

WaterBear: Asynchronous BFT with Information-Theoretic Security and Quantum Security

Haibin Zhang
Beijing Institute of Technology
haibin@bit.edu.cn

Sisi Duan*
Tsinghua University
duansisi@tsinghua.edu.cn

Boxin Zhao
Zhongguancun Laboratory
bx2965198@163.com

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

Abstract—Designing information-theoretically secure and quantum secure Byzantine fault-tolerant (BFT) protocols in asynchronous environments has been an elusive goal. Critically, all practical asynchronous BFT protocols require using common coins, but we do not have efficient setup-free common coin protocols that are unconditionally secure or quantum secure.

We design and implement WaterBear, a family of new asynchronous BFT protocols that are information-theoretically secure or quantum secure. Via extensive evaluation, we show that our protocols are efficient under both failure-free and failure scenarios, achieving comparable performance to the state-of-the-art asynchronous BFT protocols with much weaker security guarantees.

To achieve the goal, we have designed much more efficient asynchronous binary agreement (ABA) protocols from local coins and their reposable ABA counterparts. We have also built more efficient ABA protocols from weak common coins and perfect common coins. These ABA protocols can be readily used to improve various high-level Byzantine-resilient primitives, such as asynchronous distributed key generation and BFT assuming trusted setup.

I. 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 [21], [8], [71] and is used in various ways in hybrid blockchains. It is also well-known that BFT and Byzantine agreement (BA) are equivalent (possibility- and impossibility-wise), and both are fundamental building blocks for secure multi-party computation (MPC) achieving fairness and guaranteed output delivery [13], [54], [12], [27], [58]. This paper designs and implements the first practical information-theoretic (unconditionally secure) and quantum secure asynchronous BFT protocols, resolving long-standing open problems in fault-tolerant distributed computing and cryptography. Additionally, we propose a number of more efficient asynchronous binary agreement (ABA) protocols, including two ABA protocols from local coins, an ABA protocol from weak common coins, and an ABA protocol from common coins. These protocols can be readily used to improve various high-level protocols.

A. Background

Information-theoretic vs. computational security. Depending on the capacities of the adversary, any cryptographic or security protocols can be in one of the two models:

- *Computational security*, where the adversary is restricted to probabilistic polynomial-time (PPT).
- *Information-theoretic (IT) security*, where the adversary is unbounded.

Computationally secure protocols assume the hardness of some intractability problems (e.g., RSA, Diffie-Hellman). These mathematical problems may, in the future, be proven to be broken or weakened due to newly developed cryptanalysis techniques or some technological breakthrough (e.g., quantum computer). In contrast, IT security provides everlasting security without relying on any unproven intractability assumptions: IT protocols are not only quantum secure but future-proof. Moreover, IT protocols are natural candidates for building protocols secure under concurrent composition [31], [53], [29]. **PKC vs. no PKC vs. quantum security.** It is a major problem in fault-tolerant distributed computing to design various Byzantine-resilient protocols relying on no public key cryptography (PKC). For instance, the first practical partially synchronous BFT without PKC is PBFT [26]. Many other Byzantine-resilient protocols in various settings (e.g., secure causal atomic broadcast, atomic register) are known [38], [50], [40]. It is, however, an open problem if one could design practical PKC-free BFT in asynchronous environments. Existing asynchronous BFT protocols assume various primitives from public-key cryptography (e.g., common coins).

There are many reasons why one may favor symmetric cryptography. First, symmetric cryptography is based on basic and well-studied primitives (blockciphers and hashes). Second, PKC is many orders of magnitude slower than symmetric cryptography. Besides, classic symmetric cryptography is believed to be quantum-resistant; in contrast, only PKC schemes based on certain mathematical problems (e.g., lattices) are quantum-resistant, but we do not know how to efficiently realize some key building blocks (e.g., common coins) needed for existing asynchronous BFT from these problems.

Note that IT security implies quantum security; the reverse does not hold, for one could use non-quantum technologies to attack systems.

B. Challenges of IT and Quantum Secure Asynchronous BFT

We divide asynchronous BFT protocols *ever implemented* into several categories: 1) the BKR framework of Ben-Or, Kelmer, and Rabin [13], including HoneyBadgerBFT[61], BEAT [41], and EPIC [56], 2) the CKPS framework of Cachin, Kusawe, Petzold, and Shoup [17], including SINTRA [19], Dumbo [49], and Speeding Dumbo [48], 3) the CNV paradigm

*Corresponding author.

	IT secure	no pkc	quantum secure	no trusted setup	adaptive	WAN
SINTRA [19]						
RITAS [62]		✓	✓	✓	✓	
HoneyBadger [61]; BEAT [41]						✓
Dumbo [49]; Speeding Dumbo [48]						✓
EPIC [56]					✓	✓
Tusk [35]; Bullshark [46]						✓
PACE [72]						✓
WaterBear-QS (this work)		✓	✓	✓	✓	✓
WaterBear (this work)	✓	✓	✓	✓	✓	✓

Table I: Comparison of efficient asynchronous BFT protocols.

of Correia, Neves, and Veríssimo [32], with RITAS as an implementation [62], 4) the DKSS framework of Danezis, Kokoris Kogias, Sonnino, and Spiegelman [52], implemented in Tusk [35] and Bullshark [46], and 5) the PACE framework of Zhang and Duan [72]. While the frameworks mark significant milestones in developing practical asynchronous BFT protocols, the instantiations derived from these frameworks rely critically on cryptographic common coins or Byzantine agreement and are neither IT nor quantum secure. In fact, it is difficult to make them IT or just quantum secure, because they face a common hurdle—the inefficiency of IT or quantum secure common coin and asynchronous Byzantine agreement.

In the IT setting, the local coin based ABA protocols of Ben-Or [11] and Bracha [15] need an exponential expected time. While later ABA constructions using asynchronous verifiable secret sharing (AVSS) are much more efficient, these protocols are not yet practical. First, IT AVSS is notoriously difficult to build. Even in the statistical setting (allowing errors), such a primitive is overly complex. For instance, to build AVSS, the approach of Canetti and Rabin [24] needs to begin with an information checking protocol (resembling signatures but working in the IT setting), then asynchronous recoverable sharing, then asynchronous weak secret sharing, and finally AVSS. The improved approach of Patra, Choudhury, and Rangan [66] remains complex, following the route of information-checking protocol, then asynchronous weak commitment, and then AVSS. Second, the transformation from AVSS to ABA is equally expensive, requiring running n^2 AVSS instances to generate a *single* (weak) coin.

Even if relaxing to quantum security, we currently still lack efficient instantiations of common coin protocols based on quantum resistant primitives (e.g., lattices).

C. Our Contributions

A detour to IT and quantum secure BFT. One possible way of circumventing the inefficiency challenge of IT ABA is to use local coin based ABA instead. However, directly using local coin based ABA leads to protocols that terminate in exponential expected time and fail to scale [62].

This paper takes a detour and develops the PACE asynchronous BFT framework [72] in the IT setting and in the quantum security setting. PACE uses Byzantine reliable broadcast (RBC) and reproposable ABA (RABA) as building blocks and removes the two subphase bottleneck of the BKR paradigm, allowing RABA instances to run in a fully parallelizable manner. Syntactically, RABA is different from ABA in the sense that replicas are allowed to change their votes. The concrete instantiation of the PACE paradigm uses a common coin based RABA protocol built on the CKS threshold PRF scheme [18], achieving computational security

ABA (local coins)	messages/round	steps/round
Bracha’s ABA [15]	n^3	9 to 12
Cubic-ABA (this work)	n^3	5 to 7
Quadratic-ABA (this work)	n^2	4 or 5

Table II: 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.

and static security only.

PACE ensures that as long as there are just $f + 1$ replicas, instead of $2f + 1$ replicas, propose 1, then a correct replica that terminates for the instance would decide 1 and the corresponding transactions would be delivered. The powerful property implies a fast path for consensus: while PACE was designed for RABA with expected constant rounds, even with local coin based RABA, the protocol will on average terminate in a single RABA round with high probability.

Thus, our strategy is to *reduce IT BFT to RABA with local coins and then to ABA with local coins*. As reported in almost all asynchronous BFT protocols [49], [41], [72], ABA is the major performance bottleneck. Our first goal is to design efficient local coin based ABA protocols.

Efficient local coin based ABA. To our knowledge, only two local coin based ABA protocols have been proposed: Ben-Or’s ABA [11] assuming $n > 5f$, and Bracha’s ABA [14] assuming optimal resilience (the most efficient such protocol for nearly three decades).

Table II 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. It involves n parallel RBC instances in only one phase of the protocol. Accordingly, 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.

Quadratic-ABA adopts a new design, having 4 or 5 steps per round. The key is to replace n RBC instances (with $O(n^3)$ messages) used in Cubic-ABA and Bracha’s ABA with two-step all-to-all communication. As a result, Quadratic-ABA is the first local coin based ABA that achieves $O(n^2)$ message complexity per round.

Tackling a liveness issue for RABA. We have carefully designed Cubic-RABA and Quadratic-RABA that are as efficient as Cubic-ABA and Quadratic-ABA, respectively. Unlike prior transformations following a generic approach in [72], we identify and tackle a subtle liveness problem arising when transforming Quadratic-ABA to Quadratic-RABA. The issue we identify demonstrates the subtlety of transforming ABA to RABA, and once again underlines the importance of a full

protocol	reference implementation	RBC	RABA	authenticated channels
WaterBear	WaterBear-C	Bracha’s RBC [15]	Cubic-RABA (this paper)	HMAC*
	WaterBear-Q	Bracha’s RBC [15]	Quadratic-RABA (this paper)	HMAC*
WaterBear-QS	WaterBear-QS-C	CT RBC [20]	Cubic-RABA (this paper)	HMAC
	WaterBear-QS-Q	CT RBC [20]	Quadratic-RABA (this paper)	HMAC

Table III: WaterBear-QS and WaterBear instantiations. *HMAC is quantum secure but not IT secure; we simply used HMAC in our reference implementation for WaterBear to *demonstrate the overhead of WaterBear itself*, because all other protocols introduced in this paper use HMAC for authentication. As in PACE, both WaterBear-QS and WaterBear have a fast path allowing the protocol 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.

proof for fault-tolerant distributed protocols.

WaterBear: unbounded security and beyond. WaterBear develops a family of protocols utilizing the PACE paradigm. In particular, we use Cubic-RABA to build WaterBear-C and Quadratic-RABA to build WaterBear-Q. Besides the critical IT RABA protocols, WaterBear has two other major differences compared to the prior PACE instantiation. First, we use an IT RBC, the Bracha’s RBC, as the underlying RBC. Second, to achieve unconditional security and adaptive security, we use the technique of EPIC [56] to remove the threshold encryption scheme used in the PACE paradigm.

WaterBear assumes authenticated channels only and has *all* desirable properties a BFT protocol one could think of, being optimally resilient, achieving unconditional security and adaptive security, and not relying on trusted setup.

WaterBear-QS: no PKC and quantum security. We also present a family of asynchronous BFT protocols—WaterBear-QS, which does not achieve IT security but achieves quantum security for both safety and liveness properties. The only difference between WaterBear and WaterBear-QS is that WaterBear-QS additionally uses a collision-resistant hash function. 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.

What motivated us to design WaterBear-QS is a crucial observation from our evaluation that RBC is another performance bottleneck. One of our experiments shows that the size of transactions matters significantly for the throughput. For instance, we find the throughput using a 250B transaction is about half of the throughput for 100B, a somewhat surprising finding that has not been previously reported. As the ABA phase does not carry bulk data in both cases, the difference must be due to the RBC phase. Hence, WaterBear-QS uses the CT RBC leveraging hash functions to reduce the protocol communication [20]. As the hash function is the only cryptographic tool used, WaterBear-QS is quantum secure.

A new asynchronous BFT platform and extensive evaluation under failures. Starting from HoneyBadgerBFT, existing efficient asynchronous BFT protocols, including BEAT, Dumbo, and EPIC, use the HoneyBadgerBFT programming framework using Python. We instead build a new platform using Golang. Our platform implements WaterBear-C, WaterBear-Q, WaterBear-QS-C, WaterBear-QS-Q, and BEAT (one of the most efficient open-source asynchronous BFT libraries) [41], [1].

With deployment in 5 continents, we show that our protocols offer comparable performance as the state-of-the-art asynchronous BFT protocols, while achieving much stronger security (IT or quantum security, adaptive security, and no

trusted dealer needed). We also design and evaluate various failure and attack scenarios for the protocols implemented. Via extensive experiments, all of our protocols are shown to be highly robust against these failures and attacks.

While our primary goal is to build stronger BFT protocols, WaterBear-QS-Q does offer consistently and significantly better performance than BEAT (which is somewhat surprising).

ABA from weak common coins and perfect common coins. By extending the technique of Quadratic-ABA, we can obtain CC-ABA that works for both weak common coins (with a constant probability all correct replicas obtain the same coin) and perfect common coins. We describe ABA protocols using common coins in Table IV and Table V: in both cases, CC-ABA compares favorably with existing protocols. The results are important: significant research has focused on how to improve the concrete steps for ABA. Zhang and Duan recently showed that the concrete steps of ABA protocols are vital to the performance of asynchronous BFT: even a single step improvement in ABA, the resulting BFT protocol could be easily improved by, say, 2x [72]. CC-ABA with weak common coins can be used to improve asynchronous distributed key generation protocol of Abraham et al. [4] and VABA protocols [59], [45], while CC-ABA with perfect common coins can be used to improve various asynchronous BFT protocols using ABA such as PACE and Dumbo, and the recent asynchronous distributed key generation protocol (where the key is a field element) requiring the good-case-coin-free property [37].

ABA (weak common coins)	steps/round	rounds
MMR15 [64, 2nd alg]	9 to 13	$d + 1$
Crain [33, 1st alg]	5 to 7	$d + 1$
CC-ABA (this work)	4 or 5	$d + 1$

Table IV: ABA protocols using weak common coins. Steps/round denotes number of steps per round. Rounds denote the expected number of rounds. The total number of steps is a product of steps/round and rounds. By weak common coins, we mean all correct replicas output 0 with probability $1/d$ and output 1 with probability $1/d$ where d is a constant and $d \geq 2$.

Non-goals. We emphasize that our goal is not to build the most efficient asynchronous BFT; instead, we, for the first time, demonstrate IT and quantum secure protocols can offer comparable performance to practical BFT asynchronous protocols.

II. RELATED WORK

Asynchronous vs. partially synchronous BFT. Partially synchronous BFT protocols never violate safety, but they achieve liveness only when the network becomes synchronous [42]. As shown in [9], even for partially synchronous BFT protocols focusing on robustness [7], [28], their performance may reduce 78%-99% in failures or attack scenarios. Additionally, partially

ABA (common coins)	steps/round	rounds	good-case-coin-free
MMR14 [63]*	2 or 3	4	no
MMR15 [64, 2nd alg]	9 to 13	3	yes
Cobalt [60]	3 or 4	4	no
Crain [33, 1st alg]	5 to 7	3	yes
Crain [33, 2nd alg]	2 or 3 [†]	4	no
Pillar [72]	2 or 3	4	no
CC-ABA (this work)	4 or 5	3	yes

Table V: ABA protocols using perfect common coins. *The protocol suffers from a liveness issue. [†]The second algorithm of Crain relies high threshold common coins and is less efficient than Pillar. Compared to Pillar, CC-ABA has the good-case-coin-free property that is vital for the state-of-the-art asynchronous distributed key generation protocol [37].

synchronous protocols may experience zero throughput with a network scheduler [61]. Asynchronous BFT protocols, in contrast, do not rely on any timing assumptions and are intrinsically robust against various performance attacks.

IT BFT in partially synchronous environments. There exist several partially synchronous BFT protocols that are IT secure or can be made IT secure. In particular, PBFT (the journal version) [26], PBFT (described in Castro’s PhD thesis) [25], and Cachin’s formulation for PBFT [16] assume authenticated channels but use cryptographic hash functions. These protocols are quantum resistant but not IT secure. They can, however, be modified to achieve IT security if removing the usage of the hash functions. Recently, Stern and Abraham proposed IT HotStuff, an IT secure, partially synchronous BFT protocol that uses $O(1)$ persistent storage and $O(n^2)$ messages, where each message contains a constant number of words [70].

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 [62] contains an adaptively secure asynchronous BFT (atomic broadcast) protocol, but as it relies on inefficient local coin based ABA, it is less efficient than other protocols in WAN or large-size networks. DAG-Rider [52] achieves adaptive security if using adaptively secure common coin protocols.

The situation for asynchronous environments is in sharp contrast to that of partially synchronous BFT protocols, most of which attain adaptive security [26], [39], [69], [47], [51], [7], [28]. WaterBear achieves adaptive security and IT security and significantly outperforms EPIC that is adaptively secure but not IT secure.

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) [52]. DAG-Rider [52] achieves quantum safety, even if when being instantiated using a cryptographic common coin protocol (e.g., [18], [55]). The BKR protocol and their descendants (e.g., HoneyBadgerBFT [61], MiB [57], PACE [72]) achieve quantum safety if using techniques from EPIC [56]. All the above-mentioned protocols, however, do not achieve quantum liveness. Tusk [35] and Bullshark [46] are variants of DAG-Rider; they extensively use signatures and hashes, and achieve neither quantum safety nor quantum liveness.

(IT) Byzantine agreement. Byzantine agreement (BA) is a central tool for both fault-tolerant distributed computing and cryptography. The condition $n \geq 3f + 1$ is both necessary and sufficient for both synchronous and asynchronous BA protocols [67]. The celebrated impossibility result of Fischer, Lynch, and Paterson [44] implies that a randomized BA protocol must have non-terminating executions. A BA protocol may be $(1-\epsilon)$ -terminating, where correct replicas terminate the protocol with an overwhelming probability, or almost-surely terminating, where replicas terminate with probability one.

