As easy as ABC: Optimal (A)ccountable (B)yzantine (C)onsensus is easy!

Pierre Civit[1], Seth Gilbert[2], Vincent Gramoli[3,4], Rachid Guerraoui[4] and Jovan Komatovic[4]

[1]Sorbonne University, CNRS, LIP6
[2]NUS Singapore
[3]University of Sydney
[4]EPFL

***Abstract*—It is known that the agreement property of the Byzantine consensus problem among n processes can be violated in a non-synchronous system if the number of faulty processes exceeds $t_0 = \lceil n/3 \rceil - 1$ [9], [17]. In this paper, we investigate the *accountable* Byzantine consensus problem in non-synchronous systems: the problem of solving Byzantine consensus whenever possible (i.e., when the number of faulty processes does not exceed t_0) and allowing correct processes to obtain a proof of culpability of $t_0 + 1$ faulty processes whenever correct processes disagree. We present four complementary contributions:**

1) **We introduce $\mathcal{ABC}$: a simple yet efficient transformation of any Byzantine consensus protocol to an accountable one. $\mathcal{ABC}$ introduces an overhead of (1) only two all-to-all communication rounds and $O(n^2)$ additional bits in executions with up to t_0 faults, and (2) three all-to-all communication rounds and $O(n^3)$ additional bits in executions with more faults.**
2) **We prove a tight lower bound on the communication complexity needed for any accountable Byzantine consensus protocol. In particular, we show that any algorithm incurs a cubic communication complexity in an execution in which disagreement occurs and that this bound is tight by applying $\mathcal{ABC}$ to the binary DBFT consensus protocol [10].**
3) **We demonstrate that, when applied to an optimal Byzantine consensus protocol, $\mathcal{ABC}$ constructs an accountable Byzantine consensus protocol that is (1) optimal in solving consensus whenever consensus is solvable, and (2) optimal in obtaining accountability whenever disagreement happens.**
4) **We generalize $\mathcal{ABC}$ to other distributed computing problems besides the classic consensus problem. We characterize a class of agreement tasks, including reliable and consistent broadcast [4], that $\mathcal{ABC}$ renders accountable.**

I. Introduction

Ensuring both safety ("nothing bad ever happens") and liveness ("something good eventually happens") of a wide variety of distributed Byzantine problems is impossible if the number of Byzantine processes exceeds a certain predefined threshold [17]. This limitation motivated researchers to investigate *accountable* variants of these problems [8], [21]. The accountable variant of a problem $\mathcal{P}$ consists in (1) solving problem $\mathcal{P}$ under the appropriate assumptions (e.g., whenever the number of Byzantine processes does not exceed the threshold), and (2) allowing all correct participants to detect some fraction of culprits if the safety of problem $\mathcal{P}$ is violated. Accountability in distributed systems is important since it discourages bad behavior. If malicious behavior is guaranteed to result in apprehension and punishment, malicious processes are much less likely to carry out an attack in the first place, thus strengthening the security of the system.

This paper primarily focuses on obtaining accountability in Byzantine consensus protocols that operate in non-synchronous systems. The Byzantine consensus problem [17] is defined among n processes while tolerating up to $t_0 = \lceil n/3 \rceil - 1$ Byzantine (malicious) processes. A process initially *proposes* a value and eventually *decides* a value such that the following properties hold:

- (Liveness) *Termination:* Every correct process eventually decides.
- (Safety) *Agreement:* All correct processes decide the same value.
- (Safety) *Validity:* If all correct processes propose the same value, only that value can be decided by a correct process.

The conjunction of the aforementioned properties can only be ensured if the number of faulty processes does not exceed t_0 [17]. If indeed the faulty processes overpopulate the system, any of these properties might be violated. This work focuses on cases when violation of the agreement property occurs. Specifically, we take a closer look at the *accountable* Byzantine consensus problem. A process initially proposes and later decides a value (as in the Byzantine consensus problem) and *detects* some faulty processes. Formally, the accountable Byzantine problem is solved if and only if the following properties are ensured:

- *Termination:* If the number of faulty processes does not exceed t_0, then every correct process eventually decides.
- *Agreement:* If the number of faulty processes does not exceed t_0, then all correct processes decide the same value.
- *Validity:* If the number of faulty processes does not exceed t_0 and all correct processes propose the same value, only that value can be decided by a correct

pierre.civit@lip6.fr
seth.gilbert@comp.nus.edu.sg
vincent.gramoli@sydney.edu.au
rachid.guerraoui@epfl.ch
jovan.komatovic@epfl.ch

process.
- *Accountability:* If two correct processes decide different values, then every correct process eventually detects at least $t_0 + 1$ faulty processes and obtains a proof of culpability of all detected processes.

A. Contributions

The contributions of the paper are fourfold:

1) We present a generic and simple transformation - $\mathcal{ABC}$ – that allows *any* Byzantine consensus protocol to obtain accountability. Our transformation is efficient: it introduces an overhead of (1) only two all-to-all communication rounds and $O(n^2)$ exchanged bits of information in all executions with at most t_0 faulty processes (i.e., in the common case), and (2) three all-to-all communication rounds and $O(n^3)$ exchanged bits of information otherwise (i.e., in the degraded case).
$\mathcal{ABC}$ owns its simplicity and efficiency to the observation that the composition presented in Algorithm 1 solves the Byzantine consensus problem. Indeed, if the number of faults does not exceed t_0, all processes eventually broadcast and receive $n - t_0$ matching CONFIRM messages. However, the important mechanism illustrated in Algorithm 1 is that faulty processes *must* send conflicting CONFIRM messages in order to cause disagreement. Hence, whenever correct processes disagree, they only need to exchange received CONFIRM messages to obtain accountability.

Algorithm 1 Intuition Behind $\mathcal{ABC}$ Transformation

```
1: function propose(v) do
2:     ▷ bc is any Byzantine consensus protocol
3:     v' ← bc.propose(v)
4:     broadcast [CONFIRM, v']
5:     wait for n − t0 [CONFIRM, v']
6:     return v'
```

2) We show that our $\mathcal{ABC}$ transformation, despite its simplicity, suffices for achieving *optimal* communication complexity in providing accountability. Namely, we prove that any accountable Byzantine consensus incurs cubic communication complexity in an execution in which disagreement occurs and we demonstrate that the lower bound is tight by applying $\mathcal{ABC}$ to a cubic Byzantine consensus protocol (e.g., binary DBFT [10]).
3) We demonstrate that, when applied to an optimal (with respect to the communication complexity) Byzantine consensus protocol, $\mathcal{ABC}$ produces an accountable Byzantine consensus protocol that is (1) optimal in solving consensus whenever consensus is solvable, and (2) optimal in obtaining accountability whenever disagreement occurs.
4) We show that $\mathcal{ABC}$ is not limited to Byzantine consensus. Specifically, we define a class of *easily accountable agreement tasks* and we demonstrate that generalized $\mathcal{ABC}$ transformation indeed provides accountability for such tasks. Important distributed tasks, like Byzantine reliable [5] and Byzantine consistent [5] broadcast, fall into the class of easily accountable agreement tasks.

