)$, it sets iv_{r+1} as v and *vset* as $\{v\}$.
- Ln 32-34: Otherwise, p_i uses the local coin value as iv_{r+1} and *vset* as $\{0,1\}$, i.e., p_i accepts both 0 and 1 in the first phase of the following round.