For our purpose, we focus on ABA protocols in the IT setting with a computationally unbounded adversary. For almost-surely ABA, Ben-Or’s ABA requires $n \geq 5f + 1$ [11], while Bracha’s ABA [14] 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 $n \geq 4f + 1$ [43]. Abraham, Dolev, and Halpern [3] provide the first almost-surely ABA with polynomial efficiency (concretely, $O(n^2)$ expected running time) and optimal resilience. Bangalore, Choudhury, and Patra [10] improve the expected running time of [3] by a factor of n .

For $(1 - \epsilon)$ -terminating ABA, Canetti and Rabin [24] build an expected constant-round ABA protocol with optimal resilience. Patra, Choudhury, and Rangan [66] build a more efficient construction in terms of communication complexity.

Both almost-surely ABA and $(1 - \epsilon)$ -terminating ABA follow the classic framework of Feldman and Micali [43] 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 [68].) Unfortunately, the framework of using AVSS for common coins is prohibitively expensive, as we have argued in the introduction. Patra, Choudhury, and Rangan [66] also propose an approach for sharing multiple secrets simultaneously. While such an approach is useful to build more efficient multi-valued BA (MBA), it is unknown if it would yield more efficient ABA protocols. While, for instance, the CNV asynchronous BFT framework [32] does use MBA, it may run $O(n)$ consecutive MBA instances (which is inefficient).

IT and universally composable MPC. As one of the most significant results in the area of cryptography and distributed computing, Ben-Or, Goldwasser, and Wigderson [12] and Chaum, Crépeau, and Damgård [27] provide generic feasibility results for perfect (IT and error-free) MPC with adaptive security. Kushilevitz, Lindell, and Rabin establish a framework for the security of protocols in the IT setting under concurrent composition [53]. Cohen, Coretti, Garay, and Zikas [29] provide a theoretic foundation for BA and MPC with probabilistic termination in the UC framework [22]. Asynchronous protocols in the UC framework are not guaranteed to (eventually) terminate because the UC adversary can delay the computation indefinitely. In light of this, Coretti, Garay, Hirt, and Zikas generalize the UC framework to the asynchronous fault-tolerant computing setting [31].

III. SYSTEM MODEL AND DEFINITIONS

A. System and Threat Model

This section describes the system model for distributed computing protocols, where f out of n replicas may fail

arbitrarily (Byzantine failures). Unless specified otherwise, the protocols we consider 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 [42].
- **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.
- **Unbounded adversary:** Depending on the capacities of the adversary, a protocol may achieve *computational security*, where the adversary is bounded and restricted to probabilistic polynomial-time (PPT), or achieve *information-theoretic (IT) security*, where the adversary is unbounded. We have argued IT security is preferable compared to computational security, but constructing IT secure protocols is more challenging.
- **Adaptive corruptions:** Depending on how the adversary decides to corrupt parties, there are two types of corruptions: static corruptions and adaptive corruptions. In the static corruption model, the adversary is restricted to choose its set of corrupted replicas at the start of the protocol and cannot change this set later on. An adaptive adversary can choose its set of corrupted replicas at any moment during the execution of the protocol, based on the information it accumulated thus far (i.e., the messages observed and the states of previously corrupted replicas). There is a strong separation result that statically secure protocols are not necessarily adaptively secure [23], [34].

All known asynchronous BFT protocols *implemented* assume trusted setup, achieving computational security and static security only. None of them achieve quantum security either.

For our protocols, we may associate each protocol instance with a unique identifier id , tagging each messages in the instance with id . If no ambiguity arises, we may simply omit the identifiers.

B. Definitions and Background

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 .

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 [72]. 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 toward 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 $f + 1$ replicas, not simply all correct replicas, propose 1, then a correct replica that terminates *decides* 1. The property guarantees the PACE framework to have sufficient transactions delivered.

RBC. A Byzantine reliable broadcast (RBC) protocol [36], [6], [15], [20] is specified by *r-broadcast* and *r-deliver* such that the following properties hold:

- **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.

Bracha’s broadcast [14] has a bandwidth of $\mathcal{O}(n^2|m|)$ and is IT secure, and CT RBC due to Cachin and Tessaro [20]

uses hash functions (with output length λ) to reduce the bandwidth to $\mathcal{O}(n|m| + \lambda n^2 \log n)$. Recent works proposed various IT RBC and RBC protocols using hash functions (quantum secure) with lower communication [36], [5].

PACE framework. PACE uses RBC and RABA in a black-box manner to construct efficient asynchronous BFT. The framework allows all ABA instances to run in parallel, removing a well-known bottleneck in the original framework of Ben-Or, Kemler, and Rabin [13]. PACE also provides a fast path for consensus, allowing the protocol to terminate using a single RABA round. It is an interesting research problem to improve RABA (and the underlying ABA) protocols in terms of message and communication complexity.

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* [24], 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 terminates in expected exponential rounds, while ABA assuming common coins (including CC-ABA 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.

IV. ABA FROM LOCAL COINS AND COMMON COINS

The state-of-the-art local coin based ABA protocol, Bracha’s ABA [14], has $\mathcal{O}(n^3)$ 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 then present Quadratic-ABA that reduces the messages from $\mathcal{O}(n^3)$ to $\mathcal{O}(n^2)$ and reduces the number of steps to 5 in each round. *The improvement is significant, allowing WaterBear to attain the same average message complexity as PACE— $\mathcal{O}(n^3)$.* Both Cubic-ABA and Quadratic-ABA can be readily modified to be efficient RABA protocols.

As an important by-product, extending the idea of Quadratic-ABA and assuming the existence of (weak or perfect) common coins, we are able to present ABA protocols that have an expected constant rounds and outperform the state-of-the-art ABA protocols, as shown in Table IV and Table V in Sec. I.

A. Cubic-ABA

Figure 1 describes the pseudocode of Cubic-ABA and Figure 2 illustrates the workflow. Cubic-ABA uses the *broadcast* primitive of best-effort broadcast and the *r-broadcast* and *r-deliver* primitives of RBC. The protocol proceeds in rounds, beginning with round 0. Each round consists of three phases.

```

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 $_r(iv_r)$  {▷ 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      $bset_r \leftarrow bset_r \cup \{v\}$ 
12   wait until  $bset_r \neq \emptyset$ 
13     if main-vote $_r()$  has not been sent
14       broadcast main-vote $_r(v)$  where  $v \in bset_r$ 
15   upon receiving  $n - f$  main-vote $_r()$  such that for each received
    main-vote $_r(b)$ ,  $b \in bset_r$  {▷ 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 $_r(v)$ ,  $v \in bset_r$ ; for each final-vote $_r(*)$ ,  $bset_r = \{0, 1\}$ 
20     if there there are  $n - f$  final-vote $_r(v)$ 
21        $iv_{r+1} \leftarrow v$ , decide  $v$ 
22     else if there are  $f + 1$  final-vote $_r(v)$ 
23        $iv_{r+1} \leftarrow v$ 
24     else
25        $c \leftarrow \text{Random}()$  {obtain local coin}
26        $iv_{r+1} \leftarrow c$ 
27    $r \leftarrow r + 1$ 

```

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




Figure 2: The workflow of Cubic-ABA.

In the first phase, a replica p_i broadcasts a pre-vote $_r(iv_r)$ message, where $iv_r \in \{0, 1\}$ is the input value of p_i for round r (ln 07). At ln 08-09, if p_i receives $f + 1$ pre-vote $_r(v)$ for some $v \in \{0, 1\}$ and has not previously broadcast pre-vote $_r(v)$, it also broadcasts pre-vote $_r(v)$.

At ln 10-14, p_i enters the second phase. If p_i receives $2f + 1$ 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 main-vote $_r(v)$ message (ln 12-14).

In the third phase, a correct replica p_i accepts a main-vote $_r(v)$ message only if v has already been added locally to $bset_r$ (ln 15). If p_i has received $n - f$ main-vote $_r(v)$, p_i r -broadcasts a final-vote $_r(v)$ message (ln 16-17). Otherwise, 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 24-26: Otherwise, a replica generates a local random coin and uses it as input for the next round.

Intuition and discussion. Our motivation for Cubic-ABA is to reduce the number of parallel RBCs in Bracha’s ABA. We recall Bracha’s ABA in Appendix A. In each round, Bracha’s ABA has three phases, where in each phase, replicas run n parallel RBCs, with 12 steps and $O(n^3)$ messages.

In our approach, the first two phases of Cubic-ABA resemble those of common coin based ABA protocols [72], [63], [33], [60], [64], 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 $n - f$ votes (i.e., final-vote $_r(v)$) for the same v in the same phase, any correct replica will either decide v or vote for v in the following round. 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 a result, the number of n parallel RBC instances is 1 instead of 3, and the number of steps is reduced from 12 to 7.

B. Quadratic-ABA

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.

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 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-

```

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 $_r(iv_r)$  {▷ 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      $bset_r \leftarrow bset_r \cup \{v\}$ 
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 $_r(v)$ ,
 $v \in bset_r$  {▷ 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 $_r(*)$ ,  $bset_r = \{0, 1\}$  {▷ phase 4}
20     if there 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 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        $c \leftarrow Random()$  {obtain local coin}
30        $iv_{r+1} \leftarrow c$ 
31    $r \leftarrow r + 1$ 

```

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




Figure 4: The workflow of Quadratic-ABA.

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

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

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

- Ln 24-25: If p_i receives $n - f$ $\text{final-vote}_r(v)$, 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 $n - f$ $\text{final-vote}_r()$ messages that carry only value v and $*$, it uses v for round $r + 1$.
- Ln 28-30: Otherwise, p_i generates a local random coin and uses it as input for the next round.

Intuition and discussion. The $\text{pre-vote}_r()$ phase and $\text{vote}_r()$ phase are similar to the $\text{pre-vote}_r()$ phase and the $\text{main-vote}_r()$ phase in Cubic-ABA. In Quadratic-ABA, we use the $\text{main-vote}_r()$ phase and the $\text{final-vote}_r()$ phase to replace the parallel RBC phase of Cubic-ABA. Accordingly, Quadratic-ABA achieves $O(n^2)$ messages and $O(n^2)$ communication.

Our goal is to guarantee that if a correct replica receives $n - f$ $\text{final-vote}_r(v)$, any correct replica will set iv_{r+1} as v . First, if a correct replica sends $\text{main-vote}_r(v)$, no correct replica will send $\text{main-vote}_r(\bar{v})$ (which we prove in Lemma 13 in Appendix D). In particular, if a correct replica sends $\text{main-vote}_r(v)$, it must have received $n - f$ $\text{vote}_r(v)$. If another correct replica sends $\text{main-vote}_r(\bar{v})$, at least $n - f$ replicas have sent $\text{vote}_r(v)$. Therefore, at least one correct replica has sent both $\text{vote}_r(v)$ and $\text{vote}_r(\bar{v})$, contradicting the fact that each correct replica only sends a single $\text{vote}_r()$ message in each round. Furthermore, if a correct replica sends $\text{final-vote}_r(v)$, no correct replica will send $\text{final-vote}_r(\bar{v})$ or even accept $\text{final-vote}_r(\bar{v})$ from other replicas (Lemma 14 and Lemma 19). Briefly speaking, this is because if a correct replica accepts $\text{final-vote}_r(\bar{v})$, at least one correct replica has sent $\text{main-vote}_r(\bar{v})$. Meanwhile, if a correct replica accepts $\text{final-vote}_r(v)$, at least one correct replica has sent $\text{main-vote}_r(v)$, contradicting Lemma 13. Hence, at the end of each round, if a correct replica receives only $\text{final-vote}_r(v_1)$ and $\text{final-vote}_r(*)$, another correct replica receives only $\text{final-vote}_r(v_2)$ and $\text{final-vote}_r(*)$, 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 $n - f$ $\text{final-vote}_r(v)$. Among them, at least $f + 1$ correct replicas have sent $\text{final-vote}_r(v)$. With $3f + 1$ replicas in total, there are at most $2f$ $\text{final-vote}_r(\bar{v})$ or $\text{final-vote}_r(*)$. Hence, every correct replica receives at least one $\text{final-vote}_r(v)$. As no correct replica will accept $\text{final-vote}_r(\bar{v})$, every correct replica will only have $\text{final-vote}_r(v)$ and $\text{final-vote}_r(*)$. 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.

C. CC-ABA

Both Cubic-ABA and Quadratic-ABA can be transformed to ABA from weak common coins [24], [65] and perfect common coins. Here by weak common coins, we mean that all correct replicas output 0 with probability $1/d$ and output 1 with probability $1/d$ where d is a constant and $d \geq 2$, and the probability that correct replicas obtain different values is

$(d - 2)/d$. By perfect common coins, we mean that all correct replicas always output the same random coin. Note perfect coins are a special case of weak coins (by setting $d = 2$).

As Quadratic-ABA is more efficient, we here focus on Quadratic-ABA. Our main result is that by replacing local coins of Quadratic-ABA with weak (or perfect) common coins, we immediately obtain CC-ABA terminating in $O(1)$ time. CC-ABA reduces the expected number of steps of prior constructions, as shown in Table IV and Table V. Note that ABA is the major bottleneck in asynchronous BFT protocols as reported in [41], [49], [48]. The improvement is significant and has practical implications, as the recent work of PACE has shown that even a single step improvement can lead to a drastic performance improvement (for instance, easily 2x) in BFT protocols [72]. We prove the correctness of CC-ABA in Appendix E.

V. RABA FROM LOCAL COINS

As shown in PACE [72], 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 Appendix B.

A. The Subtlety of Building Quadratic-RABA

PACE introduced a general approach to converting ABA to RABA [72]. Following their approach, we present Quadratic-RABA in Figure 5. Quadratic-RABA is identical to Quadratic-ABA except the first round (round 0), where we make the following changes. First, we use a $\text{propose}()$ event and a $\text{repropose}()$ event (Ln 03-07). Upon $\text{propose}(v)$, a replica p_i starts round 0 and executes the $\text{broadcast-vote}(v)$ function. Upon $\text{repropose}(1)$ event, regardless of which round a replica is in, p_i still executes the $\text{broadcast-vote}(v)$ function. The $\text{propose}()$ and $\text{repropose}()$ 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 $\text{broadcast-vote}(v)$ function, replica p_i broadcasts a $\text{pre-vote}_0(v)$ message (Ln 09). At Ln 10-14, if $v = 1$, p_i immediately adds 1 to $bset_0$, and broadcasts $\text{vote}_0(1)$, $\text{main-vote}_0(1)$, and $\text{final-vote}_0(1)$.

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* of terminating in only one step. Namely, if $f + 1$ correct replicas propose 1, no correct replica will receive $2f + 1$ $\text{final-vote}_0(0)$. 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 $\text{vote}_0(1)$, $\text{main-vote}_0(1)$, $\text{final-vote}_0(1)$, 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 (as shown in [72]). Surprisingly and unexpectedly, we find that for Quadratic-ABA, however, there is still a subtle liveness issue for round 0. We illustrate the issue via an example. Suppose f correct replicas propose 1 and $f + 1$ correct replicas propose 0. The f replicas directly broadcast $\text{vote}_0(1)$, $\text{main-vote}_0(1)$,


```

01 initialization
02    $r \leftarrow 0$  {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-vote0( $v$ ) has not been sent, broadcast pre-vote0( $v$ )
10   if  $v = 1$ 
11      $bset_0 \leftarrow bset_0 \cup \{1\}$ 
12     if vote0() has not been sent, broadcast vote0(1)
13     if main-vote0() has not been sent, broadcast main-vote0(1)
14     if final-vote0() has not been sent, broadcast final-vote0(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      $bset_r \leftarrow bset_r \cup \{v\}$ 
21     wait until  $bset_r \neq \emptyset$ 
22     if vote $r$ () has not been sent
23       broadcast vote $r$ ( $v$ ) where  $v \in bset_r$ 
24     upon receiving  $n - f$  vote $r$ () such that for each received
    vote $r$ ( $b$ ),  $b \in bset_r$ 
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  $r = 0$ ,  $v \in bset_r$ , 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  $r = 0$ ,  $v \in bset_r$ , 2) 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        $iv_{r+1} \leftarrow v$ , decide  $v$ 
35     else if there are only final-vote $r$ ( $v$ ) and final-vote $r$ (*)
36        $iv_{r+1} \leftarrow v$ 
37     else
38       if  $r = 0$ ,  $c \leftarrow 1$  {coin in the first round is 1}
39       else  $c \leftarrow \text{Random}()$  {obtain local coin}
40        $iv_{r+1} \leftarrow c$ 
41    $r \leftarrow r + 1$ 

```

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

and final-vote₀(1). Even if the $f + 1$ correct replicas that proposed 0 may later repropose 1, they may have already sent vote₀(0), main-vote₀(0), and final-vote₀(0). In this case, no correct replica will accept final-vote₀(1) as they do not receive $f + 1$ main-vote₀(1). In summary, 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), and each correct replica accepts a final-vote _{r} (v) message only if it has previously received $f + 1$ main-vote _{r} (v).

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) (In 28) and final-vote _{r} (v) (In 32), a correct replica accepts it as long as $v \in bset_0$. With this modification, the set of $f + 1$ correct replicas that proposed 0 will repropose 1, so every correct replica will eventually put

```

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  $j \in S$ 
14      $a$ -deliver( $\cup_{j \in S} \{m_j\}$ )

```

Figure 6: The WaterBear family (WaterBear-C and WaterBear-Q). The code for replica p_i in epoch e . WaterBear uses Bracha’s broadcast as the underlying the RBC. WaterBear-C uses Cubic-RABA as the underlying RABA and WaterBear-Q uses Quadratic-RABA as the underlying RABA. WaterBear uses the technique of EPIC to select transactions.

1 in pre-vote₀(0). Hence, every correct replica will eventually accept main-vote₀(1) and final-vote₀(1).

Our result underlines the subtlety of constructing RABA from ABA and the importance of a full proof for a new protocol. We prove the correctness of Quadratic-RABA in Appendix G.