B. Related Work

The work on accountability in distributed systems was pioneered in [15]. The authors presented PeerReview - a generic accountability layer for distributed systems. Importantly, PeerReview does not allow correct processes to irrefutably detect faulty processes in non-synchronous environments, i.e., faulty processes might be *suspected forever*, but never detected. Therefore, PeerReview does not suffice for accountability in non-synchronous Byzantine consensus protocols. Same authors initiated the formal study of Byzantine failures in the context of accountability [16].

Recently, with the expansion of blockchain systems, the interest in accountable distributed protocols resurfaced once again. Polygraph [8] - the first accountable Byzantine consensus protocol - was introduced by Civit *et al.* The Polygraph protocol is based on DBFT [10], tolerates up to n faulty processes in achieving accountability[1] and has the communication complexity of $O(n^4)$, where n denotes the total number of processes. Casper [3] is another system designed around the goal of obtaining accountability. Most recently, authors of [20] investigated the possibility of obtaining accountability in protocols based on PBFT [6] in scenarios in which the system is not severely corrupted. Specifically, they present variants of PBFT [6] and HotStuff [22] that achieve accountability without worsening the communication complexity of the base consensus protocol; however, they allow for accountability only if up to $2n/3$ processes are faulty, which implies that their "accountability threshold" is lower than the one of Polygraph. The commonality between the discussed prior work is employing sophisticated mechanisms for obtaining accountability. Indeed, the prior work achieves accountability with the help of non-trivial modifications applied to the base consensus protocol. In contrast, we take a fundamentally different approach that allows us to treat the base consensus protocol as a "black box", thus obtaining simpler and more efficient accountable Byzantine consensus protocols. Table I compares accountable Byzantine consensus protocols obtained by $\mathcal{ABC}$ with the existing alternatives.

a) Roadmap: We present the system model in §II. We devote §III to our $\mathcal{ABC}$ transformation. Specifically, we first introduce the novel accountable confirmer problem (§III-A), the crucial building block of $\mathcal{ABC}$. Then, we present $\mathcal{ABC}$ and prove its correctness (§III-B). In §III-C, we demonstrate that $\mathcal{ABC}$ suffices for obtaining optimal communication complexity in accountable Byzantine consensus protocols. We define easily accountable agreement tasks and prove the applicability of generalized $\mathcal{ABC}$ to such tasks in §IV. Finally, we conclude the paper in §V.

II. Model

We consider a system with a set $\{P_1, ..., P_n\}$ of n processes that communicate by exchanging messages through a point-

[1]Note that disagreement cannot occur if the number of faulty processes exceeds $n - 2$. Hence, satisfying accountability in executions with $n - 1$ or n faults is trivial.

Base Consensus Protocol	Communication Complexity of the Base Consensus Protocol	Communication Complexity of the Accountable Variant in the Common Case	Communication Complexity of the Accountable Variant in the Degraded Case	Accountability Threshold	Paper
PBFT [6]	$O(n^4)$	$O(n^4)$	$O(n^4)$	$2n/3$	[20]
HotStuff [22]	$O(n^3)$	$O(n^3)$	$O(n^3)$	$2n/3$	[20]
Binary DBFT [10]	$O(n^3)$	$O(n^4)$	$O(n^4)$	n	[8]
Multivalue DBFT [10]	$O(n^4)$	$O(n^4)$	$O(n^4)$	n	[8]
Any	X	X	$max(X, O(n^3))$	n	**this paper**

TABLE I: Overview of the main properties of existing accountable Byzantine consensus protocols. We consider worst-case communication complexities in all columns.

to-point network. The system is *non-synchronous*: there is no bound that always holds on message delays and relative speed of processes. Non-synchronous systems include:

- asynchronous systems, where the bound does not exist, and
- partially synchronous systems [13], where the bound holds only eventually.

All our results given in the present paper assume a non-synchronous system.

Each process is assigned its *local protocol* to execute. A local protocol of a process defines steps to be taken by the process during a run of the system. The collection of all local protocols assigned to processes is referred to as a *distributed protocol* (or simply a *protocol*).

A subset of all processes might be *faulty*: these processes may arbitrarily deviate from their local protocol, i.e., we consider the Byzantine failure model. If a process is not faulty, we say that the process is *correct*. We assume that any message sent by a correct process to a correct process is eventually received, i.e., we assume that communication is *reliable*. Moreover, we assume that the order of message receptions is controlled by a computationally bounded adversary. An *execution* of the system is a single run of the system, i.e., it is a sequence of sending and receiving events, as well as the internal events of processes. We denote by t the actual number of faulty processes in an execution. Finally, we denote by $\mathbb{P}(X)$ the power set of a set X.

a) Cryptographic Primitives: We assume an idealized *public-key infrastructure* (PKI): each process is associated with its own public/private key pair that is used to sign messages and verify signatures of other processes. A message m sent by a process P_i that is properly signed with the PKI private key of P_i is said to be *properly authenticated*. We denote by m_{σ_i} a message m signed with the PKI private key of a process P_i.

Moreover, we assume a (k, n)-threshold signature scheme [18], where $k = n - \lceil n/3 \rceil + 1$. In this scheme, each process holds a distinct private key and there exists a single public key. Each process P_i can use its private key to produce a partial signature of a message m by invoking $ShareSign_i(m)$. Moreover, a partial signature $tsignature$ of a message m produced by process P_i could be verified with $ShareVerify_i(m, tsignature)$. Finally, set $S = \{tsignature_i\}$ of partial signatures, where $|S| = k$ and, for each $tsignature_i \in S$, $tsignature_i = ShareSign_i(m)$, could be combined into a *single* digital signature by invoking $Combine(S)$; a combined digital signature $tcombined$ of message m could be verified with $Verify(m, tcombined)$. In the paper, we assume that the cost of obtaining the threshold signature scheme [1] is amortized and, thus, negligible.

Crucially, we assume that the PKI private key of a correct process is *never* revealed (irrespectively of the number of faulty processes in the system). Therefore, if a message m is signed with the PKI private key of a process P_i and P_i is correct, then the message m was certainly sent by P_i. Conversely, if the number of faulty processes exceeds $n - k$, the threshold private key of a process *can* be revealed and faulty processes might forge a partial signature of a correct process.

b) Proof of Culpability: We say that a set $\mathcal{S}$ of properly authenticated messages sent by a process P_i is a *proof of culpability* of P_i if and only if there does not exist an execution e of the system where (1) P_i sends all the messages from the $\mathcal{S}$ set, and (2) P_i is correct. Observe that a proof of culpability of a process contains messages signed by the process with its PKI private key. Indeed, the PKI private key of a correct process is *never* revealed (as opposed to the threshold private key of a correct process that might be revealed if the number of faults exceeds $n - k$, where $k = n - \lceil n/3 \rceil + 1$), which implies that a proof of culpability of a correct process can *never* be obtained.

c) Complexity Measure: In this work, as in many in distributed computing, we care about the *communication complexity* which is the maximum number of authenticators sent by all correct processes combined across all executions of the system. An *authenticator* is either a partial signature or a signature.

III. $\mathcal{ABC}$ Transformation

This section presents $\mathcal{ABC}$, our transformation that enables any Byzantine consensus protocol to obtain accountability. To this end, we first introduce the *accountable confirmer* problem and give its implementation (§III-A). Then, we construct our $\mathcal{ABC}$ transformation around accountable confirmer (§III-B). In §III-C, we prove that $\mathcal{ABC}$ allows for obtaining optimal communication complexity in accountable Byzantine consensus protocols. Finally, we conclude the section with a brief discussion about the applicability of $\mathcal{ABC}$ and communication optimality it provides (§III-D).

A. Accountable Confirmer

The accountable confirmer problem is a distributed problem defined among n processes. The problem is associated

with parameter $t_0 = \lceil n/3 \rceil - 1$ emphasizing that some properties are ensured only if the number of faulty processes does not exceed t_0[2]. Accountable confirmer exposes the following interface: (1) request $submit(v)$ - a process *submits* value v; invoked at most once, (2) indication $confirm(v')$ - a process *confirms* value v'; triggered at most once, and (3) indication $detect(F, proof)$ - a process *detects* processes from the set F such that $proof$ represents a proof of culpability of all processes that belong to F; triggered at most once. The following properties are ensured:

- *Terminating Convergence:* If the number of faulty processes does not exceed t_0 and all correct processes submit the same value, then that value is eventually confirmed by every correct process.
- *Agreement:* If the number of faulty processes does not exceed t_0, then no two correct processes confirm different values.
- *Validity:* Value confirmed by a correct process was submitted by a correct process.
- *Accountability:* If two correct processes confirm different values, then every correct process eventually detects at least $t_0 + 1$ faulty processes and obtains a proof of culpability of all detected processes.

Terminating convergence ensures that, if (1) the number of faults does not exceed t_0, and (2) all correct processes submit the same value, then all correct processes eventually confirm that value[3]. Agreement stipulates that no two correct processes confirm different values if the system is not corrupted (even if submitted values of correct processes differ). Validity ensures that any confirmed value is submitted by a correct process. Finally, accountability ensures detection of $t_0 + 1$ faulty processes by every correct process whenever correct processes confirm different values.

a) Implementation: We now give an implementation of the accountable confirmer problem (Algorithm 2). The implementation takes advantage of threshold signatures (see §II) in order to obtain quadratic communication complexity in the common case (i.e., in executions with up to t_0 faulty processes). In the degraded case (i.e., in executions with more than t_0 faulty processes), the complexity is cubic.

Each process initially broadcasts the value it submitted in a SUBMIT message (line 18): the SUBMIT message contains the value and the partial signature of the value. Moreover, the entire message is signed with the PKI private key of the sender. Once a process receives such a SUBMIT message, the process (1) checks whether the message is properly signed (line 6), (2) verifies the partial signature (line 20), and (3) checks whether the received value is equal to its submitted value (line 20). If all of these checks pass, the process stores the received partial signature (line 22) and the entire message (line 23). Once a process stores partial signatures from (at least) $n-t_0$ processes (line 25), the process confirms its submitted value (line 27) and informs other processes about its confirmation by combining the received partial signatures into a *light certificate* (line 28). The role of threshold signatures in our implementation is to allow a light certificate to contain a *single* signature, thus obtaining quadratic overall communication complexity if $t \leq t_0$.

Once a process receives two conflicting light certificates (line 33), the process concludes that correct processes might have indeed confirmed different values[4]. If the process has already confirmed its value, the process broadcasts the set of (at least) $n-t_0$ properly authenticated $[\text{SUBMIT}, v, *]$ messages (line 35), where v is the value confirmed (and submitted) by the process; such set of messages is a *full certificate* for value v. Finally, once a process receives two conflicting full certificates (line 40), the process obtains a proof of culpability of (at least) $t_0 + 1$ faulty processes (line 43), which ensures accountability. Indeed, each full certificate contains $n - t_0$ properly authenticated messages: every process whose message is in both full certificates is faulty and these messages represent a proof of its misbehavior (recall that no faulty process *ever* obtains the PKI private key of a correct process).

Accountable Confirmer - Definitions for Algorithm 2

1) A combined digital signature $tsig$ is a *valid light certificate for value* v if and only if $Verify(v, tsig) = \top$.

2) A set $\mathcal{S}$ of properly authenticated $[\text{SUBMIT}, v, *]_{\sigma_*}$ messages is a *valid full certificate for value* v if and only if:
a) $|\mathcal{S}| \geq n - t_0$
b) Each message m is sent (i.e., signed) by a distinct process.

3) Let $tsig_v$ be a valid light certificate for value v and let $tsig_{v'}$ be a valid light certificate for value v'. $tsig_v$ *conflicts* with $tsig_{v'}$ if and only if $v \neq v'$.

4) Let $\mathcal{S}_v$ be a valid full certificate for value v and let $\mathcal{S}_{v'}$ be a valid full certificate for value v'. $\mathcal{S}_v$ *conflicts* with $\mathcal{S}_{v'}$ if and only if $v \neq v'$.

5) Let (m_1, m_2) be a pair of properly authenticated messages sent by the same process P_i. (m_1, m_2) is a *proof of culpability of* P_i if and only if:
a) $m_1 = [\text{SUBMIT}, v, share_1]_{\sigma_i}$
b) $m_2 = [\text{SUBMIT}, v', share_2]_{\sigma_i}$
c) $v \neq v'$.

Theorem 1. *Algorithm 2 solves the accountable confirmer problem with:*

- $O(n^2)$ *communication complexity in the common case, and*
- $O(n^3)$ *communication complexity in the degraded case.*

Proof. We start by proving the terminating convergence property. Indeed, if $t \leq t_0$ and all correct processes submit the same value v, then the rule at line 25 eventually triggers at each correct process. Since each correct process confirms

[2]Recall that $t_0 = \lceil n/3 \rceil - 1$ is the number of faulty processes tolerated by the Byzantine consensus problem.

[3]Note that it is *not* guaranteed that any correct process eventually confirms a value if correct processes submit different values (even if the number of faulty processes does not exceed t_0).

[4]Observe that the process is not certain that correct processes have confirmed different values because light certificates could be sent by faulty processes (possible only if $t > t_0$).