VI. 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 are additionally information-theoretically secure.

A. The WaterBear Protocols

WaterBear follows the PACE paradigm but uses the trick in EPIC [56] 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, and *propose*, *repropose* and *decide* primitives of WaterBear RABA. Figure 6 depicts the pseudocode of WaterBear. In terms of transaction selection strategy, we follow EPIC and ask replicas to select random transactions *in plaintext* for most epochs and periodically switch to the FIFO selection, where replicas maintain a log of transactions according to the order transactions are received and replicas select the first group of transactions in the buffer as input. As shown in EPIC, the approach shares similar performance as the random selection approach used in HoneyBadgerBFT and BEAT. 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 repropose 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

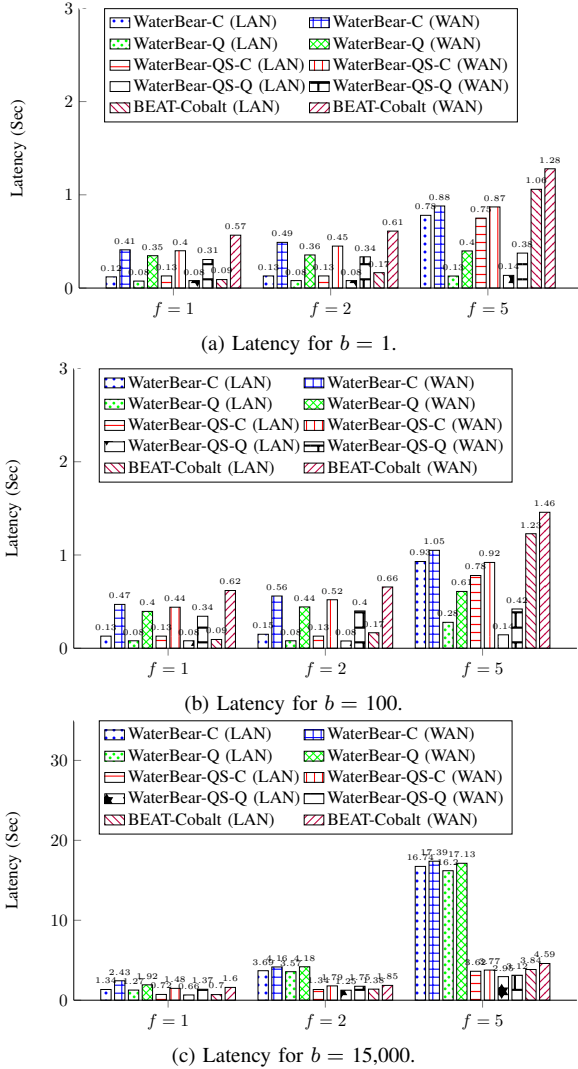


Figure 7: Latency of the protocols.

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.

B. 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. WaterBear-QS does not achieve information-theoretic security but achieves quantum security for both safety and liveness properties. Prior to our work, no such BFT protocol has been implemented. The difference between WaterBear and WaterBear-QS is that WaterBear-QS leverages CT RBC [20] to reduce the communication complexity of RBC. Jumping ahead, we show the modification leads to a dramatic performance improvement compared to WaterBear protocols.

VII. IMPLEMENTATION AND EVALUATION

Implementation. We implemented WaterBear-C, WaterBear-Q, WaterBear-QS-C, and WaterBear-QS-Q in Golang. The library involves more than 10,000 LOC for the protocol

implementations and about 1,000 LOC for evaluation.

All the protocols use authenticated channels, and WaterBear-QS additionally uses a hash function. We use HMAC to realize the authenticated channel. HMAC is quantum secure but not IT secure; we used HMAC in our reference implementation for WaterBear to *demonstrate the overhead of WaterBear itself*, because all other protocols introduced in this paper use HMAC for authentication. We use SHA256 as the hash function. We use gRPC as the underlying communication library.

All the protocols directly use RBC in their RBC phases. WaterBear-C and WaterBear-QS-C additionally use RBC in the ABA phase. For WaterBear-C and WaterBear-Q, we use Bracha’s broadcast (which is IT secure) in the RBC phase. For WaterBear-QS-C and WaterBear-QS-Q, we use CT RBC [20] (using erasure coding and hash functions) in the RBC phase. In ABA (the third phase) of WaterBear-C and WaterBear-QS-C, we directly use Bracha’s broadcast, because in the ABA phase, there is no bulk data (and no need to use erasure coding). To implement CT RBC, we use a Golang Reed-Solomon code library [2].

For comparison, we choose to implement BEAT-Cobalt in our Golang library. BEAT [1], [41] was originally implemented in Python 2.7 using MMR ABA [63]. We implement BEAT-Cobalt, replacing MMR with Cobalt-ABA, as Cobalt ABA addressed the liveness issue of MMR. There are several reasons we chose BEAT-Cobalt as the baseline protocol. First, BEAT-Cobalt is one of the most efficient open-source asynchronous BFT implementations. As shown in PACE [72], BEAT-Cobalt is more efficient than Dumbo [49] for $n \leq 46$. Second, all WaterBear protocols achieve adaptive security and EPIC is the only known adaptively secure asynchronous BFT protocol implemented. EPIC has shown that BEAT-Cobalt significantly outperforms EPIC in both LAN and WAN settings. Hence, as long as we demonstrate the performance difference among BEAT-Cobalt and our protocols, we can reasonably argue which is the most efficient adaptively secure asynchronous BFT protocol among EPIC, WaterBear, and WaterBear-QS-C. Of course, EPIC neither achieves quantum security nor IT security, and EPIC requires trusted setup.

Here are some other reasons why we choose not to (and do not need to) compare with Dumbo and Tusk. Neither protocols achieve any properties our protocols can achieve: no adaptive security, relying on PKC, no quantum resistance, and assuming trusted setup. *Indeed, our goal is not to claim the performance benefit of WaterBear protocols, but we aim at refuting the conventional wisdom that quantum secure and IT secure asynchronous BFT are not (yet) practical.*

Overview of evaluation. We evaluate the performance of our protocols on Amazon EC2 utilizing up to 61 virtual machines (VMs) from different regions in five continents. 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 deploy our protocols in both “LAN” and “WAN” settings. In the LAN setting, the replicas are run in the same region of EC2 (e.g., US Virginia), but these replicas may be located in different physical datacenters. In the WAN setting, the replicas are evenly distributed across different continents.

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 report the average performance (for both throughput and latency). We use two different transaction sizes. We evaluate the performance of the protocols for transactions with 100 bytes by default and evaluate that using 250 bytes per transaction.

We assess the performance of 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:

- WaterBear-QS-C and WaterBear-QS-Q are about as efficient as BEAT-Cobalt. The three protocols, however, offer interesting trade-offs for different n 's. When $n \geq 16$, WaterBear-QS-Q is strictly better than WaterBear-QS-C and BEAT-Cobalt (higher throughput and lower latency).
- We confirm that the bandwidth of RBC is one of the major bottlenecks for the BKR framework and PACE framework. WaterBear-QS-C, WaterBear-QS-Q, and BEAT-Cobalt (using bandwidth-efficient CT RBC) are about twice as efficient as WaterBear (using bandwidth-expensive Bracha's broadcast). Moreover, for all the protocols we evaluate, the throughput with a transaction size of 100 bytes is more than 2x the throughput with a transaction size of 250 bytes. To put it differently, WaterBear-C and WaterBear-Q can be easily made more efficient if using a more efficient IT RBC.
- All the four protocols we propose are highly robust against various crash and Byzantine failures, just as BEAT-Cobalt.

A. Performance in Failure-Free Cases

Latency. We report the latency of the protocols in both LAN and WAN settings for $f = 1, 2$, and 5 in Figure 7 for $b = 1, 100$, and 15,000. For $b = 1$ and $b = 500$, WaterBear-C, WaterBear-Q, WaterBear-QS-C, and WaterBear-QS-Q achieve lower latency than BEAT-Cobalt in both LAN and WAN settings. Among the protocols, WaterBear-QS-Q consistently achieves the lowest latency, mainly due to the lower communication complexity in the RBC phase and lower message complexity in the ABA phase.

For $b = 15,000$, WaterBear-QS-C, WaterBear-QS-Q, and BEAT-Cobalt share similar performance. When $f = 1$, the latency for WaterBear-QS-C, WaterBear-QS-Q, and BEAT-Cobalt are about half of that for WaterBear-C and WaterBear-Q in both LAN and WAN settings. As f increases, the difference in latency among these protocols increases. When $f = 5$, the latency of WaterBear-C and WaterBear-Q is around four times as much as that for WaterBear-QS-C, WaterBear-QS-Q and BEAT-Cobalt. This is expected, since as f increases, the network bandwidth consumption for WaterBear is significantly higher than the other protocols.

WaterBear-QS-Q has consistently lower latency than WaterBear-QS-C, and WaterBear-QS-C has lower latency than BEAT-Cobalt in both LAN and WAN settings. When $f = 5$ and $b = 15000$ in the WAN environment, the latency of WaterBear-QS-C and WaterBear-QS-Q are 17.9% and 32.0% lower than BEAT-Cobalt, respectively. The better performance

for WaterBear-QS is that WaterBear-QS utilizes a biased ABA, causing the protocol to terminate faster. Furthermore, WaterBear-QS-Q outperforms WaterBear-QS-C, as Quadratic-ABA has lower message complexity than Cubic-ABA.

It is not surprising that all protocols achieve higher latency in WANs than in LANs. We find the difference is more visible when f is smaller and when b is smaller. For instance, for $b = 1$, the latency of WaterBear-C is 241.7% higher for $f = 1$ in WANs compared to that in LAN but is only 12.8% higher for $f = 5$. For $b = 15,000$, the latency of WaterBear-C is 81.3% higher for $f = 1$ in WANs compared to that in LANs but is only 3.9% higher for $f = 5$.

Throughput and scalability. We evaluate the throughput and scalability for WaterBear-C, WaterBear-Q, WaterBear-QS-C, WaterBear-QS-Q and BEAT-Cobalt by varying the network size f from 1 to 20. Unless otherwise specified, all experiments are conducted in the WAN setting running on *m5.xlarge* type instances. We also report throughput vs. latency in Figure 8.

First of all, similar to the results for latency, 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.

We assess the throughput of all five protocols for $f = 1$ as depicted in Figure 8b: the performance of WaterBear-QS-C, WaterBear-QS-Q, and BEAT-Cobalt in WANs are very close, though the peak throughput of WaterBear-QS-Q is slightly higher in most experiments. We also conduct a separate experiment in the LAN setting and evaluate the throughput, as shown in Figure 8a. Unlike the results in WANs, the throughput of BEAT-Cobalt in LANs is marginally higher than WaterBear-QS-C and the throughput of WaterBear-QS-Q is marginally higher than BEAT-Cobalt: the peak throughput of BEAT-Cobalt is 2.3% higher than WaterBear-QS-C and the peak throughput of WaterBear-QS-Q is 3.9% higher than BEAT-Cobalt. 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.

When f increases, the performance trend is slightly different from the case for $f = 1$. In particular, when $f = 5$ and $f = 10$, the throughput of WaterBear-QS-C is consistently higher than BEAT-Cobalt. When $f = 20$, the peak throughput of the two protocols is again very close, the latency of WaterBear-QS-C, however, is consistently higher than BEAT-Cobalt. The performance of WaterBear-QS-Q is consistently higher than the other protocols (higher throughput and lower latency). The peak throughput of WaterBear-QS-Q is 23.0%, 47.4%, and 27.4% higher than BEAT-Cobalt for $f = 5, 10$, and 20, respectively. Meanwhile, the peak throughput of WaterBear-QS-C is 6.0% and 17.4% higher than BEAT-Cobalt for $f = 5$ and 10, respectively. When $f = 20$, the peak throughput of WaterBear-QS-C is only slightly lower than BEAT-Cobalt. The evaluation results are mainly due to two factors: 1) the PACE paradigm employed by WaterBear-QS-Q and WaterBear-QS-C is more efficient and the BKR framework employed by BEAT; 2) The Quadratic-ABA protocol used in WaterBear-QS-Q achieves lower communication complexity than Cubic-ABA (used in WaterBear-QS-C) and fewer expected number of steps than Cobalt used in BEAT.

Performance on different types of VMs. Different from prior

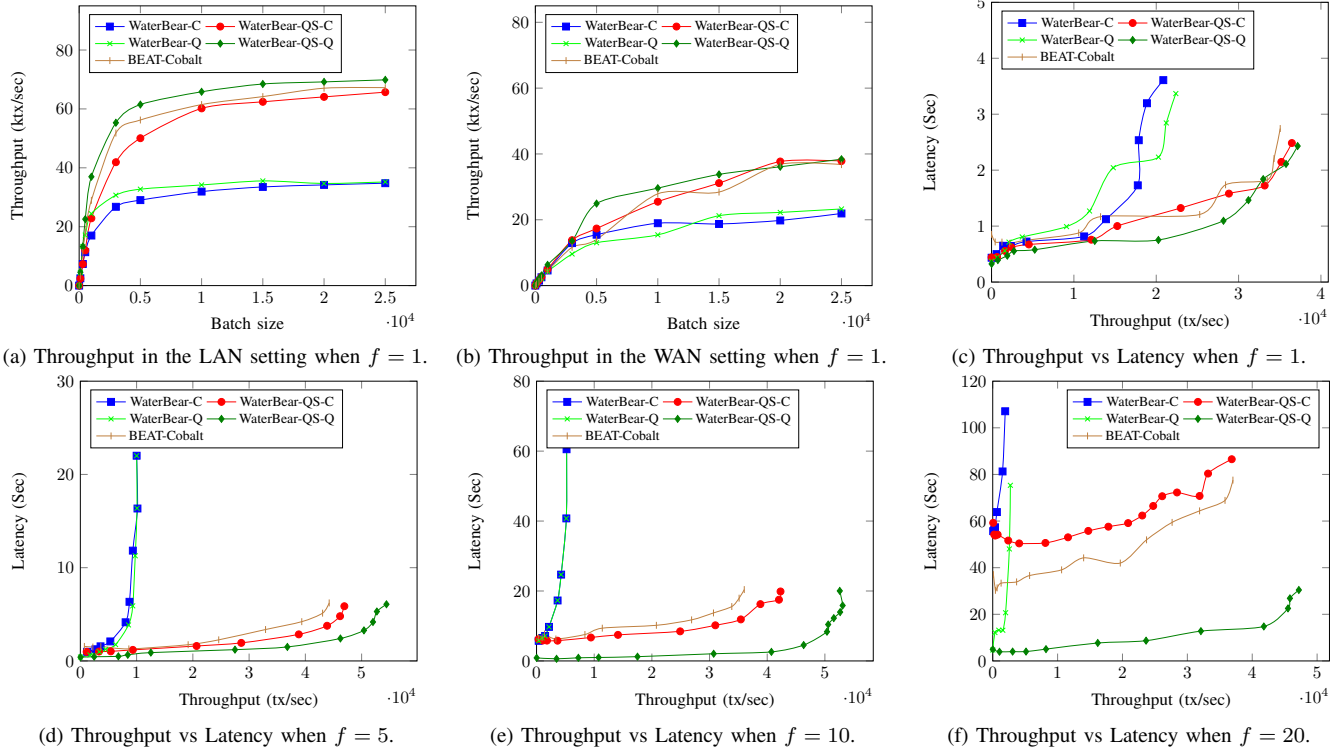
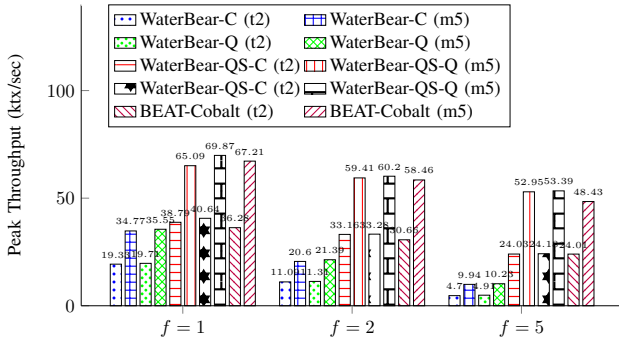
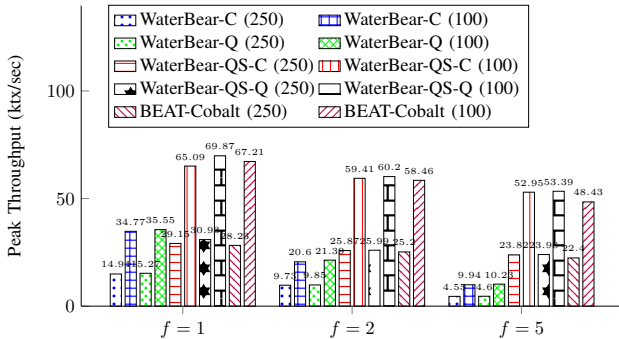


Figure 8: Throughput vs latency on *m5.xlarge* instances for $f = 1$ to $f = 20$.



(a) Peak throughput of protocols running on different EC2 instances.



(b) Peak throughput for transaction size of 100 bytes and 250 bytes.

Figure 9: Performance of the protocols for $f = 1, 2,$ and 5 .

protocols (HoneyBadgerBFT, BEAT, Dumbo, EPIC) that all evaluate the performance on *t2.medium* instances, we evaluate the performance of the protocols using both *t2.medium* (t2 in the figures) and *m5.xlarge* (m5 in the figures) instances. In particular, we evaluate the throughput with $b = 15,000$ for $f = 1, f = 2,$ and $f = 5,$ the results of which are shown


in Figure 9a. For all the protocols, the peak throughput on *m5.xlarge* instances is about $2\times$ that on *t2.medium*.

Performance with different transaction sizes. We also report the throughput of the protocols by fixing b to 15,000 but using different sizes of transactions (100 bytes and 250 bytes), the results of which are shown in Figure 9b. For all the five protocols, the performance using transaction size of 100 bytes is consistently higher, being at least twice as efficient as that with 250 bytes. The finding highlights the main bottleneck for the protocols for large transaction sizes is RBC.

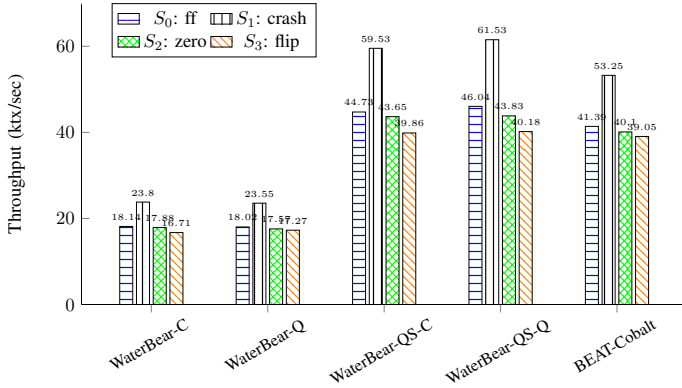
B. Performance under Failures

To assess the protocol performance under failures and attacks, we carefully design various experiments as follows.

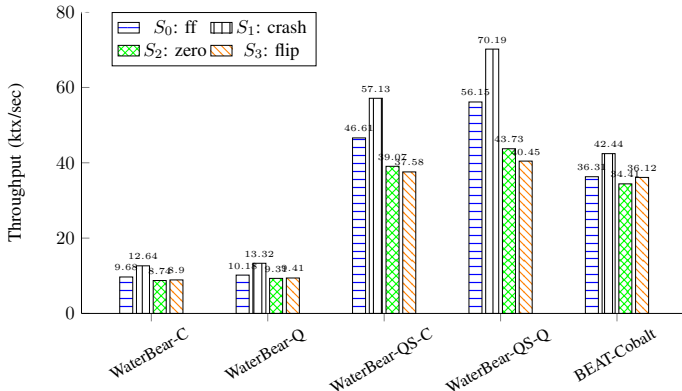
- S_0 : **(failure-free)** In this scenario, all replicas are correct. S_0 is the baseline scenario we use to compare with failure scenarios.
- S_1 : **(crash)** In this scenario, we let f replicas crash by not participating in the protocols.
- S_2 : **(Byzantine; keep voting 0)** In this scenario, we control all f faulty replicas to keep voting for 0 in each step of (R)ABA. For all the five protocols, doing so would intuitively make fewer ABA and RABA instances decide 1 and would likely decrease the throughput of the protocols. We aim to observe the throughput reduction in this scenario compared to the failure-free scenario.
- S_3 : **(Byzantine; flipping the (R)ABA input)** In this scenario, 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 three protocols. For WaterBear and WaterBear-QS, the strategy would, at



(a) Throughput under failure scenarios in WAN for $f = 1$ and $b = 15,000$.



(b) Throughput under failure scenarios in WAN for $f = 2$ and $b = 15,000$.



(c) Throughput under failure scenarios in WAN for $f = 5$ and $b = 15,000$.

Figure 10: Performance of the protocols in failure scenarios.

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 to decide in round 0 and force them to enter the second round of RABA, where the two protocols start to query the local coins. Jumping ahead, our experiment actually shows that the flipping strategy S_3 works slightly better than S_2 , but in general is not destructive.

Note that we do not attempt to attack the RBC phase for all these protocols, because RBC is highly robust during failures and attacks (see, e.g., [30]).

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

Performance under crash failures (S_1). When $f = 1$, WaterBear family achieves higher throughput under crash failures compared to that in the failure-free case. For WaterBear-QS-C, WaterBear-QS-Q, and BEAT-Cobalt, the performance is slightly lower under crash failures than those in the failure-free scenario. The throughput of WaterBear-QS-C, WaterBear-QS-Q, and BEAT-Cobalt are 3.0% lower, 6.7% lower, and 8.2% lower in the crash failure scenario than in the failure-free scenario. For all other cases (other f), all five protocols achieve higher performance under crash failures compared to that in the failure-free case. For instance, when $f = 5$, the throughput of WaterBear-QS-C is 22.6% higher, the throughput of WaterBear-QS-Q is 25.0% higher, and the throughput of BEAT-Cobalt is 16.9% higher in the crash failure scenario. This is mainly because under crash failures, the network bandwidth consumption is much lower (about 33% lower) than in the failure-free case. When $f = 1$, as the network bandwidth consumption does not dominate the overhead, the performance of WaterBear-QS-C, WaterBear-QS-Q, and BEAT-Cobalt under crash failures is similar to that in failure-free cases. When $f = 2$ and $f = 5$, the performance improvement under crash failures for WaterBear-QS-C is higher than that of WaterBear-QS-Q and BEAT-Cobalt. We believe this is partly because WaterBear-QS-C involves multiple RBC instances (in both the RBC phase and the ABA phase) and uses more bandwidth than WaterBear-QS-Q and BEAT-Cobalt.

Performance under Byzantine failures. The performance of all the protocols is slightly lower under Byzantine failures compared to the failure-free scenario and crash failure scenario. The performance degradation of WaterBear-QS-C, WaterBear-QS-Q, and BEAT-Cobalt may appear higher than that of WaterBear-C and WaterBear-Q. This is actually because the performance of WaterBear-C and WaterBear-Q is much lower than the other two protocols. For WaterBear-QS-C, the throughput is 10.9%-19.4% lower under Byzantine failures compared to failure-free scenario. For WaterBear-QS-Q, the throughput is 12.7%-27.9% lower under Byzantine failures. In comparison, the throughput of BEAT-Cobalt is 0.5%-6.2% under Byzantine failures. The higher performance degradation for WaterBear-QS-C and WaterBear-QS-Q 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.

S_2 vs. S_3 . 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-Cobalt, 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. In contrast, for WaterBear-QS-C and WaterBear-QS-Q, the performance is consistently lower in S_3 than S_2 . This is expected: first, Cobalt ABA uses common coin and has $O(1)$ time complexity, thereby being less sensitive to the attack in S_3 . Second, as Cubic-ABA and Quadratic-ABA use local coins, forcing replicas to receive flipped values may drive replicas to receive both 0 and 1 such that the local coin value will be used; when Cubic-ABA and Quadratic-ABA use local coins, the protocols may have more rounds.

VIII. CONCLUSION

This paper designs and implements a family of asynchronous BFT protocols that are information-theoretically secure or quantum secure. Our protocols are also the first to defend against adaptive corruptions while assuming no trusted

dealers in completely asynchronous settings. Our experiments demonstrate that protocols are efficient in both failure and failure-free scenarios. In particular, one of our quantum secure protocols, WaterBear-QC-Q, consistently outperforms the state-of-the-art protocols that have much weaker security guarantees. We also contribute in different settings more efficient ABA and RABA protocols that can be used to improve various high-level Byzantine-resilient protocols.

REFERENCES

- [1] “BEAT library,” <https://github.com/fififish/beat>, 2021.
- [2] “Reed-Solomon library,” <https://github.com/klauspost/reedsolomon>, 2021.
- [3] I. Abraham, D. Dolev, and J. Y. Halpern, “An almost-surely terminating polynomial protocol for asynchronous byzantine agreement with optimal resilience,” ser. PODC, 2008.
- [4] I. Abraham, P. Jovanovic, M. Maller, S. Meiklejohn, G. Stern, and A. Tomescu, “Reaching consensus for asynchronous distributed key generation,” in *PODC*, 2021.
- [5] N. Alhaddad, S. Das, S. Duan, L. Ren, M. Varia, Z. Xiang, and H. Zhang, “Balanced byzantine reliable broadcast with near-optimal communication and improved computation,” in *PODC*, 2022.
- [6] N. Alhaddad, M. Varia, and H. Zhang, “High-threshold avss with optimal communication complexity,” in *FC*, 2021.
- [7] Y. Amir, B. Coan, J. Kirsch, and J. Lane, “Prime: Byzantine replication under attack,” *TDSC*, vol. 8, no. 4, pp. 564–577, 2011.
- [8] E. Androutaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, S. Muralidharan, C. Murthy, B. Nguyen, M. Sethi, G. Singh, K. Smith, A. Sorniotti, C. Stathakopoulou, M. Vukolić, S. W. Cocco, and J. Yellick, “Hyperledger fabric: A distributed operating system for permissioned blockchains,” 2018.
- [9] P. Aublin, S. B. Mokhtar, and V. Quéma, “Rbft: Redundant byzantine fault tolerance,” in *ICDCS*, July 2013, pp. 297–306.
- [10] L. Bangalore, A. Choudhury, and A. Patra, “The power of shunning: Efficient asynchronous byzantine agreement revisited*,” *J. ACM*, 2020.
- [11] M. Ben-Or, “Another advantage of free choice: Completely asynchronous agreement protocols (extended abstract),” in *PODC*, 1983, pp. 27–30.
- [12] M. Ben-Or, S. Goldwasser, and A. Wigderson, “Completeness theorems for non-cryptographic fault-tolerant distributed computation,” ser. STOC, 1988.
- [13] M. Ben-Or, B. Kelmer, and T. Rabin, “Asynchronous secure computations with optimal resilience,” in *PODC*. ACM, 1994, pp. 183–192.
- [14] G. Bracha, “An asynchronous $(n-1)/3$ -resilient consensus protocol,” in *PODC*. ACM, 1984, pp. 154–162.
- [15] —, “Asynchronous byzantine agreement protocols,” *Information and Computation*, vol. 75, no. 2, pp. 130–143, 1987.
- [16] C. Cachin, “Yet another visit to paxos,” ser. IBM Research Report RZ 3754, 2010.
- [17] C. Cachin, K. Kursawe, F. Petzold, and V. Shoup, “Secure and efficient asynchronous broadcast protocols,” in *CRYPTO*. Springer, 2001, pp. 524–541.
- [18] C. Cachin, K. Kursawe, and V. Shoup, “Random oracles in constantino-ple: Practical asynchronous Byzantine agreement using cryptography,” *Journal of Cryptology*, vol. 18, no. 3, pp. 219–246, 2005.
- [19] C. Cachin and J. A. Poritz, “Secure intrusion-tolerant replication on the internet,” in *DSN*. IEEE, 2002, pp. 167–176.
- [20] C. Cachin and S. Tessaro, “Asynchronous verifiable information dispersal,” in *SRDS*. IEEE, 2005, pp. 191–201.
- [21] C. Cachin and M. Vukolic, “Blockchain consensus protocols in the wild,” in *DISC*, 2017.
- [22] R. Canetti, “Universally composable security,” *J. ACM*, 2020.
- [23] R. Canetti, U. Feige, O. Goldreich, and M. Naor, “Adaptively secure multi-party computation,” in *STOC*, 1996.
- [24] R. Canetti and T. Rabin, “Fast asynchronous byzantine agreement with optimal resilience,” in *STOC*, vol. 93. Citeseer, 1993, pp. 42–51.
- [25] M. Castro, “Practical byzantine fault tolerance,” ser. PhD thesis, 2001.
- [26] M. Castro and B. Liskov, “Practical Byzantine fault tolerance and proactive recovery,” *TOCS*, vol. 20, no. 4, pp. 398–461, 2002.
- [27] D. Chaum, C. Crépeau, and I. Damgård, “Multiparty unconditionally secure protocols,” in *STOC*, 1988.
- [28] A. Clement, E. L. Wong, L. Alvisi, M. Dahlin, and M. Marchetti, “Making byzantine fault tolerant systems tolerate byzantine faults,” in *NSDI*, vol. 9, 2009, pp. 153–168.
- [29] R. Cohen, S. Coretti, J. Garay, and V. Zikas, “Probabilistic termination and composability of cryptographic protocols,” *J. Cryptol.*, p. 690–741, 2019.
- [30] D. Collins, R. Guerraoui, J. Komatovic, P. Kuznetsov, M. Monti, M. Pavlovic, Y.-A. Pignolet, D.-A. Seredinschi, A. Tonkikh, and A. Xytkis, “Online payments by merely broadcasting messages,” in *DSN*. IEEE, 2020, pp. 26–38.
- [31] S. Coretti, J. Garay, M. Hirt, and V. Zikas, “Constant-round asynchronous multi-party computation based on one-way functions,” in *ASIACRYPT*, 2016.
- [32] M. Correia, N. F. Neves, and P. Veríssimo, “From consensus to atomic broadcast: Time-free byzantine-resistant protocols without signatures,” *The Computer Journal*, vol. 49, no. 1, pp. 82–96, 2006.
- [33] T. 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*, vol. abs/2002.08765, 2020. [Online]. Available: <https://arxiv.org/abs/2002.08765>
- [34] R. Cramer, I. Damgård, S. Dziembowski, M. Hirt, and T. Rabin, “Efficient multiparty computations secure against an adaptive adversary,” in *EUROCRYPT*, 1999.
- [35] G. Danezis, E. K. Kogias, A. Sonnino, and A. Spiegelman, “Narwhal and tusk: A dag-based mempool and efficient bft consensus,” ser. arxiv.org/abs/2105.11827.
- [36] S. Das, Z. Xiang, and L. Ren, “Asynchronous data dissemination and its applications,” in *CCS*, 2021, pp. 2705–2721.
- [37] S. Das, T. Yurek, Z. Xiang, A. Miller, L. Kokoris-Kogias, and L. Ren, “Practical asynchronous distributed key generation,” *IEEE Symposium on Security and Privacy*, 2022.
- [38] D. Dobre, G. O. Karame, W. Li, M. Majuntke, N. Suri, and M. Vukolic, “Proofs of writing for robust storage,” *IEEE Trans. Parallel Distributed Syst.*, vol. 30, no. 11, pp. 2547–2566, 2019.
- [39] S. Duan, H. Meling, S. Peisert, and H. Zhang, “BChain: Byzantine replication with high throughput and embedded reconfiguration,” in *OPDIS*, 2014, pp. 91–106.
- [40] S. Duan, M. K. Reiter, and H. Zhang, “Secure causal atomic broadcast, revisited,” in *DSN*.
- [41] —, “BEAT: Asynchronous bft made practical,” in *CCS*. ACM, 2018, pp. 2028–2041.
- [42] C. Dwork, N. Lynch, and L. Stockmeyer, “Consensus in the presence of partial synchrony,” *JACM*, vol. 35, no. 2, pp. 288–323, 1988.
- [43] P. Feldman and S. Micali, “Optimal algorithms for byzantine agreement,” in *STOC*, 1988.
- [44] M. J. Fischer, N. A. Lynch, and M. S. Paterson, “Impossibility of distributed consensus with one faulty process.” Massachusetts Inst of Tech Cambridge lab for Computer Science, Tech. Rep., 1982.
- [45] Y. Gao, Y. Lu, Z. Lu, Q. Tang, J. Xu, and Z. Zhang, “Efficient asynchronous byzantine agreement without private setups,” *ICDCS*, 2022.
- [46] N. Giridharan, L. Kokoris-Kogias, A. Sonnino, and A. Spiegelman, “Bullshark: Dag bft protocols made practical,” *ACM CCS*, 2022.
- [47] R. Guerraoui, N. Knežević, V. Quéma, and M. Vukolić, “The next 700 bft protocols,” *ACM Transactions on Computer Systems*, vol. 32, no. 4, pp. 12:1–12:45, 2015.
- [48] B. Guo, Y. Lu, Z. Lu, Q. Tang, J. Xu, and Z. Zhang, “Speeding dumbo: Pushing asynchronous bft closer to practice,” *NDSS*, 2022.
- [49] B. Guo, Z. Lu, Q. Tang, J. Xu, and Z. Zhang, “Dumbo: Faster asynchronous bft protocols,” in *CCS*, 2020.
- [50] J. Hendricks, G. R. Ganger, and M. K. Reiter, “Low-overhead byzantine fault-tolerant storage,” in *SOSP*, 2007.

- [51] J. Hendricks, S. Sinnamohideen, G. R. Ganger, and M. K. Reiter, “Zzyx: Scalable fault tolerance through byzantine locking,” in *DSN*. IEEE, 2010, pp. 363–372.
- [52] I. Keidar, E. Kokoris-Kogias, O. Naor, and A. Spiegelman, “All you need is dag,” in *PODC*, 2021.
- [53] E. Kushilevitz, Y. Lindell, and T. Rabin, “Information-theoretically secure protocols and security under composition,” ser. STOC, 2006.
- [54] L. Lamport, R. Shostak, and M. Pease, “The Byzantine generals problem,” *ACM Transactions on Programming Languages and Systems (TOPLAS)*, vol. 4, no. 3, pp. 382–401, 1982.
- [55] B. Libert, M. Joye, and M. Yung, “Born and raised distributively: Fully distributed non-interactive adaptively-secure threshold signatures with short shares,” *Theoretical Computer Science*, vol. 645, pp. 1–24, 2016.
- [56] C. Liu, S. Duan, and H. Zhang, “Epic: Efficient asynchronous bft with adaptive security,” in *DSN*, 2020.
- [57] —, “MiB: Asynchronous BFT with more replicas,” *CoRR*, vol. abs/2108.04488, 2021.
- [58] D. Lu, T. Yurek, S. Kulshreshtha, R. Govind, A. Kate, and A. Miller, “Honeybadgermpc and asynchromix: Practical asynchronous mpc and its application to anonymous communication,” in *CCS*, 2019.
- [59] Y. Lu, Z. Lu, Q. Tang, and G. Wang, “Dumbo-mvba: Optimal multi-valued validated asynchronous byzantine agreement, revisited,” in *Proceedings of the 39th Symposium on Principles of Distributed Computing*, 2020, pp. 129–138.
- [60] E. MacBrough, “Cobalt: Bft governance in open networks,” *arXiv preprint arXiv:1802.07240*, 2018.
- [61] A. Miller, Y. Xia, K. Croman, E. Shi, and D. Song, “The honey badger of bft protocols,” in *CCS*. ACM, 2016, pp. 31–42.
- [62] H. Moniz, N. F. Neves, M. Correia, and P. Verissimo, “Ritas: Services for randomized intrusion tolerance,” *TDSC*, vol. 8, no. 1, pp. 122–136, 2008.
- [63] A. Mostefaoui, H. Moumen, and M. Raynal, “Signature-free asynchronous Byzantine consensus with $t \leq n/3$ and $o(n^2)$ messages,” in *PODC*. ACM, 2014, pp. 2–9.
- [64] A. Mostefaoui, H. Moumen, and M. Raynal, “Signature-free asynchronous binary byzantine consensus with $t < n/3$, $o(n^2)$ messages, and $O(1)$ expected time,” *J. ACM*, vol. 62, no. 4, pp. 31:1–31:21, 2015.
- [65] A. Mostefaoui, H. Moumen, and M. Raynal, “Signature-free asynchronous binary byzantine consensus with $t_j n/3$, $o(n^2)$ messages, and $o(1)$ expected time,” *Journal of the ACM (JACM)*, vol. 62, no. 4, pp. 1–21, 2015.
- [66] A. Patra, A. Choudhury, and C. Rangan, “Asynchronous byzantine agreement with optimal resilience,” *Distrib. Comput.*, vol. 27, pp. 111–146, 2014.
- [67] M. Pease, R. Shostak, and L. Lamport, “Reaching agreement in the presence of faults,” *JACM*, vol. 27, no. 2, p. 228–234, Apr. 1980.
- [68] M. O. Rabin, “Randomized byzantine generals,” in *SFCS*. IEEE, 1983, pp. 403–409.
- [69] J. Sousa, E. Alchieri, and A. Bessani, “State machine replication for the masses with bft-smart,” in *DSN*, 2014, pp. 355–362.
- [70] G. Stern and I. Abraham, “Information theoretic hotstuff,” in *OPODIS*, 2018.
- [71] M. Vukolić, “The quest for scalable blockchain fabric: Proof-of-work vs. bft replication,” in *International workshop on open problems in network security*. Springer, 2015, pp. 112–125.
- [72] H. Zhang and S. Duan, “Pace: Fully parallelizable bft from reposable byzantine agreement,” in *CCS 2022*.