Algorithm 2 Accountable Confirmer - Code for Process P_i

1: **Implements:**
2: Accountable Confirmer, **instance** ac
3: **Uses:**
4: Best-Effort Broadcast [4], **instance** beb ▷ Simple broadcast without any guarantees if the sender is faulty.
5: **Rules:**
6: 1) Any SUBMIT message that is not properly authenticated is discarded.
7: 2) Rules at lines 25, 33, 34 and 40 are activated at most once.
8: **upon event** $\langle ac, Init \rangle$ **do**
9: $value_i \leftarrow \bot$
10: $confirmed_i \leftarrow false$
11: $from_i \leftarrow \emptyset$
12: $lightCertificate_i \leftarrow \emptyset$
13: $fullCertificate_i \leftarrow \emptyset$
14: $obtainedLightCertificates_i \leftarrow \emptyset$
15: $obtainedFullCertificates_i \leftarrow \emptyset$
16: **upon event** $\langle ac, Submit \mid v \rangle$ **do**
17: $value_i \leftarrow v$
18: **trigger** $\langle beb, Broadcast \mid [\text{SUBMIT}, v, ShareSign_i(v)]_{\sigma_i} \rangle$
19: **upon event** $\langle beb, Deliver \mid P_j, [\text{SUBMIT}, value, share]_{\sigma_j} \rangle$ **do**
20: **if** $ShareVerify_j(value, share) = \top$ and $value = value_i$ and $P_j \notin from_i$ **then**
21: $from_i \leftarrow from_i \cup \{P_j\}$
22: $lightCertificate_i \leftarrow lightCertificate_i \cup \{share\}$
23: $fullCertificate_i \leftarrow fullCertificate_i \cup \{[\text{SUBMIT}, value, share]_{\sigma_j}\}$
24: **end if**
25: **upon** $|from_i| \geq n - t_0$ **do**
26: $confirmed_i \leftarrow true$
27: **trigger** $\langle ac, Confirm \mid value_i \rangle$
28: **trigger** $\langle beb, Broadcast \mid [\text{LIGHT-CERTIFICATE}, value_i, Combine(lightCertificate_i)] \rangle$
29: **upon event** $\langle beb, Deliver \mid P_j, [\text{LIGHT-CERTIFICATE}, value_j, lightCertificate_j] \rangle$ **do**
30: **if** $lightCertificate_j$ is a valid light certificate **then**
31: $obtainedLightCertificates_i \leftarrow obtainedLightCertificates_i \cup \{lightCertificate_j\}$
32: **end if**
33: **upon** $certificate_1, certificate_2 \in obtainedLightCertificates_i$ where $certificate_1$ conflicts with $certificate_2$
34: and $confirmed_i = true$ **do**
35: **trigger** $\langle beb, Broadcast \mid [\text{FULL-CERTIFICATE}, value_i, fullCertificate_i] \rangle$
36: **upon event** $\langle beb, Deliver \mid P_j, [\text{FULL-CERTIFICATE}, value_j, fullCertificate_j] \rangle$ **do**
37: **if** $fullCertificate_j$ is a valid full certificate **then**
38: $obtainedFullCertificates_i \leftarrow obtainedFullCertificates_i \cup \{fullCertificate_j\}$
39: **end if**
40: **upon** $certificate_1, certificate_2 \in obtainedFullCertificates_i$ where $certificate_1$ conflicts with $certificate_2$ **do**
41: $proof \leftarrow$ extract a proof of culpability of (at least) $t_0 + 1$ processes from $certificate_1$ and $certificate_2$
42: $F \leftarrow$ set of processes detected via $proof$
43: **trigger** $\langle ac, Detect \mid F, proof \rangle$

only the value it has submitted (line 27), the property is satisfied by Algorithm 2.

We prove agreement by contradiction. Let a correct process P_i confirm value v, let another correct process P_j confirm value $v' \neq v$ and let $t \leq t_0$. Hence, P_i (resp., P_j) has received $n - t_0$ SUBMIT messages for value v (resp., v'). Given that $t_0 < n/3$, we conclude that number of processes that have sent the SUBMIT message for both values must be greater than t_0. This implies that there are more than t_0 faulty processes, which contradicts the fact that $t \leq t_0$. Therefore, the agreement property is ensured.

Validity trivially follows from the fact that each correct process confirms only the value it has submitted (line 27).

We now prove accountability. Let a correct process P_i confirm value v and let another correct process P_j confirm value $v' \neq v$. The rule at lines 33 and 34 is eventually triggered at each correct process that confirms a value. Once the rule is triggered at P_i (resp., P_j), the process broadcasts its full certificate to all processes (line 35). Eventually, the rule at line 40 is triggered at each correct process, which ensures

accountability. Indeed, every process whose SUBMIT messages belong to both conflicting full certificates is detected; moreover, such process is indeed faulty since no correct process submits different values, hence, no correct process ever sends different SUBMIT messages.

Finally, we prove the claimed communication complexity:

- If $t \leq t_0$, the communication complexity of the algorithm is quadratic because (1) light certificates are sent only once and they contain a single signature, and (2) no correct process ever sends its full certificate.
- If $t > t_0$, the communication complexity is cubic. Indeed, broadcasting of a full certificate (that contains $O(n)$ authenticators) dominates the communication complexity in this case. Therefore, each correct process sends $O(n)$ authenticators to all processes (line 35), which results in the cubic overall communication complexity.

□

B. $\mathcal{ABC}$: Byzantine Consensus + Accountable Confirmer = Accountable Byzantine Consensus

We now define our $\mathcal{ABC}$ transformation (Algorithm 3), the main contribution of our work. $\mathcal{ABC}$ is built on the observation that any Byzantine consensus protocol paired with accountable confirmer solves the accountable Byzantine consensus problem.

Algorithm 3 $\mathcal{ABC}$ Transformation - Code For Process P_i

```
1:  Implements:
2:      Accountable Byzantine Consensus, instance abc
3:  Uses:
4:      ▷ Byzantine consensus protocol to be transformed
5:      Byzantine Consensus, instance bc
6:      Accountable Confirmer, instance ac
7:  upon event ⟨abc, Propose | proposal⟩ do
8:      trigger ⟨bc, Propose | proposal⟩
9:  upon event ⟨bc, Decide | decision⟩ do
10:     trigger ⟨ac, Submit | decision⟩
11: upon event ⟨ac, Confirm | confirmation⟩ do
12:     trigger ⟨abc, Decide | confirmation⟩
13: upon event ⟨ac, Detect | F, proof⟩ do
14:     trigger ⟨abc, Detect | F, proof⟩
```

The following theorem shows that Algorithm 3 solves the accountable Byzantine consensus problem, which implies that $\mathcal{ABC}$ indeed allows Byzantine consensus protocols to obtain accountability.

Theorem 2. *Algorithm 3 solves the accountable Byzantine consensus problem.*

Proof. Consider an execution where $t \leq t_0$. All correct processes eventually decide the same value v from Byzantine consensus at line 9 (by termination and agreement of Byzantine consensus). Moreover, if all correct processes have proposed the same value (line 7), then the proposed value is indeed v (ensured by validity of Byzantine consensus). Terminating convergence of accountable confirmer ensures that all correct processes eventually confirm v (line 11) and decide from the accountable Byzantine consensus (line 12). Hence, Algorithm 3 satisfies termination, agreement and validity if $t \leq t_0$.

If correct processes disagree (i.e., decide different values at line 12), then these processes have confirmed different values from accountable confirmer (line 11). Thus, accountability is ensured by Algorithm 3 since accountability is ensured by accountable confirmer, i.e., every correct process eventually detects faulty processes from accountable confirmer (line 13). Thus, accountability is satisfied by Algorithm 3, which concludes the theorem. □

Finally, we note that $\mathcal{ABC}$ does not worsen the communication complexity of any Byzantine consensus protocol. It is well-known that any protocol that solves the Byzantine consensus problem incurs quadratic communication complexity due to the lower bound set by Dolev *et al.* [12]. Given the fact that accountable confirmer has quadratic communication complexity in the common case (Theorem 1), every Byzantine consensus protocol *retains* its complexity after our transformation.