APPENDIX A BRACHA’S ABA

We present Bracha’s ABA [14]. 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 $vset$

```

01 Initialization
02  $r \leftarrow 0$  {round}
03 func propose( $v$ )
04  $iv_0 \leftarrow v$ 
05  $vset \leftarrow \{0, 1\}$  {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 {▷ phase 2}
pre-vote $_r(v)$ ,  $v \in vset$ 
10 if there are  $n - f$  pre-vote $_r(v)$ 
11 decide  $v$ 
12  $iv_{r+1} \leftarrow v$ 
13  $vset \leftarrow \{v\}$ 
14 else
15  $v \leftarrow$  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 {▷ phase 3}
main-vote $_r(v)$ ,  $v \in vset$ 
18 if there are at least  $n/2$  main-vote $_r(v)$ 
19  $vset \leftarrow \{v\}$ 
20 else
21  $v \leftarrow \{\perp\}$ 
22  $vset \leftarrow \{0, 1\}$ 
23  $r$ -broadcast final-vote $_r(v)$ 
24 upon  $r$ -delivering  $n - f$  final-vote $_r()$  such that for each
final-vote $_r(v)$ ,  $v \in vset$ ; for each final-vote $_r(*)$ ,  $vset = \{0, 1\}$ 
25 if there are at least  $2f + 1$  final-vote $_r(v)$ 
26 decide  $v$ 
27  $iv_{r+1} \leftarrow v$ 
28  $vset \leftarrow \{v\}$ 
29 else if there are  $f + 1$  final-vote $_r(v)$ 
30  $iv_{r+1} \leftarrow v$ 
31  $vset \leftarrow \{0, 1\}$ 
32 else
33  $c \leftarrow Random()$  {obtain local coin}
34  $iv_{r+1} \leftarrow c$ 
35  $vset \leftarrow \{0, 1\}$ 
32  $r \leftarrow r + 1$ 

```

Figure 11: The Bracha’s ABA protocol [14]. The code for p_i .

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 last phase of the previous round). In the first phase of the round 0, both 0 and 1 are considered valid. In the second phase and the third phase, 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 *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 $n - f$ *pre-vote* $_r()$ messages such that for each *pre-vote* $_r(v)$, $v \in vset$. There are two cases:

- Ln 10-13: If p_i has received $n - f$ *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).
- 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 $\text{main-vote}_r(v)$ message (Ln 16).

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

- Ln 18-19: If p_i receives at least $n/2$ $\text{main-vote}_r(v)$, 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 $\text{final-vote}_r(v)$ message (Ln 23). Then every replica waits for $n - f$ valid $\text{final-vote}_r()$ messages (Ln 24). Note that $\text{final-vote}_r(*)$ is considered valid only if $vset = \{0, 1\}$. There are three cases:

- Ln 25-28: If p_i receives at least $2f + 1$ $\text{final-vote}_r(v)$, 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 $f + 1$ $\text{final-vote}_r(v)$, it sets iv_{r+1} as v and $vset$ as $\{v\}$.
- Ln 32-35: 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.

APPENDIX B CUBIC-RABA

The pseudocode of Cubic-RABA protocol is shown in Figure 12. Cubic-RABA is identical to Cubic-ABA, except for round 0 (the first round). We have made the following changes for round 0. First, both $\text{propose}()$ and $\text{repropose}()$ events are allowed. Upon the $\text{propose}(v)$ event (Ln 03), a replica p_i executes the $\text{broadcast-vote}(v)$ function and starts round 0. Upon the $\text{repropose}(v)$ function (Ln 06), p_i executes $\text{broadcast-vote}(v)$. Note that upon a $\text{repropose}()$ event, p_i must have already started the protocol and may even proceed to a round greater than 0. In this case, regardless of which round the replica is in, it execute the $\text{broadcast-vote}(v)$ function and broadcasts a $\text{pre-vote}_0(v)$ message.

Second, in the $\text{broadcast-vote}(v)$ function (Ln 08-13), p_i broadcasts a $\text{pre-vote}_0(v)$ message. If $v = 1$, p_i adds 1 to $bset_0$ (Ln 11). If p_i has not previously broadcast $\text{main-vote}_0()$, it broadcasts $\text{main-vote}_0(1)$ (Ln 12). If p_i has not r -broadcast $\text{final-vote}_0()$, it r -broadcasts $\text{final-vote}_0(1)$ (Ln 13). Furthermore, in round 0, if p_i receives $f + 1$ $\text{pre-vote}_0(1)$ messages and has not broadcast $\text{main-vote}_0()$ or $\text{final-vote}_0()$ (Ln 18), it also broadcasts $\text{main-vote}_0(1)$ (Ln 20-21) and r -broadcasts $\text{final-vote}_0(1)$ (Ln 22-23).

Finally, the coin value for round 0 is set to 1 (Ln 39). In round $r \geq 1$, Cubic-RABA is identical to Cubic-ABA.

Analysis. The proof of Cubic-RABA is shown in Appendix F. We show that the changes we have made on top of Cubic-ABA can transform Cubic-ABA into a RABA protocol. The first change can ensure the *biased termination* property. In particular, it guarantees that if a quorum of correct replicas either directly propose 1 or propose 0 and later on repropose 1, the protocol will terminate. The second and the third changes ensure the biased validity property. If $f + 1$ correct replicas propose 1, they will directly add 1 to $bset_0$, broadcast $\text{pre-vote}_0(1)$, $\text{main-vote}_0(1)$, and r -broadcast $\text{final-vote}_0(1)$. Namely, no correct replica can receive $n - f$ $\text{main-vote}_0(0)$ or r -broadcast $\text{final-vote}_0(0)$. Furthermore, no correct replica can receive $n - f$ $\text{final-vote}_0(0)$ or $f + 1$ $\text{final-vote}_0(0)$. Furthermore, for the case where a correct replica uses the local coin to enter the next round, the coin value is also

```

01 initialization
02    $r \leftarrow 0$ 
03 func  $\text{propose}(v)$ 
04    $\text{broadcast-vote}(v)$ 
05   start round 0
06 func  $\text{repropose}(v)$ 
07    $\text{broadcast-vote}(v)$ 
08 func  $\text{broadcast-vote}(v)$ 
09   if  $\text{pre-vote}_0(v)$  has not been sent, broadcast  $\text{pre-vote}_0(v)$ 
10   if  $v = 1$ 
11      $bset_0 \leftarrow bset_0 \cup \{1\}$ 
12     if  $\text{main-vote}_0()$  has not been sent, broadcast  $\text{main-vote}_0(1)$ 
13     if  $\text{final-vote}_0()$  has not been sent,  $r$ -broadcast  $\text{final-vote}_0(1)$ 
14 round r
15   if  $r > 0$ , broadcast  $\text{pre-vote}_r(iv_r)$ 
16   upon receiving  $\text{pre-vote}_r(v)$  from  $f + 1$  replicas
17     if  $\text{pre-vote}_r(v)$  has not been sent, broadcast  $\text{pre-vote}_r(v)$ 
18     if  $r = 0$  and  $v = 1$ 
19        $bset_0 \leftarrow bset_0 \cup \{1\}$ 
20       if  $\text{main-vote}_0()$  has not been sent
21         broadcast  $\text{main-vote}_0(1)$ 
22       if  $\text{final-vote}_0()$  has not been sent
23          $r$ -broadcast  $\text{final-vote}_0(1)$ 
24     upon receiving  $\text{pre-vote}_r(v)$  from  $2f + 1$  nodes
25        $bset_r \leftarrow bset_r \cup \{v\}$ 
26     wait until  $bset_r \neq \emptyset$ 
27     if  $\text{main-vote}_r()$  has not been sent
28       broadcast  $\text{main-vote}_r(v)$  where  $v \in bset_r$ 
29     upon receiving  $n - f$   $\text{main-vote}_r()$  such that 1)  $\text{final-vote}_r()$ 
has not been sent; 2) for each received  $\text{main-vote}_r(b)$ ,  $b \in$ 
30     if there are  $n - f$   $\text{main-vote}_r(v)$ 
31        $r$ -broadcast  $\text{final-vote}_r(v)$ 
32     else  $r$ -broadcast  $\text{final-vote}_r(*)$ 
33     upon r-delivering  $n - f$   $\text{final-vote}_r()$  such that for each
 $\text{final-vote}_r(v)$ ,  $v \in bset_r$ ; for each  $\text{final-vote}_r(*)$ ,  $bset_r = \{0, 1\}$ 
34     if there are  $n - f$   $\text{final-vote}_r(v)$ 
35       decide  $v$ 
36     else if there are  $f + 1$   $\text{final-vote}_r(v)$ 
37        $iv_{r+1} \leftarrow v$ 
38     else
39       if  $r = 0$ ,  $iv_{r+1} \leftarrow 1$ 
40       else  $iv_{r+1} \leftarrow \text{Random}()$ 
41    $r \leftarrow r + 1$ 

```

Figure 12: Cubic-RABA. The code for p_i .

1. Accordingly, Cubic-RABA achieves biased validity. Other properties of Cubic-RABA simply follow from Cubic-ABA, as we only modify round 0 of the protocol.

APPENDIX C PROOF OF CUBIC-ABA

We show that Cubic-ABA achieves validity, agreement, termination, and integrity.

Lemma 1. *If all correct replicas propose $iv_r = v$ in round r , then any correct replica that enters round $r + 1$ sets iv_{r+1} as v .*

Proof: If all correct replicas propose v in round r , every correct replica broadcasts $\text{pre-vote}_r(v)$. No correct replica will forward $\text{pre-vote}_r(\bar{v})$, as there are no more than $f + 1$ $\text{pre-vote}_r(\bar{v})$ messages. Hence, no correct replica will add \bar{v} to $bset_r$. Furthermore, all correct replicas will eventually send $\text{main-vote}_r(v)$ and r -broadcast $\text{final-vote}_r(v)$. No correct

replica accepts $\text{final-vote}_r(\bar{v})$ or $\text{final-vote}_r(*)$, since they only have v in their $bset_r$. Hence, any correct replica that enters round $r + 1$ sets iv_{r+1} as v . ■

Note that the lemma above holds for the case where a correct replica decides v in round r .

Lemma 2. *If all correct replicas propose v in round r , then for any $r' > r$, any correct replica that enters round r' sets $iv_{r'}$ as v .*

Proof: The proof is by induction on the round number. The base case holds for r according to Lemma 1. For the induction step, we show that the lemma holds for round $r' + 1$. In other words, if all correct replicas propose $iv_{r'} = v$ in round r' , then in round $r' + 1$, any correct replica sets $iv_{r'+1}$ as v .

In round r' , as no correct replica sends $\text{pre-vote}_{r'}(\bar{v})$, no correct replica can receive $f + 1$ $\text{pre-vote}_{r'}(\bar{v})$ messages. In other words, no correct replica will forward $\text{pre-vote}_{r'}(\bar{v})$. Meanwhile, no correct replica will accept $\text{final-vote}_{r'}(\bar{v})$ or $\text{final-vote}_{r'}(*)$ since correct replicas only have v in their $bset_{r'}$. Furthermore, every correct replica will r -broadcast $\text{final-vote}_{r'}(v)$. Therefore, any correct replica that enters round $r' + 1$ sets $iv_{r'+1}$ as v . ■

Theorem 3 (Validity). *If all correct replicas propose v , then any correct replica that terminates decides v .*

Proof: We assume that a correct replica p_i terminates and decides \bar{v} and prove the correctness by contradiction.

If p_i terminates and decides \bar{v} in round 0, it will enter round 1 with $iv_1 = \bar{v}$. This is a contradiction with Lemma 1, as if all correct replicas propose v , any correct replica that enters round 1 sets iv_1 as v . If p_i terminates and decides \bar{v} in round $r > 0$, it r -delivers $n - f$ $\text{final-vote}_r(\bar{v})$. Similarly, it has put \bar{v} in its $bset_r$. Therefore, at least one correct replicas has set $iv_r = \bar{v}$ and broadcast $\text{pre-vote}_r(\bar{v})$. This is a contradiction with Lemma 2 since any correct replica that enters round r sets iv_r as v . This completes the proof of the theorem. ■

Lemma 4. *If a correct replica p_i decides v in round r , any correct replica that enters round $r + 1$ sets iv_{r+1} as v .*

Proof: If p_i decides v in round r , it r -delivers $n - f$ $\text{final-vote}_r(v)$. In other words, at least $f + 1$ correct replicas r -broadcast $\text{final-vote}_r(v)$. We assume that a correct replica p_k enters round $r + 1$ using value $iv_{r+1} = \bar{v}$ and prove the lemma by contradiction. If p_k sets iv_{r+1} as \bar{v} , there are three conditions: A) p_k r -delivers at least $n - f$ $\text{final-vote}_r(\bar{v})$; B) p_k r -delivers $f + 1$ $\text{final-vote}_r(\bar{v})$; C) none of the conditions holds. In other words, p_k has received fewer than $f + 1$ $\text{final-vote}_r(v)$ and fewer than $f + 1$ $\text{final-vote}_r(\bar{v})$. We now show that none of the three conditions is possible.

Condition A): Replica p_k r -delivers $n - f$ $\text{final-vote}_r(\bar{v})$. We already know that at least $n - f$ replicas r -broadcast $\text{final-vote}_r(v)$. Therefore, at least one correct replica r -broadcast both $\text{final-vote}_r(v)$ and $\text{final-vote}_r(\bar{v})$, a contradiction.

Condition B): Replica p_k r -delivers $f + 1$ $\text{final-vote}_r(\bar{v})$. We already know that p_i r -delivers $n - f$ $\text{final-vote}_r(v)$. Therefore, at least one replica (correct or Byzantine) r -broadcast both $\text{final-vote}_r(\bar{v})$ and $\text{final-vote}_r(v)$ such that p_k r -delivers $\text{final-vote}_r(\bar{v})$ and p_i r -delivers $\text{final-vote}_r(v)$. This is a violation of the agreement property or RBC.

Condition C): Replica p_k r -delivers $n - f$ $\text{final-vote}_r()$ messages (let the set of replicas be S_1). Among the messages

from S_1 , fewer than $f + 1$ are $\text{final-vote}_r(\bar{v})$ and fewer than $f + 1$ are $\text{final-vote}_r(v)$. Other messages can only be $\text{final-vote}_r(*)$. We already know that p_i r -delivers $n - f$ $\text{final-vote}_r(v)$ (let the set of replicas be S_2). S_1 and S_2 have at least $n - 2f \geq f + 1$ replicas in common. Therefore, at least one replica r -broadcasts both v and $*$ (or \bar{v}) such that p_i r -delivers $\text{final-vote}_r(v)$ and p_k has r -delivers $\text{final-vote}_r(*)$ (or $\text{final-vote}_r(\bar{v})$), a violation of agreement property of RBC. ■

Theorem 5 (Agreement). *If a correct replica decides v , then any correct replica that terminates decides v .*

Proof: We assume that a correct replica p_i decides v and another correct replica p_j decides \bar{v} and prove the theorem by contradiction. There are two cases: 1) p_i and p_j decide in the same round r ; 2) p_i and p_j decide in different rounds.

We first prove case 1). If replica p_i decides v in round r , it r -delivers $n - f$ $\text{final-vote}_r(v)$. If p_j decides \bar{v} , it r -delivers $n - f$ $\text{final-vote}_r(\bar{v})$. The two sets of $n - f$ replicas have at least $f + 1$ replicas in common. Among the $f + 1$ replicas, at least one is correct. Therefore, at least one correct replica must have r -broadcast both $\text{final-vote}_r(v)$ and $\text{final-vote}_r(\bar{v})$, a contradiction.

We now prove case 2) by assuming that p_i decides value v in round r and p_j decides \bar{v} in round r' where $r' > r$.

According to Lemma 4, any correct replica enters round $r + 1$ sets iv_{r+1} as v . Furthermore, according to Lemma 2, for any round $r'' \geq r + 1$, any correct replica sets enters round r'' sets $iv_{r''}$ as v . If replica p_j decides value \bar{v} in round r' , at least one correct replica has set $iv_{r'}$ as \bar{v} and sent $\text{pre-vote}_{r'}(\bar{v})$, a contradiction with Lemma 2. ■

Lemma 6. *Let $v_1 \in \{0, 1\}$ and $v_2 \in \{0, 1\}$. If a correct replica p_i r -delivers $f + 1$ $\text{final-vote}_r(v_1)$ and enters round $r + 1$, another correct replica p_j r -delivers $f + 1$ $\text{final-vote}_r(v_2)$ and enters round $r + 1$, then it holds that $v_1 = v_2$.*

Proof: If p_i r -delivers $f + 1$ $\text{final-vote}_r(v_1)$, at least one correct replica r -broadcasts $\text{final-vote}_r(v_1)$. According to the protocol, the correct replica has received $n - f$ $\text{main-vote}_r(v_1)$. Therefore, for any other correct replicas, among the $n - f$ $\text{main-vote}_r()$ messages, at least one must be $\text{main-vote}_r(v_1)$. They either receive $n - f$ $\text{main-vote}_r(v_1)$ and r -broadcast $\text{final-vote}_r(v_1)$, or receive both $\text{main-vote}_r(v_1)$ and $\text{main-vote}_r(\bar{v}_1)$ and r -broadcast $\text{final-vote}_r(*)$. No correct replica will r -broadcast $\text{final-vote}_r(\bar{v}_1)$. For replica p_j , if it r -delivers $f + 1$ $\text{final-vote}_r(v_2)$, at least one correct replica r -broadcasts $\text{final-vote}_r(v_2)$. Therefore, it must hold that $v_1 = v_2$. ■

Theorem 7 (Termination). *Every correct replica eventually decides some value.*

Proof: The proof consists of two parts. First, in each round r , correct replicas will enter the next round. Second, the value iv_r used by any correct replica cannot be manipulated by the adversary.

We first show that in round r , correct replicas will enter the next round. In each round, every replica sets iv_r as either 0 or 1 in Cubic-ABA. Accordingly, at least $f + 1$ correct replicas have the same $iv_r = v$. Therefore, all correct replicas will eventually receive $2f + 1$ $\text{pre-vote}_r(v)$ for some v and send $\text{main-vote}_r()$ message. Correct replicas will have at least v in

their $bset_r$ and r -broadcast either $\text{final-vote}_r(v)$ for some v or $\text{final-vote}_r(*)$. Similarly, any correct replica will eventually r -deliver $n-f$ $\text{final-vote}_r()$ messages and enter the next round.

We then show that if a correct replica p_i does not decide in round r , the value $iv_{r+1} = v$ cannot be manipulated by a malicious network scheduler such that correct replicas always enter the next round with inconsistent values. If p_i does not decide in round r , there are two conditions: A) p_i r -delivers $f+1$ $\text{final-vote}_r(v)$; B) p_i r -delivers $n-f$ $\text{final-vote}_r()$ messages. In the $\text{final-vote}_r()$ messages, fewer than $f+1$ are $\text{final-vote}_r(v)$ and fewer than $f+1$ are $\text{final-vote}_r(\bar{v})$. For condition B, a correct replica enters the next round with its local coin c . The c value is independent with the value chosen by any correct replica. We now prove that the value v in condition A cannot be manipulated.

According to Lemma 6, if a correct replica receives $f+1$ $\text{final-vote}_r(v_1)$ and another correct replica receives $f+1$ $\text{final-vote}_r(v_2)$, then it holds that $v_1 = v_2$. If correct replicas use local coins to enter the next round, with a probability of $\frac{1}{2^{n-f}}$, replicas will enter the next round with the same value. The protocol will reach a state where agreement can be reached in 2^{n-f} expected rounds. After that, it takes another round for each replica to terminate, i.e., the protocol terminates in $2^{n-f} + 1$ expected rounds. ■

Theorem 8 (Integrity). *No correct replica decides twice.*

Proof: According to the protocol, after a correct replica decides some value, it participates in one more round of the protocol. However, it terminates the protocol after it r -broadcasts a $\text{final-vote}_r()$ message. Thus, the replica does not decide again in the following round. This completes the proof of the theorem. ■

APPENDIX D PROOF OF QUADRATIC-ABA

We show that Quadratic-ABA achieves validity, agreement, termination, and integrity.

Lemma 9. *If all correct replicas propose $iv_r = v$ in round r , then any correct replica that enters round $r+1$ sets iv_{r+1} as v .*

Proof: If all correct replicas propose $iv_r = v$ in round r , every correct replica broadcasts $\text{pre-vote}_r(v)$. No correct replica will receive more than $f+1$ $\text{pre-vote}_r(\bar{v})$ messages. Hence, no correct replica will add \bar{v} to $bset_r$. Furthermore, all correct replicas will eventually send $\text{vote}_r(v)$ and receive $n-f$ $\text{vote}_r(v)$. As no correct replica ever has \bar{v} in $bset_r$, all correct replica will not accept $\text{vote}_r(\bar{v})$. Therefore, all correct replicas will send $\text{main-vote}_r(v)$. No correct replica will accept $\text{main-vote}_r(\bar{v})$ or $\text{main-vote}_r(*)$ as $\bar{v} \notin bset_r$ and it cannot receive more than $f+1$ $\text{vote}_r(\bar{v})$. Accordingly, every correct replicas will send $\text{final-vote}_r(v)$ and receive $n-f$ $\text{final-vote}_r(v)$. No correct replica accepts $\text{final-vote}_r(\bar{v})$ as they only have v in their $bset_r$. Hence, any correct replica that enters round $r+1$ sets iv_{r+1} as v . ■

Note that the lemma above holds for the case where a correct replica decides v in round r .

Lemma 10. *If all correct replicas propose $iv_r = v$ in round r , then for any $r' > r$, any correct replica that enters round r' sets $iv_{r'}$ as v .*

Proof: The proof is by induction on the round number.

The base case holds for r according to Lemma 9. For the induction step, we show that the lemma holds for round $r'+1$. In other words, if all correct replicas propose $iv_{r'} = v$ in round r' , then in round $r'+1$, any correct replica sets $iv_{r'+1}$ as v .

In round r' , as no correct replica sends $\text{pre-vote}_{r'}(\bar{v})$, no correct replica can receive $f+1$ $\text{pre-vote}_{r'}(\bar{v})$ messages. In other words, no correct replica will put \bar{v} to $bset_{r'}$. Therefore, all correct replicas will send $\text{vote}_{r'}(v)$ and no correct replicas will receive $f+1$ $\text{vote}_{r'}(\bar{v})$. Accordingly, all correct replicas will send $\text{main-vote}_{r'}(v)$ and will not accept $\text{main-vote}_{r'}(\bar{v})$. Any correct replica then only sends $\text{final-vote}_{r'}(v)$. Meanwhile, no correct replica will accept $\text{final-vote}_{r'}(\bar{v})$ or $\text{final-vote}_{r'}(*)$ since correct replicas only have v in their $bset_{r'}$ and no correct replica can receive $f+1$ $\text{main-vote}_{r'}(v)$. Furthermore, every correct replica will receive $n-f$ $\text{final-vote}_{r'}(v)$. It is now clear that any correct replica that enters round $r'+1$ sets $iv_{r'+1}$ as v . ■

Lemma 11. *If a correct replica p_i sends $\text{final-vote}_r(v)$, at least one correct replica has proposed $iv_r = \bar{v}$ and broadcast $\text{pre-vote}_r(\bar{v})$.*

Proof: If p_i sends $\text{final-vote}_r(v)$, it must have received $n-f$ $\text{main-vote}_r(\bar{v})$. Among the replicas that sent $\text{main-vote}_r(\bar{v})$, at least $f+1$ are correct. The correct replicas must have sent $\text{vote}_r(\bar{v})$ and put \bar{v} to $bset_r$. Each replica puts \bar{v} to $bset_r$ only if it receives $n-f$ $\text{pre-vote}_r(\bar{v})$. Therefore, at least one correct replicas has proposed $iv_r = \bar{v}$ and broadcast $\text{pre-vote}_r(\bar{v})$. ■

Theorem 12 (Validity). *If all correct replicas propose v , then any correct replica that terminates decides v .*

Proof: We assume that a correct replica p_i terminates and decides \bar{v} and prove the correctness by contradiction.

If p_i terminates and decides \bar{v} in round 0, it will enter round 1 with $iv_1 = \bar{v}$. This is a contradiction with Lemma 9. If p_i terminates and decides \bar{v} in round $r > 0$, it receives $n-f$ $\text{final-vote}_r(\bar{v})$. Among the replicas that sent $\text{final-vote}_r(\bar{v})$, at least $f+1$ are correct. According to Lemma 11, at least one correct replica has broadcast $\text{pre-vote}_r(\bar{v})$. This is a contradiction with Lemma 10 since any correct replica that enters round r sets iv_r as v . ■

Lemma 13. *If a correct replica p_i sends $\text{main-vote}_r(v)$, any correct replica p_j only sends $\text{main-vote}_r(v)$ or $\text{main-vote}_r(*)$.*

Proof: If p_i sends $\text{main-vote}_r(v)$, it has received $n-f$ $\text{vote}_r(v)$. We assume that p_j sends $\text{main-vote}_r(\bar{v})$ and prove the lemma by contradiction. If p_j sends $\text{main-vote}_r(\bar{v})$, it has received $n-f$ $\text{vote}_r(\bar{v})$. According to the protocol, every correct replica only sends $\text{vote}_r()$ message once and each replica only sends either $\text{vote}_r(v)$ or $\text{vote}_r(\bar{v})$. Therefore, at least one correct replica has sent $\text{vote}_r(v)$ to p_i and sent $\text{vote}_r(\bar{v})$ to p_j , a contradiction. ■

Lemma 14. *If a correct replica p_i sends $\text{final-vote}_r(v)$, any correct replica p_j only sends $\text{final-vote}_r(v)$ or $\text{final-vote}_r(*)$.*

Proof: If p_i sends $\text{final-vote}_r(v)$, it has received $n-f$ $\text{main-vote}_r(v)$. We assume that p_j sends $\text{final-vote}_r(\bar{v})$ and prove the lemma by contradiction. If p_j sends $\text{final-vote}_r(\bar{v})$, it has received $n-f$ $\text{main-vote}_r(\bar{v})$. According to the protocol, every correct replica only sends $\text{main-vote}_r()$ message once. Therefore, at least one correct replica has sent $\text{main-vote}_r(v)$ to p_i and sent $\text{main-vote}_r(\bar{v})$ to p_j , a contradiction. ■

Lemma 15. *If a correct replica p_i decides v in round r , any correct replica that enters round $r + 1$ sets iv_{r+1} as v .*

Proof: If p_i decides v in round r , it receives $n - f$ $\text{final-vote}_r(v)$. In other words, at least $f + 1$ correct replicas have broadcast $\text{final-vote}_r(v)$. We assume that a correct replica p_k enters round $r + 1$ sets $iv_{r+1} = \bar{v}$ and prove the lemma by contradiction. If p_k sets iv_{r+1} as \bar{v} , there are three conditions: A) p_k receives at least $n - f$ $\text{final-vote}_r(\bar{v})$; B) p_k only receives $\text{final-vote}_r(\bar{v})$ and $\text{final-vote}_r(*)$; C) none of the above holds. In other words, p_k receives only $\text{final-vote}_r(*)$ or receives both $\text{final-vote}_r(v)$ and $\text{final-vote}_r(\bar{v})$. We now show that none of the three conditions is possible.

Condition A): Replica p_k receives $n - f$ $\text{final-vote}_r(\bar{v})$. We already know that at least $n - f$ replicas have sent $\text{final-vote}_r(v)$ as p_i receives $n - f$ $\text{final-vote}_r(v)$. Therefore, at least one correct replica has sent both $\text{final-vote}_r(v)$ and $\text{final-vote}_r(\bar{v})$, a contradiction.

Condition B): Replica p_k receives $n - f$ $\text{final-vote}_r(*)$ and $\text{final-vote}_r(\bar{v})$ and has not received $\text{final-vote}_r(v)$. We already know that p_i receives $n - f$ $\text{final-vote}_r(v)$. Therefore, at least one correct replica has sent $\text{final-vote}_r(v)$ to p_i and either $\text{final-vote}_r(*)$ or $\text{final-vote}_r(\bar{v})$ to p_k , a contradiction.

Condition C): Replica p_k receives only $\text{final-vote}_r(*)$ or receives both $\text{final-vote}_r(v)$ and $\text{final-vote}_r(\bar{v})$. We know that p_i receives $n - f$ $\text{final-vote}_r(v)$. Therefore, at least $f + 1$ correct replicas have sent $\text{final-vote}_r(v)$. If p_k receives $n - f$ $\text{final-vote}_r(*)$ messages, at least one of them must be $\text{final-vote}_r(v)$. In this case, if p_k enters round $r + 1$ with iv_{r+1} as \bar{v} , p_k must have received at least one $\text{final-vote}_r(\bar{v})$, as if p_k only receives $\text{final-vote}_r(v)$ and $\text{final-vote}_r(*)$, it will set iv_{r+1} as \bar{v} . If p_k accepts $\text{final-vote}_r(v)$, it has received $f + 1$ $\text{main-vote}_r(v)$, among which at least one is sent by a correct replica. If p_k accepts $\text{final-vote}_r(\bar{v})$, it has received $f + 1$ $\text{main-vote}_r(\bar{v})$, among which at least one is sent by a correct replica. This is a contradiction with Lemma 13. ■

Theorem 16 (Agreement). *If a correct replica decides v , then any correct replica that terminates decides v .*

Proof: We assume that a correct replica p_i decides v and another correct replica p_j decides \bar{v} and prove the theorem by contradiction. There are two cases: 1) p_i and p_j decide in the same round r ; 2) p_i and p_j decide in different rounds.

We first prove case 1). If replica p_i decides v in round r , it receives $n - f$ $\text{final-vote}_r(v)$. If p_j decides \bar{v} , it receives $n - f$ $\text{final-vote}_r(\bar{v})$. The two sets of $n - f$ replicas have at least $f + 1$ replicas in common. Among the $f + 1$ replicas, at least one is correct. Therefore, at least one correct replica must have sent both $\text{final-vote}_r(v)$ and $\text{final-vote}_r(\bar{v})$, a contradiction.

We now prove case 2) by assuming that p_i decides value v in round r and p_j decides \bar{v} in round r' where $r' > r$.

According to Lemma 15, if p_i decides v , any correct replica enters round $r + 1$ sets iv_{r+1} as v . Furthermore, according to Lemma 10, for any round $r'' \geq r + 1$, any correct replica that enters round r'' sets $iv_{r''}$ as v . If replica p_j decides value \bar{v} in round r' , it has received $n - f$ $\text{final-vote}_r(v)$ so at least $f + 1$ correct replicas have sent $\text{final-vote}_r(v)$. According to Lemma 11, at least one correct replica has set $iv_{r'}$ as \bar{v} and sent $\text{pre-vote}_{r'}(\bar{v})$, a contradiction with Lemma 10. ■

Lemma 17. *If a correct replica p_i sends $\text{vote}_r(v)$ for $v \in \{0, 1\}$, any correct replica eventually accepts $\text{vote}_r(v)$.*

Proof: If p_i sends $\text{vote}_r(v)$ message, it has received $n - f$ $\text{pre-vote}_r(v)$, among which at least $f + 1$ are sent by correct replicas. Accordingly to the protocol, any correct replica that has not sent $\text{pre-vote}_r(v)$ will also send $\text{pre-vote}_r(v)$ upon receiving $f + 1$ $\text{pre-vote}_r(v)$. Therefore, every correct replica eventually sends $\text{pre-vote}_r(v)$, receives $n - f$ $\text{pre-vote}_r(v)$, and then adds v to $bset_r$. Hence, every correct replica eventually accepts $\text{vote}_r(v)$. ■

Lemma 18. *If a correct replica p_i broadcasts a $\text{main-vote}_r(v)$ or a $\text{main-vote}_r(*)$ message given that $v \in \{0, 1\}$, any correct replica accepts the $\text{main-vote}_r(v)$ message.*

Proof: If p_i sends a $\text{main-vote}_r(v)$ message, it has received and accepted $n - f$ $\text{vote}_r(v)$, among which at least $f + 1$ are sent by correct replicas. Therefore, any correct replica eventually receives $f + 1$ $\text{vote}_r(v)$ and accept $\text{vote}_r(v)$.

If p_i sends a $\text{main-vote}_r(*)$ message, it must have received and accepted both $\text{vote}_r(v)$ and $\text{vote}_r(\bar{v})$, or it has received at least one $\text{vote}_r(*)$. In any of the cases, p_i has put both 0 and 1 to $bset_r$. If p_i puts v to $bset_r$, it has received $2f + 1$ $\text{pre-vote}_r(v)$, among which at least $f + 1$ are sent by correct replicas. Then any correct replica eventually receives $f + 1$ $\text{pre-vote}_r(v)$ and sends $\text{pre-vote}_r(v)$. Every correct replica eventually receives $n - f$ $\text{pre-vote}_r(v)$ and adds v to $bset_r$. Therefore, every correct replica eventually accepts $\text{main-vote}_r(*)$. ■

Lemma 19. *If a correct replica p_i broadcasts a $\text{final-vote}_r(v)$ or a $\text{final-vote}_r(*)$ message given that $v \in \{0, 1\}$, any correct replica accepts the $\text{final-vote}_r(v)$ message.*

Proof: The lemma can be proved similarly as in Lemma 18. ■

Lemma 20. *Let $v_1 \in \{0, 1\}$ and $v_2 \in \{0, 1\}$. If a correct replica p_i receives only $n - f$ $\text{final-vote}_r(*)$ and $\text{final-vote}_r(v_1)$ messages, another correct replica p_j only receives $n - f$ $\text{final-vote}_r(v_2)$ and $\text{final-vote}_r(*)$ messages, $v_1 = v_2$.*

Proof: If p_i accepts $\text{final-vote}_r(v_1)$, it has previously received $f + 1$ $\text{main-vote}_r(v_1)$, among which at least one is sent by a correct replica. If p_j accepts $\text{final-vote}_r(v_1)$, it has previously received $f + 1$ $\text{main-vote}_r(v_2)$, among which at least one is sent by a correct replica. According to Lemma 13, it holds that $v_1 = v_2$. ■

Theorem 21 (Termination). *Every correct replica eventually decides some value.*

Proof: The proof consists of two parts. First, in each round r , correct replicas will enter the next round. Second, the value iv_r used by any correct replica cannot be manipulated by the adversary.