Corollary 1. Let Π be a Byzantine consensus protocol with the communication complexity X_Π. Let Π^A be a protocol obtained by applying $\mathcal{ABC}$ to Π. Then, Π^A solves the Byzantine consensus problem with the communication complexity X_Π.

C. $\mathcal{ABC}$ Suffices For Optimal Accountability

This subsection proves that any distributed protocol that solves the accountable Byzantine consensus problem incurs cubic communication cost. Moreover, we show that the lower bound is tight by applying $\mathcal{ABC}$ (§III-B) to any cubic (or subcubic) Byzantine consensus protocol (Corollary 2). Therefore, we show that our simple transformation allows Byzantine consensus protocols to obtain accountability *optimally* with respect to the communication complexity.

Let Π^A be a distributed protocol that solves the accountable Byzantine consensus problem among n processes. If up to $t_0 = \lceil n/3 \rceil - 1$ processes are faulty, Π^A ensures termination, agreement and validity; if disagreement occurs, each correct process eventually detects at least $t_0 + 1$ faulty processes (and obtains a proof of culpability of all detected processes). Without loss of generality, let $n = 3t_0 + 1$.

We start by separating processes that execute Π^A into three disjoint groups: (1) group A, where $|A| = t_0$, (2) group B, where $|B| = t_0 + 1$, and (3) group C, where $|C| = t_0$. Given that Π^A solves the Byzantine consensus problem, the following two executions exist:

- e_1: All processes from group C are faulty and silent throughout the entire execution. Moreover, all processes from the $A \cup B$ set propose value v. Since $|C| = t_0$, Π^A ensures that all processes from the $A \cup B$ set eventually decide the same value v (because of the validity property) by some global time t_1.
- e_2: All processes from group A are faulty and silent throughout the entire execution. Moreover, all processes from the $B \cup C$ set propose value $v' \neq v$. Since $|A| = t_0$,

Π^A ensures that all processes from the $B \cup C$ set eventually decide the value $v' \neq v$ (because of the validity property) by some global time t_2.

Now, we can devise another execution e where:

- Processes from group A and processes from group C are correct, whereas processes from group B are faulty. Moreover, all processes from group A propose v, whereas all processes from group C propose $v' \neq v$.
- Processes from group B behave towards processes from group A as in execution e_1 and processes from group B behave towards processes from group C as in e_2.
- All messages between processes from group A and group C are delayed until time $max(t_1, t_2)$.

Execution e is indistinguishable from execution e_1 to processes from group A, which implies that all processes from group A decide value v by time t_1. Similarly, all processes from group C decide value $v' \neq v$ by time t_2.

Finally, we denote by *partitioningExecution* the prefix of execution e up to time $max(t_1, t_2)$ (Part (a) of Figure 1 depicts *partitioningExecution*). Observe that the following holds for *partitioningExecution*:

- All processes from group A decide v in *partitioningExecution*.
- All processes from group C decide $v' \neq v$ in *partitioningExecution*.
- No message is exchanged between any two processes $(a \in A, c \in C)$.

We are now ready to prove the cubic lower bound on communication complexity for solving the accountable Byzantine consensus protocol.

Theorem 3. *The communication complexity of Π^A is $\Omega(n^3)$.*

Proof. The proof is built on top of *partitioningExecution* we constructed above. Namely, *partitioningExecution* is convenient for proving the cubic lower bound since the only way for correct processes (i.e., processes from the $A \cup C$ set) to ensure accountability is to exchange information among themselves. Indeed, faulty processes (i.e., processes from group B) appear correct to all processes from group A (resp., group C). Therefore, no faulty process is detected in *partitioningExecution* because of the fact that no communication is established between groups A and C.

Recall that each correct process needs to obtain a proof of culpability of (at least) $t_0 + 1 = O(n)$ faulty processes. If processes $a \in A$ and $c \in C$ aim to collaboratively obtain a proof of culpability of $t_0 + 1$ processes, both a and c need to send (at least) $t_0 + 1 = \Omega(n)$ authenticators. Moreover, a proof of culpability of $t_0 + 1$ processes must contain (at least) $\Omega(n)$ authenticators.

We now devise a continuation of *partitioningExecution* which ensures that correct processes do send $\Omega(n^3)$ authenticators. We start by stating that there is a single correct process in group A - we denote this process by a. Other processes from group A are Byzantine and they do not send any message to a in the continuation of *partitioningExecution*. All processes from group C are correct and all processes from group B are faulty and silent. Finally, all messages sent between processes from group C that are not received in *partitioningExecution* are delayed.

Let $c_1 \in C$ be a process from group C; recall that c_1 is correct. In the continuation of *partitioningExecution*, c_1 eventually obtains a proof of culpability Σ of $t_0 + 1$ processes by communicating with a *single* process $a_1 \in A$. Specifically, no message is received by c_1 from any process that belongs to group A before c_1 obtains Σ by communicating with a_1. Importantly, c_1 cannot distinguish the current execution from one in which only correct processes are c_1 and a_1 (recall that accountability must be ensured even in the scenario with only two correct processes). However, after the communication with a_1, process c_1 cannot distinguish the current execution from one in which (1) a_1 is faulty (and just behaves correctly towards c_1), and (2) there exist other processes from group A that are correct, disagree with c_1 and need to detect faulty processes. Therefore, c_1 needs to communicate with other processes from group A. The aforementioned construction of the continuation of *partitioningExecution* is repeated for all processes $a_i \in A$: (1) process c_1 sends $\Omega(n)$ authenticators in order to allow process $a_i \in A$ to obtain Σ, and (2) before process a_i obtains Σ, process c_1 does not hear from any other process from group A from which it has not heard yet (i.e., process c_1 communicates with processes from group A in "one-by-one" fashion). Recall that process c_1 does not hear from any process from group C until c_1 has "helped" each process from group A to obtain Σ (i.e., all processes from group C might be faulty as seen from the perspective of c_1). Finally, we conclude that c_1 communicates quadratic number of authenticators in the execution ($\Omega(n)$ authenticators per $t_0 = O(n)$ processes).

We apply the same reasoning for correct process $c_2 \in C$. First, process c_2 hears from any process from group A only after c_1 has already ensured that all processes from group A obtain Σ. Moreover, any message sent by c_2 to a is received after c_1 has already ensured that all processes from group A obtain Σ. All $t_0 - 1$ faulty processes from group A behave towards c_2 as if they receive the information from group C for the first time. Moreover, c_2 communicates with all processes from group A in "one-by-one" fashion. Note that a, the only correct process from group A, might not behave towards c_2 as if it hears the information from group C for the first time. However, process c_2 cannot be certain that neither a nor processes from group C that communicated with a are indeed correct; as seen from the perspective of c_2, all processes from group C that communicated with a might be faulty (recall that c_2 has not heard from other processes from group C in the continuation of *partitioningExecution* thus far), which implies that a might be faulty. Hence, c_2 cannot rely on process a or processes from group C that communicated with a to ensure that all processes from group A obtain a proof of culpability, which results in the fact that c_2 also sends quadratic number of authenticators.