We first show that in round r , correct replicas will enter the next round. In each round, every replica sets iv_r to either 0 or 1 in Quadratic-ABA. Accordingly, at least $f + 1$ correct replicas have the same $iv_r = v$. All correct replicas will eventually receive $2f + 1$ $\text{pre-vote}_r(v)$ for some v and send a $\text{vote}_r(v)$ message. Correct replicas will send either $\text{vote}_r(0)$ or $\text{vote}_r(1)$ and receive at least $n - f$ $\text{main-vote}_r(v)$ messages. For any correct replica, if it sends $\text{vote}_r(v)$ for $v \in \{0, 1\}$, any correct replica will eventually accept $\text{vote}_r(v)$, according to Lemma 17. All correct replicas will then send either $\text{main-vote}_r(v)$ for $v \in \{0, 1\}$ or $\text{main-vote}_r(*)$. According

to Lemma 18, every correct replica eventually accepts any $\text{main-vote}_r()$ message sent by a correct replica. Then every correct replica either sends $\text{final-vote}_r(v)$ or $\text{final-vote}_r(*)$. According to Lemma 19, every correct replica accepts any $\text{final-vote}_r()$ message from a correct replica. Therefore, any correct replica will eventually receive $n - f$ $\text{final-vote}_r()$ messages and enter the next round.

We then show that if a correct replica p_i does not decide in round r , the value $iv_{r+1} = v$ cannot be manipulated by a malicious network scheduler such that correct replicas always enter the next round with inconsistent values. If p_i does not decide in round r , there are two conditions: A) p_i receives $n - f$ $\text{final-vote}_r(v)$ and $\text{final-vote}_r(*)$; B) p_i receives both $\text{final-vote}_r(v)$ and $\text{final-vote}_r(\bar{v})$ messages, or receives $n - f$ $\text{final-vote}_r(*)$. For condition B, a correct replica enters the next round with its local coin c . The c value is independent with the value chosen by any correct replica. We now prove that the value v in condition A cannot be manipulated.

According to Lemma 20, if p_i receives $n - f$ $\text{final-vote}_r(v_1)$ and $\text{final-vote}_r(*)$ and p_j receives $n - f$ $\text{final-vote}_r(v_1)$ and $\text{final-vote}_r(*)$, $v_1 = v_2$. In other words, the value v used by any correct replica cannot be manipulated by the network scheduler.

If correct replicas use local coins to enter the next round, with a probability of $\frac{1}{2^{n-f}}$, replicas will enter the next round with the same value. Replicas will reach a state where agreement can be reached in 2^{n-f} expected rounds and execute the protocol for another round before terminating the protocol. Therefore, the protocol will terminate in $2^{n-f} + 1$ expected rounds. ■

Theorem 22 (Integrity). *No correct replica decides twice.*

Proof: According to the protocol, after a correct replica decides some value, it participates in one more round of the protocol. However, it terminates the protocol after it receives a $\text{final-vote}_r()$ message. Hence, the replica does not decide again in the following round. ■

APPENDIX E PROOF OF CC-ABA

We prove the correctness of CC-ABA that simply replaces the local coins of Quadratic-ABA with weak common coins (or perfect common coins). According to the proofs of Quadratic-ABA, the validity, agreement, and integrity properties do not depend on the values of the coins. Therefore, validity, agreement and integrity follow those of Quadratic-ABA. We now present the following lemma and prove termination.

Lemma 23. *If a correct replica receives and accepts both $\text{final-vote}_r(v_1)$ and $\text{final-vote}_r(v_2)$ such that $v_1, v_2 \in \{0, 1\}$, $v_1 = v_2$.*

Proof: If a correct replica accepts $\text{final-vote}_r(v_1)$, it has previously received at least $f + 1$ $\text{main-vote}_r(v_1)$. If the replica accepts $\text{final-vote}_r(v_2)$, it has previously received at least $f + 1$ $\text{main-vote}_r(v_2)$. Therefore, at least one correct replica has sent $\text{main-vote}_r(v_1)$ and at least one correct replica has sent $\text{main-vote}_r(v_2)$. According to Lemma 13, if a correct replicas sends $\text{main-vote}_r(v_1)$, any correct replicas will only send $\text{main-vote}_r(v_1)$ or $\text{main-vote}_r(*)$. Therefore, we conclude that $v_1 = v_2$. ■

The proof consists of two parts. First, in each round r , correct replicas will enter the next round. Second, the value

iv_r used by any correct replica cannot be manipulated by the adversary.

We first show that in round r , correct replicas will enter the next round. In each round, every replica sets iv_r as either 0 or 1. Accordingly, at least $f + 1$ correct replicas have the same $iv_r = v$. All correct replicas will eventually receive $2f + 1$ $\text{pre-vote}_r(v)$ for some v and send $\text{vote}_r()$ message. Correct replicas will send either $\text{vote}_r(0)$ or $\text{vote}_r(1)$ and receive at least $n - f$ $\text{main-vote}_r()$ messages. For any correct replica, if it sends $\text{vote}_r(v)$ such that $v \in \{0, 1\}$, any correct replica will eventually accept $\text{vote}_r(v)$, according to Lemma 17. All correct replicas will then send either $\text{main-vote}_r(v)$ ($v \in \{0, 1\}$) or $\text{main-vote}_r(*)$. According to Lemma 18, every correct replica eventually accepts any $\text{main-vote}_r()$ message sent by a correct replica. Then every correct replica either sends $\text{final-vote}_r(v)$ or $\text{final-vote}_r(*)$. According to Lemma 19, every correct replica accepts any $\text{final-vote}_r()$ message from a correct replica. Therefore, any correct replica will eventually receives $n - f$ $\text{final-vote}_r()$ messages and enter the next round.

We then show that if a correct replica p_i does not decide in round r , the value $iv_{r+1} = v$ cannot be manipulated by a malicious network scheduler such that correct replicas always enter the next round with inconsistent values. If p_i does not decide in round r , there are two conditions: A) p_i receives $n - f$ $\text{final-vote}_r()$ messages with only $\text{final-vote}_r(v)$ and $\text{final-vote}_r(*)$; B) p_i receives both $\text{final-vote}_r(v)$ and $\text{final-vote}_r(\bar{v})$ messages, or receives $n - f$ $\text{final-vote}_r(*)$.

If condition A applies to at least two correct replicas, according to Lemma 20, if p_i receives $n - f$ $\text{final-vote}_r(v_1)$ and $\text{final-vote}_r(*)$ and p_j receives $n - f$ $\text{final-vote}_r(v_1)$ and $\text{final-vote}_r(*)$, $v_1 = v_2$. In other words, the value v used by any correct replica cannot be manipulated by an adversary.

If condition B applies to at least two correct replicas, the correct replicas enter the next round with the weak common coin. With a probability of $2/d$, all correct replicas will have the same iv_{r+1} value. This value cannot be manipulated by an adversary.

We now show that if condition A applies to a correct replica p_i and condition B applies to a correct replica p_j , the values cannot be manipulated by an adversary. If p_j sets iv_{r+1} as the weak common coin value, it has either received $n - f$ $\text{final-vote}_r(*)$ or both $\text{final-vote}_r(v)$ and $\text{final-vote}_r(\bar{v})$. According to Lemma 23, the latter case is impossible. Therefore, p_j receives $n - f$ $\text{final-vote}_r(*)$. Accordingly, at least $f + 1$ correct replicas have sent $\text{final-vote}_r(*)$. The correct replicas have previously sent either $\text{main-vote}_r(v)$ or $\text{main-vote}_r(*)$ for some $v \in \{0, 1\}$ according to Lemma 13. No correct replica will send $\text{main-vote}_r(\bar{v})$. If condition A applies to p_i and p_i sets iv_{r+1} as v_1 ($v_1 \in \{0, 1\}$), p_i has received at least $f + 1$ $\text{main-vote}_r(v_1)$. Since at least one correct replica has sent $\text{main-vote}_r(v_1)$, this value v_1 can only be v as we already know that no correct replica will send $\text{main-vote}_r(\bar{v})$. In other words, the value iv_{r+1} cannot be manipulated by an adversary.

CC-ABA uses weak common coins. If correct replicas begin the protocol with different input values, replicas will reach a state where decisions can be made in expected $1 - \sum_{r=1}^{\infty} \frac{r}{d} (1 - \frac{1}{d})^{r-1} = d$ rounds. After that, it takes another round for replicas to terminate the protocol, so the expected number of rounds is $d + 1$. For the special case that uses perfect common coins, the expected number of rounds is 3.

APPENDIX F
PROOF OF CUBIC-RABA

We now show that Cubic-RABA achieves validity, unanimous termination, agreement, biased validity, biased termination, and integrity.

Lemma 24. *If all correct replicas propose v in round 0 and never repropose \bar{v} , then any correct replica enters the round 1 sets iv_1 as v .*

Proof: In round 0, all replicas send $\text{pre-vote}_0(v)$. No correct replica will receive $f + 1$ $\text{pre-vote}_0(\bar{v})$ and send $\text{pre-vote}_0(\bar{v})$. Similarly, all correct replicas will send $\text{main-vote}_0(v)$ and will never accept $\text{main-vote}_0(\bar{v})$. All correct replicas will r -broadcast $\text{final-vote}_0(v)$ and will never accept $\text{final-vote}_0(\bar{v})$. Therefore, any correct replica that enters round 1 sets iv_1 as v . ■

Theorem 25 (Validity). *If all correct replicas propose v and never repropose \bar{v} , then any correct replica that terminates decides v .*

Proof: We assume that a correct replica p_i terminates and decides \bar{v} and prove the correctness by contradiction. If p_i terminates and decides \bar{v} in round 0, correctness follows from Lemma 24. We now prove the case where p_i decides in round $r > 0$.

Since Cubic-RABA follows Cubic-ABA starting from round 1, Lemma 2 holds for $r > 0$. If p_i terminates and decides \bar{v} in round $r > 0$, it r -delivers $n - f$ $\text{final-vote}_r(\bar{v})$. Additionally, p_i has added \bar{v} to its $bset_r$. Therefore, at least one correct replica has set iv_r as \bar{v} and broadcast $\text{pre-vote}_r(\bar{v})$. This is a contradiction with Lemma 2 since any correct replica that enters round r sets iv_r as v . This completes the proof of the theorem. ■

Theorem 26 (Unanimous termination). *If all correct replicas propose v and never repropose \bar{v} , then all correct replicas eventually terminate.*

Proof: If all correct replicas propose v and never repropose \bar{v} , all correct replicas only send $\text{pre-vote}_0(v)$. No correct replica will add \bar{v} to $bset_0$. Furthermore, no correct replica will accept $\text{main-vote}_0(\bar{v})$ or $\text{final-vote}_0(\bar{v})$. Eventually all correct replicas will receive $2f + 1$ $\text{pre-vote}_0(v)$, add v to $bset_0$, and broadcast $\text{main-vote}_0(v)$. Similarly, all correct replicas will eventually receive $n - f$ $\text{main-vote}_0(v)$ and r -broadcast $\text{final-vote}_0(v)$. All correct replicas will r -deliver $n - f$ $\text{final-vote}_0(v)$. In other words, all correct replicas will terminate and decide v . ■

Lemma 27. *If p_i decides v in round 0, any correct replica that enters round 1 sets iv_1 as v .*

Proof: If p_i decides v in round 1, it r -delivers $n - f$ $\text{final-vote}_0(v)$, among which at least $f + 1$ replicas are correct. We assume that a correct replica p_k enters round 1 with $iv_1 = \bar{v}$ and prove the correctness by contradiction. If p_k enters round $r + 1$ and sets iv_1 as \bar{v} , there are three conditions: A) p_k r -delivers at least $n - f$ $\text{final-vote}_r(\bar{v})$; B) p_k r -delivers $f + 1$ $\text{final-vote}_0(\bar{v})$; C) p_k has not received more than $f + 1$ $\text{final-vote}_0(v)$ and p_k has not received more than $f + 1$ $\text{final-vote}_0(\bar{v})$. We now show that none of the three conditions is possible.

Condition A): Replica p_i r -delivers $n - f$ $\text{final-vote}_0(\bar{v})$.

We already know that at least $f + 1$ correct replicas r -broadcast $\text{final-vote}_0(v)$. Therefore, at least one correct replica r -broadcasts both $\text{final-vote}_0(v)$ and $\text{final-vote}_0(\bar{v})$, a contradiction.

Condition B): Replica p_k r -delivers $f + 1$ $\text{final-vote}_0(\bar{v})$. We already know that p_i r -delivers $n - f$ $\text{final-vote}_0(v)$. Therefore, at least one replica (correct or Byzantine) r -broadcasts both $\text{final-vote}_0(\bar{v})$ and $\text{final-vote}_0(v)$ such that p_k r -delivers $\text{final-vote}_0(\bar{v})$ and p_i r -delivers $\text{final-vote}_0(v)$, a violation of the agreement property of RBC.

Condition C): Replica p_k r -delivers $n - f$ $\text{final-vote}_0()$ messages (let the set of replicas be S_1). In the messages, fewer than $f + 1$ are $\text{final-vote}_0(\bar{v})$ and fewer than $f + 1$ are $\text{final-vote}_0(v)$. Other messages must be $\text{final-vote}_0(*)$. We already know that p_i r -delivers $n - f$ $\text{final-vote}_0(v)$ (let the set of replicas be S_2). S_1 and S_2 have at least $n - 2f \geq f + 1$ replicas in common. In other words, at least one replica r -broadcasts a $\text{final-vote}_0()$ message such that p_i r -delivers $\text{final-vote}_0(v)$ and p_k r -delivers $\text{final-vote}_0(\bar{v})$ (or $\text{final-vote}_0(*)$), a violation of the agreement property of RBC. ■

Theorem 28 (Agreement). *If a correct replica decides v , then any correct replica that terminates decides v .*

Proof: We assume that a correct replica p_i decides v and a correct replica p_j decides \bar{v} and prove the theorem by contradiction. Since Cubic-RABA follows Cubic-ABA starting from round $r > 0$, if both p_i and p_j decide in round $r > 0$, correctness follows from the agreement property of Cubic-ABA. We now show the correctness in the following cases: 1) both p_i and p_j decide in round 0; 2) p_i decides in round 0 and p_j decides in round $r > 0$.

Case 1): If p_i decides v , it r -delivers $n - f$ $\text{final-vote}_0(v)$. If p_j decides \bar{v} , it r -delivers $n - f$ $\text{final-vote}_0(\bar{v})$. The two quorum of replicas have at least $n - 2f$ replicas in common. Among the $n - 2f$ replicas, at least one is correct since $n - 2f \geq f + 1$. Therefore, at least one correct replica r -broadcasts both $\text{final-vote}_0(v)$ and $\text{final-vote}_0(\bar{v})$, a contradiction since each replica only r -broadcasts a $\text{final-vote}_r()$ message once in each round.

Case 2): If p_j decides \bar{v} in round $r = 1$, it has received at least $2f + 1$ $\text{pre-vote}_1(\bar{v})$, where at least one correct replica has sent $\text{pre-vote}_1(\bar{v})$, a contradiction with Lemma 27. Starting from round 1, Cubic-RABA follows Cubic-ABA so that Lemma 2 holds. If p_j decides \bar{v} in round $r > 1$, at least one correct replica must have sent $\text{pre-vote}_r(\bar{v})$, a contradiction with Lemma 2 since any correct replica sets iv_r as v .

This completes the proof of the theorem. ■

Lemma 29. *If $f + 1$ correct replicas propose 1 in round 0, every replica either directly decides 1 in round 0 or/and enters round 1 with $iv_1 = 1$.*

Proof: If a correct replica p_i enters round 1, there are three conditions: A) p_i r -delivers $n - f$ $\text{final-vote}_0(v)$ with the same v ; B) p_i r -delivers at least $f + 1$ $\text{final-vote}_0(v)$ for some v ; C) none of condition A or B holds. We show that $v = 1$ for all three conditions and replicas will set iv_1 as $v = 1$.

For condition A, we already know that at least $f + 1$ correct replicas have broadcast $\text{final-vote}_0(1)$. Therefore, p_i must have received $n - f$ $\text{final-vote}_0(1)$. This is because if p_i receives $n - f$ $\text{final-vote}_0(0)$, at least one correct replica r -broadcasts both $\text{final-vote}_0(1)$ and $\text{final-vote}_0(0)$, a contradiction. In other

words, p_i decides 1.

For condition B, we assume p_i r -delivers $f + 1$ final-vote₀(0) and prove the correctness by contradiction. If p_i r -delivers $f + 1$ final-vote₀(0), at least one correct replica r -broadcasts final-vote₀(0). If the correct replica r -broadcasts final-vote₀(0), the replica must have received $n - f$ main-vote₀(0). We already know that at least $f + 1$ correct replicas have sent main-vote₀(1). Any correct replica broadcasts main-vote₀(0) message once. In other words, at least one correct replica has broadcast both main-vote₀(0) and main-vote₀(1), a contradiction. Therefore, in this condition, p_i must have r -delivered $f + 1$ final-vote₀(1). It is now clear that any correct replica uses $iv_1 = 1$ to enter round 1.