Finally, the construction mechanism we presented for c_2 is repeated for all other processes from group C. Therefore, each process from group C sends quadratic number of authenticators. Since $|C| = t_0 = O(n)$, the total communication complexity of Π^A is $\Omega(n^3)$ (Part (b) of Figure 1

provides a visual depiction of the execution considered in the proof). □

The consequence of theorems 1 and 3 is that $\mathcal{ABC}$ allows any cubic (or sub-cubic) Byzantine consensus protocol (e.g., DBFT binary consensus [10]) to obtain accountability optimally (whenever disagreement occurs).

Corollary 2. Let Π be a Byzantine consensus protocol with the communication complexity $X_\Pi \leq O(n^3)$. Let Π^A be a protocol obtained by applying $\mathcal{ABC}$ to Π. Then, Π^A obtains accountability (whenever disagreement occurs) optimally, i.e., with the cubic communication complexity.

We conclude this subsection by stating the direct consequence of corollaries 1 and 2: $\mathcal{ABC}$, when applied to an optimal (with respect to the communication complexity) Byzantine consensus protocol, constructs a protocol that solves consensus optimally (whenever possible) and obtains accountability optimally (whenever disagreement occurs).

Corollary 3. Let Π_{opt} be a Byzantine consensus protocol with the optimal communication complexity X_{opt}, where $O(n^2) \leq X_{opt} \leq O(n^3)$. Let Π^A_{opt} be a protocol obtained by applying $\mathcal{ABC}$ to Π_{opt}. The following holds for Π^A_{opt}:

1) Π^A_{opt} solves the Byzantine consensus problem with the optimal communication complexity X_{opt}.
2) Π^A_{opt} obtains accountability (whenever disagreement occurs) with the optimal communication complexity $O(n^3)$.

D. Discussion

The (accountable) Byzantine consensus problem (as defined in §I) specifies the validity property which ensures that, if all correct processes propose the same value, then only that value could be decided by a correct process. In the literature, there are many variants of the validity property; the one we use is traditionally called *strong validity*. Throughout the rest of this subsection, we refer to "our" validity property as strong validity. Other most notable variants of the validity property include:

- *Weak Validity:* If all processes are correct and if a correct process decides value v, then v is proposed by a (correct) process [2], [19], [22].
- *External Validity:* A value decided by a correct process satisfies the predefined $valid$ predicate [5].

Importantly, the correctness of $\mathcal{ABC}$ *does not* depend on the specific variant of the validity property.

However, the specific variant of the considered validity property plays a role in showing that our transformation allows for optimal solution to the accountable Byzantine consensus problem. As seen in §III-C, our proof of the cubic lower bound relies on the possibility of devising $partitioningExecution$. Indeed, $partitioningExecution$ could be obtained as a consequence of the strong validity property (see §III-C). Still, if one assumes weak or external validity, there is no guarantee that such execution exist. Thus, the lower bound presented in §III-C *does not* apply to Byzantine consensus problems that do not ensure strong validity, but some other variant of the property.

IV. Generalized $\mathcal{ABC}$ Transformation

We have shown that $\mathcal{ABC}$ enables Byzantine consensus protocols to obtain accountability. This section generalizes our $\mathcal{ABC}$ transformation and defines its applicability. Namely, we specify a class of distributed computing problems named *easily accountable agreement tasks* and we prove that generalized $\mathcal{ABC}$ enables accountability in such tasks.

We introduce agreement tasks in §IV-A. Then, we define the class of easily accountable agreement tasks (§IV-B) and prove the correctness of generalized $\mathcal{ABC}$ transformation applied to such agreement tasks (§IV-C).

A. Agreement Tasks

Agreement tasks represent an abstraction of distributed input-output problems executed in a Byzantine environment. Specifically, each process has its *input value*. We assume that "$\perp$" denotes the special input value of a process that specifies that the input value is non-existent. A process may eventually halt; if a process halts, it produces its *output value*. The "$\perp$" output value of a process means that the process has not yet halted (and produced its output value). We denote by I_i (resp., O_i) the input (resp., output) value of process P_i. We note that some processes might never halt if permitted by the definition of an agreement task (we provide the formal explanation in the rest of the subsection).

An agreement task $\mathcal{A}$ is parameterized with the upper bound $t_\mathcal{A}$ on number of faulty processes that are tolerated. In other words, the specification of an agreement task assumes that no more than $t_\mathcal{A}$ processes are faulty in any execution.

Any agreement task could be defined as a relation between input and output values of processes. Since we assume that processes might fail, we only care about input and output values of correct processes. Hence, an agreement task could be defined as a relation between input and output values of *correct* processes.

An *input configuration* of an agreement task $\mathcal{A}$ is $\nu_I = \{(P_i, I_i) \text{ with } P_i \text{ is correct}\}$, where $|\nu_I| \geq n - t_\mathcal{A}$: an input configuration consists of input values of (all and exclusively) correct processes. Similarly, an *output configuration* of an agreement task is denoted by $\nu_O = \{(P_i, O_i) \text{ with } P_i \text{ is correct}\}$, where $|\nu_O| \geq n - t_\mathcal{A}$: it contains output values of correct processes. We denote by $\theta(\nu_O) = |\{O_i \,|\, (P_i, O_i) \in \nu_O \wedge O_i \neq \perp\}|$ the number of distinct non-$\perp$ values in the ν_O output configuration.

Finally, we define an agreement task $\mathcal{A}$ as tuple $(\mathcal{I}, \mathcal{O}, \Delta, t_\mathcal{A})$, where:

- $\mathcal{I}$ denotes the set of all input configurations of $\mathcal{A}$.
- $\mathcal{O}$ denotes the set of all output configurations of $\mathcal{A}$ such that, for every $\nu_O \in \mathcal{O}$, $\theta(\nu_O) \leq 1$.
- $\Delta : \mathcal{I} \to 2^\mathcal{O}$, where $\nu_O \in \Delta(\nu_I)$ if and only if the output configuration $\nu_O \in \mathcal{O}$ is valid given the input configuration $\nu_I \in \mathcal{I}$.
- $t_\mathcal{A} \leq \lceil n/3 \rceil - 1$ denotes the maximum number of faulty processes the task assumes.

As seen from the definition, correct processes that halt always output the same value in agreement tasks. Moreover, we define agreement tasks to tolerate less than $n/3$ faults.

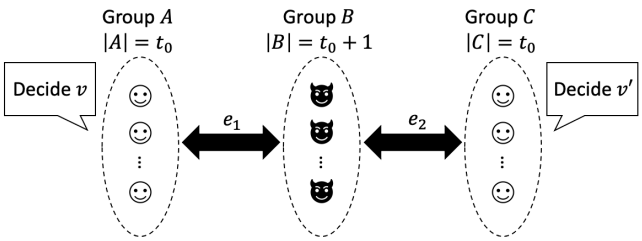

(a) The *partitioningExecution* execution in which processes from groups A and C disagree.

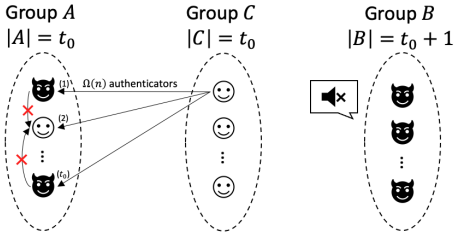

(b) The continuation of *partitioningExecution* that incurs cubic number of authenticators sent by correct processes (the communication pattern of a single process from group C is depicted).

Fig. 1: Illustration of Theorem 3.

Without loss of generality, we assume that $\Delta(\nu_I) \neq \emptyset$, for every input configuration $\nu_I \in \mathcal{I}$. Moreover, for every $\nu_O \in \mathcal{O}$, there exists $\nu_I \in \mathcal{I}$ such that $\nu_O \in \Delta(\nu_I)$.

We note that some problems that are traditionally considered as "agreement" problems do not fall into our classification of agreement tasks. For instance, Byzantine lattice agreement [11] or k-set agreement [7] are *not* agreement tasks per our definition since the number of distinct non-$\perp$ values that can be outputted is greater than 1.

a) Solutions: We say that a distributed protocol $\Pi_{\mathcal{A}}$ *solves* an agreement task $\mathcal{A} = (\mathcal{I}, \mathcal{O}, \Delta, t_{\mathcal{A}})$ if and only if, in every execution with up to $t_{\mathcal{A}}$ faults, there exists (an unknown) time T_D such that $\nu_O \in \Delta(\nu_I)$, where $\nu_I \in \mathcal{I}$ denotes the input configuration that consists of input values of all correct processes and $\nu_O \in \mathcal{O}$ denotes the output configuration that (1) consists of output values (potentially $\perp$) of all correct processes, and (2) no correct process P_i with $O_i = \perp$ updates its output value after T_D.

Finally, we say that a distributed protocol $\Pi^A_{\mathcal{A}}$ solves an *accountable* agreement task $\mathcal{A} = (\mathcal{I}, \mathcal{O}, \Delta, t_{\mathcal{A}})$ if and only if the following holds:

- *$\mathcal{A}$-Solution:* $\Pi^A_{\mathcal{A}}$ solves $\mathcal{A}$.
- *Accountability:* If two correct processes output different values, then every correct process eventually detects at least $t_{\mathcal{A}} + 1$ faulty processes and obtains a proof of culpability of all detected processes.

B. Easily Accountable Agreement Tasks

Fix an agreement task $\mathcal{A} = (\mathcal{I}, \mathcal{O}, \Delta, t_{\mathcal{A}})$. We say that $\mathcal{A}$ is an *easily accountable agreement task* if and only if one of the following conditions is satisfied:

1) *"All-or-None-Decidability":* There does not exist $\nu_O \in \mathcal{O}$ such that $(P_i, O_i \neq \perp) \in \nu_O$ and $(P_j, O_j = \perp) \in \nu_O$; or
2) *"Partial-Decidability":* For every $\nu_I \in \mathcal{I}$ such that there exists $\nu_O \in \Delta(\nu_I)$ where $(P_i, O_i = v \neq \perp) \in \nu_O$ and $(P_j, O_j = \perp) \in \nu_O$, the following holds:

   for every $c \in \mathbb{P}(\{P_i \mid (P_i, I_i) \in \nu_I\}), \nu'_O \in \Delta(\nu_I)$, where $\forall P_i \in c : (P_i, O_i = v) \in \nu'_O$ and $\forall P_j \in \{P_k \mid (P_k, I_k) \in \nu_I\} \setminus c : (P_j, O_j = \perp) \in \nu'_O$.

"All-or-None-Decidability" characterizes all the problems in which either every process halts or none does. For instance, Byzantine consensus [17] and Byzantine reliable broadcast [4] satisfy "All-or-None-Decidability".

On the other hand, some agreement tasks permit that some processes halt, whereas others do not. We say that these tasks satisfy "Partial-Decidability" if and only if it is allowed for *any* subset of correct processes to halt (and output a value). Note that "Partial-Decidability" covers the case in which no correct process ever halts. Byzantine consistent broadcast [4] is the single agreement task we are aware of that satisfies "Partial-Decidability" (in the case with a Byzantine sender). However, the significance of Byzantine consistent broadcast (e.g., for implementing cryptocurrencies [14]) motivated us to consider the "Partial-Decidability" property.

Algorithm 4 Generalized $\mathcal{ABC}$ Transformation - Code For Process P_i

```
1:  Implements:
2:      Accountable Agreement Task A, instance a − A
3:  Uses:
4:      ▷ Protocol to be transformed
5:      Protocol that solves agreement task A, instance Π_A
6:      Accountable Confirmer, instance ac
7:  upon event ⟨a − A, Input | input⟩ do
8:      trigger ⟨Π_A, Input | input⟩
9:  upon event ⟨Π_A, Output | output⟩ do
10:     trigger ⟨ac, Submit | output⟩
11: upon event ⟨ac, Confirm | confirmation⟩ do
12:     trigger ⟨a − A, Output | confirmation⟩
13: upon event ⟨ac, Detect | F, proof⟩ do
14:     trigger ⟨a − A, Detect | F, proof⟩
```

C. Correctness of Generalized $\mathcal{ABC}$ Transformation

We now prove the correctness of our generalized $\mathcal{ABC}$ transformation (Algorithm 4). First, we show that Algorithm 4 solves an easily accountable agreement problem $\mathcal{A}$ if $\mathcal{A}$ satisfies "All-or-None-Decidability".

Lemma 1. Let $\mathcal{A} = (\mathcal{I}, \mathcal{O}, \Delta, t_{\mathcal{A}})$ be an easily accountable agreement task that satisfies "All-or-None-Decidability". Then, Algorithm 4 solves $\mathcal{A}$.

Proof. If no correct process ever outputs a value at line 9, then the lemma trivially holds.

Otherwise, each correct process eventually outputs a value at line 9. Moreover, all correct processes output the exact same value v (since $\mathcal{A}$ is an agreement task). Therefore, all correct processes submit the same value v to accountable confirmer (line 10). By terminating convergence of accountable confirmer, all correct processes eventually confirm value v (line 11) and output it (line 12). Once this happens, the agreement task $\mathcal{A}$ is solved, which concludes the lemma. □

Now, we prove that Algorithm 4 solves an easily accountable agreement task $\mathcal{A}$ if $\mathcal{A}$ satisfies "Partial-Decidability".

Lemma 2. Let $\mathcal{A} = (\mathcal{I}, \mathcal{O}, \Delta, t_{\mathcal{A}})$ be an easily accountable agreement task that satisfies "Partial-Decidability". Then, Algorithm 4 solves $\mathcal{A}$.