For condition C, any correct replica will use 1 as iv_1 since the local coin value is set as 1 in round 0. This completes the proof of the lemma. ■

Theorem 30 (Biased validity). *If $f + 1$ correct replicas propose 1, then any correct replica that terminates decides 1.*

Proof: We assume that a correct replica p_i decides 0 and prove the correctness by contradiction. If p_i decides in round 0, correctness follows from Lemma 29. If p_i decides 0 in round $r > 0$, at least one correct replica has set iv_r as 0 and broadcast pre-vote_r(0). Since Cubic-RABA follows Cubic-ABA starting from round 1, Lemma 2 holds. Therefore, the claim that at least one correct replica has set iv_r as 0 is a contradiction with Lemma 2. This completes the proof of the theorem. ■

Theorem 31 (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.*

Proof: The proof consists of two parts. First, every replica correct eventually enters the next round. Second, if a correct replica enters the next round with input v , v cannot be manipulated by the adversary.

We first prove that every replica eventually enters the next round. Since Cubic-RABA follows Cubic-ABA starting from round 1, this part follows from termination of Cubic-ABA. We only need to prove that every correct replica eventually enters round 1. For replicas in Q_1 , they broadcast pre-vote₀(1) and add 1 to $bset_0$. For replicas in Q_2 , they broadcast pre-vote₀(0) upon the propose(0) function, broadcast pre-vote₀(1) upon the repropose(1) function, and eventually add 1 to $bset_0$. There are two cases: 1) the size of Q_1 is greater than $f + 1$; 2) the size of Q_1 is smaller than $f + 1$.

For the first case, at least $f + 1$ replicas in Q_1 will directly broadcast main-vote₀(1) and r -broadcast final-vote₀(1). For any correct replica p_i in Q_2 , it may send main-vote₀(1) or main-vote₀(0). There are two sub-cases: none of the correct replicas send main-vote₀(0); at least one correct replica has sent main-vote₀(0). For the first sub-case, it is clear that every correct replica eventually receives and accepts $n - f$ main-vote₀(1), as every correct replica has 1 in its $bset_0$. Similarly, every correct replica will r -broadcast final-vote₀(1) and accept $n - f$ final-vote₀(1). For the second sub-case, if a correct replica p_i sends main-vote₀(0), it receives $2f + 1$ pre-vote₀(0), among which at least $f + 1$ are sent by correct replicas. Therefore, every correct replica will eventually receive $f + 1$ pre-vote₀(0) and broadcast pre-vote₀(0). Every

replica eventually adds 0 to $bset_0$. Since every correct replica has both 1 and 0 in $bset_0$, every correct replica accepts both main-vote₀(0) and main-vote₀(1). Similarly, every correct replica accepts both final-vote₀(0) and final-vote₀(1). In other words, every correct replica eventually enters the next round.

For the second case, replicas in Q_2 will send pre-vote₀(0) upon propose(0). They will send pre-vote₀(1) upon repropose(1) and add 1 to $bset_0$. Since the size of Q_2 is greater than $f + 1$ (the size of Q_1 is smaller than $f + 1$ and $Q = Q_1 \cup Q_2$), every replica will receive $f + 1$ pre-vote₀(0), send pre-vote₀(0), and add 0 to $bset_0$. Furthermore, every correct replica in Q_2 broadcasts pre-vote₀(1) upon repropose(1). Since the size of Q_2 is greater than $f + 1$, it holds that every correct replica eventually adds 1 to $bset_0$. Therefore, every replica will accept main-vote₀(0) and main-vote₀(1), final-vote₀(0), and final-vote₀(1). In other words, every correct replica eventually enters the next round.

We now prove the second part that the value iv used by any correct replica cannot be manipulated by the adversary. Since Cubic-RABA follows Cubic-ABA starting from round 1, correctness follows from Lemma 6 and termination of Cubic-ABA. ■

Theorem 32 (Integrity). *No correct replica decides twice.*

Proof: In each round, every replica only sends a main-vote_r(0) message and a final-vote_r(0) message once. Hence, only one value will be decided and integrity thus follows. ■

APPENDIX G PROOF OF QUADRATIC-RABA

We now show that Quadratic-RABA achieves validity, unanimous termination, agreement, biased validity, biased termination, and integrity.

Lemma 33. *If all correct replicas propose v in round 0 and never repropose \bar{v} , then any correct replica enters the round 1 sets iv_1 as v .*

Proof: If all correct replicas propose v in round 0, every correct replica broadcasts pre-vote₀(v). No correct replica will receive more than $f + 1$ pre-vote₀(\bar{v}) messages. Hence, no correct replica will add \bar{v} to $bset_0$. Furthermore, all correct replicas will eventually send vote₀(v) and receive $n - f$ vote₀(v). As no correct replica ever has \bar{v} in $bset_0$, all correct replica will not accept vote₀(\bar{v}). Therefore, all correct replicas will send main-vote₀(v). No correct replica will accept main-vote₀(\bar{v}) or main-vote₀(*) as $\bar{v} \notin bset_0$ and there are no more than $f + 1$ vote₀(\bar{v}). Accordingly, every correct replicas will send final-vote₀(v) and receive $n - f$ final-vote₀(v). No correct replica accepts final-vote_r(\bar{v}) as they only have v in their $bset_r$ and no correct replica can receive more than $f + 1$ final-vote₀(\bar{v}). Hence, any correct replica that enters round $r + 1$ sets iv_{r+1} as v . ■

Theorem 34 (Validity). *If all correct replicas propose v and never repropose \bar{v} , then any correct replica that terminates decides v .*

Proof: We assume that a correct replica p_i terminates and decides \bar{v} and prove the correctness by contradiction. If p_i terminates and decides \bar{v} in round 0, correctness follows from Lemma 33. We now prove the case where p_i decides in round $r > 0$.

Since Quadratic-RABA follows Quadratic-ABA starting from round 1, Lemma 10 holds for $r > 0$. If p_i terminates and decides \bar{v} in round $r > 0$, it receives $n - f$ final-vote $_r(\bar{v})$. Among the replicas that sent final-vote $_r(\bar{v})$, at least $f + 1$ are correct. According to Lemma 11, at least one correct replica has broadcast pre-vote $_r(\bar{v})$. This is a contradiction with Lemma 10 since any correct replica that enters round r sets iv_r as v . This completes the proof of the theorem. ■

Theorem 35 (Unanimous termination). *If all correct replicas propose v and never repropose \bar{v} , then all correct replicas eventually terminate.*

Proof: If all correct replicas propose v and never repropose \bar{v} , all correct replicas only send pre-vote $_0(v)$. No correct replica will add \bar{v} to $bset_0$. Furthermore, no correct replica will accept vote $_0(\bar{v})$, main-vote $_0(\bar{v})$, or final-vote $_0(\bar{v})$. Eventually all correct replicas will receive $n - f$ pre-vote $_0(v)$, add v to $bset_0$, and broadcast vote $_0(v)$. Similarly, all correct replicas will eventually receive $n - f$ vote $_0(v)$ and broadcast main-vote $_0(v)$. All correct replicas will receive $n - f$ main-vote $_0(v)$ and broadcast final-vote $_0(v)$. In other words, all correct replicas will eventually receive $n - f$ final-vote $_0(v)$ and decide v . ■

Lemma 36. *If p_i decides v in round 0, any correct replica that enters round 1 sets iv_1 as v .*

Proof: If p_i decides v in round 1, it receives $n - f$ final-vote $_0(v)$, among which at least $f + 1$ are sent by correct replicas. We assume that a correct replica p_k enters round 1 with $iv_1 = \bar{v}$ and prove the correctness by contradiction. If p_k enters round $r + 1$ and sets iv_1 as \bar{v} , there are three conditions: A) p_k receives at least $n - f$ final-vote $_r(\bar{v})$; B) p_k receive only final-vote $_0(\bar{v})$ and final-vote $_0(*)$; C) none of the above applies. In case C), as p_j will use the common coin value 1 as iv_1 , the case is impossible. We now show that none of the first two conditions is possible.

Condition A): Replica p_i receives $n - f$ final-vote $_0(\bar{v})$. We already know that at least $f + 1$ correct replicas have sent final-vote $_0(v)$. Therefore, at least one correct replica sends both final-vote $_0(v)$ and final-vote $_0(\bar{v})$, a contradiction.

Condition B): Replica p_k receives final-vote $_0(\bar{v})$ and final-vote $_0(*)$. We already know that p_i receives $n - f$ final-vote $_0(v)$. Therefore, at least one replica has sent final-vote $_0(v)$ to p_i and a final-vote $_0(\bar{v})$ (or final-vote $_0(*)$ message) to p_j , a contradiction. ■

Theorem 37 (Agreement). *If a correct replica decides v , then any correct replica that terminates decides v .*

Proof: We assume that a correct replica p_i decides v and a correct replica p_j decides \bar{v} and prove the theorem by contradiction. Since Quadratic-RABA follows Quadratic-ABA starting from round $r > 0$, if both p_i and p_j decide in round $r > 0$, correctness follows from the agreement property of Quadratic-ABA. We now show the correctness in the following cases: 1) both p_i and p_j decide in round 0; 2) p_i decides in round 0 and p_j decides in round $r > 0$.

Case 1): If p_i decides v , it receives $n - f$ final-vote $_0(v)$. If p_j decides \bar{v} , it receives $n - f$ final-vote $_0(\bar{v})$. The two quorum of replicas have at least $n - 2f$ replicas in common. Among the $n - 2f$ replicas, at least one is correct since $n - 2f \geq f + 1$.

Therefore, at least one correct replica sends both final-vote $_0(v)$ and final-vote $_0(\bar{v})$, a contradiction since each replica only sends a final-vote $_r()$ message once in each round.

Case 2): If p_j decides \bar{v} in round $r = 1$, it has received at least $n - f$ pre-vote $_1(\bar{v})$, where at least one correct replica has sent pre-vote $_1(\bar{v})$, a contradiction with Lemma 36. Starting from round 1, Quadratic-RABA follows Quadratic-ABA so that Lemma 10 holds. If p_j decides \bar{v} in round $r > 1$, at least one correct replica must have sent pre-vote $_r(\bar{v})$, a contradiction with Lemma 10 since any correct replica sets iv_r as v . ■

Lemma 38. *If $f + 1$ correct replicas propose 1 in round 0, every replica either directly decides 1 in round 0 or/and enters round 1 with $iv_1 = 1$.*

Proof: If a correct replica p_i enters round 1, there are three conditions: A) p_i receives $n - f$ final-vote $_0(v)$ with the same v ; B) p_i receives at least a final-vote $_0(v)$ message for some v ; C) none of condition A or B holds. We show that $v = 1$ for all three conditions and replicas will set iv_1 as $v = 1$.

For condition A, we already know that at least $f + 1$ correct replicas have broadcast final-vote $_0(1)$. If p_i receives $n - f$ final-vote $_0(0)$, at least one correct replica has sent both final-vote $_0(1)$ and final-vote $_0(0)$, a contradiction. In other words, in this condition p_i decides 1.

For condition B, we p_i receives only final-vote $_0(0)$ and final-vote $_0(*)$. We already know that at least $f + 1$ correct replicas have sent final-vote $_0(1)$. Therefore, at least one correct replica must have sent both final-vote $_0(1)$ and final-vote $_0(0)$ (or final-vote $_0(*)$), a contradiction.

For condition C, any correct replica will use 1 as input for round 1 since the local coin value is set as 1 in round 0. ■

Theorem 39 (Biased validity). *If $f + 1$ correct replicas propose 1, then any correct replica that terminates decides 1.*

Proof: If p_i decides in round 0, correctness follows from Lemma 38. If p_i decides 0 in round $r > 0$, at least one correct replica has set iv_r as 0 and broadcast pre-vote $_r(0)$. Since Quadratic-RABA follows Quadratic-ABA starting from round 1, Lemma 10 holds. Therefore, the claim that at least one correct replica has set iv_r as 0 is a contradiction with Lemma 10. This completes the proof of the theorem. ■

Lemma 40. *If $f + 1$ correct replicas propose 1 in round 0, every correct replica eventually accepts final-vote $_0(1)$.*

Proof: If $f + 1$ correct replicas propose 1, they will directly broadcast pre-vote $_0(1)$, vote $_0(1)$, main-vote $_0(1)$, and final-vote $_0(1)$. Every correct replica will eventually receive $f + 1$ pre-vote $_0(1)$. For those correct replicas that have not sent pre-vote $_0(1)$, they will also broadcast pre-vote $_0(1)$. Therefore, every correct replica eventually adds 1 to $bset_0$. As $f + 1$ correct replicas broadcast vote $_0(1)$, every correct replica eventually accepts main-vote $_0(1)$ message. Similarly, as $f + 1$ correct replicas broadcast main-vote $_0(1)$, every correct replica eventually accepts final-vote $_0(1)$. ■

Lemma 41. *If a correct replica sends final-vote $_r(0)$ or final-vote $_r(*)$, every correct replica eventually accepts the final-vote $_r()$ message.*

Proof: Case 1: If a correct replica sends final-vote $_r(0)$, it has received $n - f$ main-vote $_r(0)$, among which at least $f + 1$ are sent by correct replicas. Furthermore, the correct replica

has put 0 in its $bset_r$, so it receives $n-f$ $\text{pre-vote}_r(0)$. As $f+1$ correct replicas have sent $\text{pre-vote}_r(0)$, every correct replica eventually receives $f+1$ $\text{pre-vote}_r(0)$ and send $\text{pre-vote}_r(0)$. Accordingly, every correct replica puts 0 in $bset_r$. As every correct replica also eventually receives $f+1$ $\text{main-vote}_r(0)$, every correct replica will accept $\text{final-vote}_r(0)$.

Case 2: If a correct replica sends $\text{final-vote}_r(*)$, its $bset_r$ is $\{0, 1\}$, i.e., it has received both $n-f$ $\text{pre-vote}_r(0)$ and $n-f$ $\text{pre-vote}_r(1)$. Following the prior case, every correct replica eventually has $bset_r = \{0, 1\}$, so every correct replica accepts $\text{final-vote}_r(*)$. ■

Theorem 42 (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.*

Proof: The proof consists of two parts. First, every replica correct eventually enters the next round. Second, if a correct replica enters the next round with input v , v cannot be manipulated by the adversary.

We first prove that every replica eventually enters the next round. Since Quadratic-RABA follows Quadratic-ABA starting from round 1, this part follows from termination of Cubic-ABA. We only need to prove that every correct replica eventually moves to round 1. For replicas in Q_1 , they broadcast $\text{pre-vote}_0(1)$ and add 1 to $bset_0$. For replicas in Q_2 , they broadcast $\text{pre-vote}_0(0)$ upon the $\text{propose}(0)$ event, broadcast $\text{pre-vote}_0(1)$ upon the $\text{repropose}(1)$ event, and eventually add 1 to $bset_0$. There are two cases: 1) the size of Q_1 is greater than $f+1$; 2) the size of Q_1 is smaller than $f+1$.

For the first case, at least $f+1$ replicas in Q_1 will directly broadcast $\text{vote}_0(1)$, $\text{main-vote}_0(1)$, and $\text{final-vote}_0(1)$. For any correct replica p_i in Q_2 , it may send $\text{vote}_0(1)$ or $\text{vote}_0(0)$. There are two sub-cases: none of the correct replicas send $\text{vote}_0(0)$; at least one correct replica has sent $\text{vote}_0(0)$. For the first sub-case, it is straightforward to see that every correct replica eventually receives and accepts $n-f$ $\text{vote}_0(1)$, as every correct replica has 1 in its $bset_0$. Similarly, every correct replica will send $\text{main-vote}_0(1)$ and accept $n-f$ $\text{main-vote}_0(1)$. Similarly, every correct replica will send $\text{final-vote}_0(1)$. According to Lemma 40, every correct replica eventually accepts $\text{final-vote}_0(1)$ so correct replicas will enter the next round. For the second sub-case, if a correct replica p_i sends $\text{vote}_0(0)$, it receives $n-f$ $\text{pre-vote}_0(0)$, among which at least $f+1$ are sent by correct replicas. Therefore, every correct replica will eventually receive $f+1$ $\text{pre-vote}_0(0)$ and broadcast $\text{pre-vote}_0(0)$. Every replica eventually adds 0 to $bset_0$. Since every correct replica has both 1 and 0 in $bset_0$, every correct replica accepts both $\text{vote}_0(0)$ and $\text{vote}_0(1)$. Similarly, every correct replica accepts both $\text{main-vote}_0(0)$ and $\text{main-vote}_0(1)$. In other words, every correct replica may send a $\text{final-vote}_0()$ message with any value, i.e., 1, 0, or *. According to Lemma 40, every correct replica eventually accepts $\text{final-vote}_0(1)$. According to Lemma 41, any correct replica accepts the $\text{final-vote}_0()$ message sent by any correct replica. Therefore, every correct replica eventually enters the next round.

For the second case, replicas in Q_2 will send $\text{pre-vote}_0(0)$ upon $\text{propose}(0)$. They will send $\text{pre-vote}_0(1)$ upon $\text{repropose}(1)$ and add 1 to $bset_0$. Since the size of Q_2 is

greater than $f+1$ (the size of Q_1 is smaller than $f+1$ and $Q = Q_1 \cup Q_2$), every replica will receive $f+1$ $\text{pre-vote}_0(0)$, send $\text{pre-vote}_0(0)$, and add 0 to $bset_0$. Furthermore, every correct replica in Q_2 broadcasts $\text{pre-vote}_0(1)$ upon $\text{repropose}(1)$. Since the size of Q_2 is greater than $f+1$, every correct replica eventually adds 1 to $bset_0$. According to the protocol, in round 0, every correct replica accepts $\text{main-vote}_0(v)$ and $\text{final-vote}_0(v)$ if v is added to $bset_0$. Therefore, every replica will accept both $\text{vote}_0(0)$ and $\text{vote}_0(1)$, and $\text{main-vote}_0()$ and $\text{final-vote}_0()$ with any value. Accordingly, every correct replica eventually enters the next round.

We not prove the second part where the value iv used by any correct replica cannot be manipulated by the adversary. Since Quadratic-RABA follows Quadratic-ABA starting from round 1, correctness follows from Lemma 20 and termination of Quadratic-ABA. ■

Theorem 43 (Integrity). *No correct replica decides twice.*

Proof: In each round, every replica only sends a $\text{final-vote}_r()$ message once. Hence, only one value will be decided and integrity thus follows. ■