Proof. Let ν_I denotes the specific input configuration of $\mathcal{A}$. We consider two cases:

- There does not exist $\nu_O \in \Delta(\nu_I)$ such that $(P_i, O_i \neq \bot) \in \nu_O$ and $(P_j, O_j = \bot) \in \nu_O$: In this case, the proof is identical to the proof of Lemma 1.
- Otherwise: Since $\mathcal{A}$ is an agreement task, we conclude that all processes that output a value at line 9 output the same value v. Therefore, any process that outputs a value at line 12 outputs the value v (ensured by validity of accountable confirmer). Finally, once the system stabilizes at time T_S (the system stabilizes at time T_S if and only if no correct process P_i with $O_i = \bot$ updates its output value after T_S), the fact that any subset of processes could halt and that all halted processes output v implies that Algorithm 4 solves $\mathcal{A}$.

The lemma holds since it is satisfied in all possible cases. □

Finally, we are ready to prove that Algorithm 4 solves an accountable agreement task $\mathcal{A}$, where $\mathcal{A}$ is an easily accountable agreement task, which means that generalized $\mathcal{ABC}$ is correct.

Theorem 4. *Let $\mathcal{A} = (\mathcal{I}, \mathcal{O}, \Delta, t_{\mathcal{A}})$ be an easily accountable agreement task. Then, Algorithm 4 solves the accountable agreement task $\mathcal{A}$.*

Proof. Algorithm 4 satisfies $\mathcal{A}$-solution by lemmas 1 and 2. Furthermore, Algorithm 4 ensures accountability because of the fact that accountable confirmer ensures accountability and $t_{\mathcal{A}} \leq t_0$. Thus, the theorem holds. □

V. Conclusion

We presented $\mathcal{ABC}$, the generic and simple transformation that allows Byzantine consensus protocols to obtain accountability. Besides its simplicity, $\mathcal{ABC}$ is efficient: it suffices for obtaining an accountable Byzantine consensus protocol that is (1) optimal in solving consensus whenever consensus is solvable, and (2) optimal in obtaining accountability whenever disagreement occurs. Finally, we show that $\mathcal{ABC}$ can easily be generalized to other agreement problems (e.g., Byzantine reliable broadcast, Byzantine consistent broadcast). Future work includes (1) designing similarly simple and efficient transformation for problems not covered by our generalized $\mathcal{ABC}$ transformation, like Byzantine lattice and k-set agreement problems, and (2) circumventing the cubic lower bound using randomization techniques.

References

[1] Abraham, I., Jovanovic, P., Maller, M., Meiklejohn, S., Stern, G., and Tomescu, A. Reaching consensus for asynchronous distributed key generation. In *PODC '21: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, July 26-30, 2021* (2021), A. Miller, K. Censor-Hillel, and J. H. Korhonen, Eds., ACM, pp. 363–373.

[2] Buchman, E., Kwon, J., and Milosevic, Z. The latest gossip on bft consensus. *arXiv preprint arXiv:1807.04938* (2018).

[3] Buterin, V., and Griffith, V. Casper the friendly finality gadget. *arXiv preprint arXiv:1710.09437* (2017).

[4] Cachin, C., Guerraoui, R., and Rodrigues, L. *Introduction to reliable and secure distributed programming.* Springer Science & Business Media, 2011.

[5] Cachin, C., Kursawe, K., Petzold, F., and Shoup, V. Secure and efficient asynchronous broadcast protocols. In *Annual International Cryptology Conference* (2001), Springer, pp. 524–541.

[6] Castro, M., Liskov, B., et al. Practical byzantine fault tolerance. In *OSDI* (1999), vol. 99, pp. 173–186.

[7] Chaudhuri, S. More choices allow more faults: Set consensus problems in totally asynchronous systems. *Information and Computation 105*, 1 (1993), 132–158.

[8] Civit, P., Gilbert, S., and Gramoli, V. Brief announcement: Polygraph: Accountable byzantine agreement. In *34th International Symposium on Distributed Computing, DISC 2020, October 12-16, 2020, Virtual Conference* (2020), H. Attiya, Ed., vol. 179 of *LIPIcs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 45:1–45:3.

[9] Civit, P., Gilbert, S., and Gramoli, V. Polygraph: Accountable byzantine agreement. In *2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS)* (2021), pp. 403–413.

[10] Crain, T., Gramoli, V., Larrea, M., and Raynal, M. Dbft: Efficient leaderless byzantine consensus and its application to blockchains. In *2018 IEEE 17th International Symposium on Network Computing and Applications (NCA)* (2018), IEEE, pp. 1–8.

[11] de Souza, L. F., Kuznetsov, P., Rieutord, T., and Tucci Piergiovanni, S. Accountability and reconfiguration: Self-healing lattice agreement. *CoRR abs/2105.04909* (2021).

[12] Dolev, D., and Reischuk, R. Bounds on information exchange for byzantine agreement. *Journal of the ACM (JACM) 32*, 1 (1985), 191–204.

[13] Dwork, C., Lynch, N., and Stockmeyer, L. Consensus in the presence of partial synchrony. *Journal of the ACM (JACM) 35*, 2 (1988), 288–323.

[14] Guerraoui, R., Kuznetsov, P., Monti, M., Pavlovic, M., and Seredinschi, D.-A. At2: asynchronous trustworthy transfers. *arXiv preprint arXiv:1812.10844* (2018).

[15] Haeberlen, A., Kouznetsov, P., and Druschel, P. Peerreview: practical accountability for distributed systems. In *Proceedings of the 21st ACM Symposium on Operating Systems Principles 2007, SOSP 2007, Stevenson, Washington, USA, October 14-17, 2007* (2007), T. C. Bressoud and M. F. Kaashoek, Eds., ACM, pp. 175–188.

[16] Haeberlen, A., and Kuznetsov, P. The fault detection problem. In *Principles of Distributed Systems, 13th International Conference, OPODIS 2009, Nîmes, France, December 15-18, 2009. Proceedings* (2009), T. F. Abdelzaher, M. Raynal, and N. Santoro, Eds., vol. 5923 of *Lecture Notes in Computer Science*, Springer, pp. 99–114.

[17] Lamport, L., Shostak, R., and Pease, M. The byzantine generals problem. In *Concurrency: the Works of Leslie Lamport*. 2019, pp. 203–226.

[18] Libert, B., Joye, M., and Yung, M. Born and raised distributively: Fully distributed non-interactive adaptively-secure threshold signatures with short shares. *Theor. Comput. Sci. 645* (2016), 1–24.

[19] Milosevic, Z., Hutle, M., and Schiper, A. Unifying byzantine consensus algorithms with weak interactive consistency. In *International Conference On Principles Of Distributed Systems* (2009), Springer, pp. 300–314.

[20] Sheng, P., Wang, G., Nayak, K., Kannan, S., and Viswanath, P. BFT protocol forensics. *CoRR abs/2010.06785* (2020).

[21] Sheng, P., Wang, G., Nayak, K., Kannan, S., and Viswanath, P. Bft protocol forensics. In *Computer and Communication Security (CCS)* (Nov 2021).

[22] Yin, M., Malkhi, D., Reiter, M. K., Golan-Gueta, G., and Abraham, I. HotStuff: BFT consensus with linearity and responsiveness. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing* (2019), pp. 347–356.