

# Early Stopping for Any Number of Corruptions

Julian Loss<sup>1,\*</sup> and Jesper Buus Nielsen<sup>2,\*\*</sup>

<sup>1</sup> CISA Helmholtz Center for Information Security, Germany

<sup>2</sup> Aarhus University, Denmark

**Abstract.** Minimizing the round complexity of byzantine broadcast is a fundamental question in distributed computing and cryptography. In this work, we present the first *early stopping* byzantine broadcast protocol that tolerates up to  $t = n - 1$  malicious corruptions and terminates in  $\mathcal{O}(\min\{f^2, t + 1\})$  rounds for any execution with  $f \leq t$  *actual corruptions*. Our protocol is deterministic, adaptively secure, and works assuming a plain public key infrastructure. Prior early-stopping protocols all either require honest majority or tolerate only up to  $t = (1 - \epsilon)n$  malicious corruptions while requiring either trusted setup or strong number theoretic hardness assumptions. As our key contribution, we show a novel tool called a *polariser* that allows us to transfer certificate-based strategies from the honest majority setting to settings with a dishonest majority.

## 1 Introduction

In the problem of byzantine broadcast [21], a sender  $P_s$  holds a value  $v$  that it wants to share among  $n$  parties  $P_1, \dots, P_n$  using a distributed protocol  $\Pi$  with the following properties: 1) *validity*: if the sender is honest (i.e., it follows the protocol description of  $\Pi$  correctly), all honest parties output  $v$  2) *agreement*: all honest parties output the same value  $v'$  from  $\Pi$ . Broadcast is an integral building block in many cryptographic and distributed protocols, e.g., multi-party computation, verifiable secret sharing, and state-machine replication. One of the most important efficiency metrics for a broadcast protocol is its *round complexity*: how many rounds does the protocol run for until all parties have terminated?

A seminal result of Dolev and Strong [12] shows that any broadcast protocol tolerating  $t < n$  malicious parties runs for at least  $t + 1$  rounds in some runs. In the same work, they also give a protocol that shows the tightness of their bound. However, their bound is known to be loose for protocol executions where the number of corruptions  $f$  is less than  $t$ , i.e.,  $f < t$ , and where eventual agreement is allowed. In [11] the authors define two notions of agreement, simultaneous agreement (SA) and eventual agreement (EA). In SA the parties must give output in the same round. In EA they are allowed to give output in different rounds, but must of course still agree on the output itself. They show that for SA any protocol will have runs with  $t + 1$  rounds even when  $f = 0$ . But for EA they only show a lower bound of  $f + 2$  rounds.<sup>3</sup> They did not give a protocol matching this bound. Intrigued by this gap, a long line of work has studied so-called *early stopping* protocols which terminate

---

\* This work is funded by the European Union, ERC-2023-STG, Project ID: 101116713. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

\*\* Funded by the Danish Independent Research Council under Grant-ID DFF-3103-00077B (CryptoDigi).

<sup>3</sup> The proof uses a hybrid argument appealing to agreement  $e^{\text{poly}(t)}$  times, so holds only for protocols with agreement error  $e^{-\text{poly}(t)}$ . The argument appeals to validity only twice so the validity error may be any constant  $< 1/2$ .

in  $O(g(f)) = o(t)$  rounds (for some function  $g$ ) in any execution where the actual number of corrupted parties  $f$  is sufficiently small compared to  $t$ . Early stopping protocols are known for both the information theoretic setting with  $t < n/3$  malicious corruptions [6, 17, 2, 22] and the authenticated setting with  $t < n/2$  corruptions [27]. To the best of our knowledge, however, little is known about early stopping protocols for the setting of  $n/2 \leq t < n$  malicious corruptions. On one hand, several randomized protocols achieve sublinear (in  $n$ ) round complexity for broadcast in the dishonest majority setting [16, 15, 7, 30, 29]. However, these protocols require the maximum number  $t$  of corruptions to be at most a constant fraction of  $n$  in order to stop early (some require  $t$  to be much smaller).

All in all, for the full corruption threshold  $t = n - 1$ , the tightest lower bound says that any protocol must run for at least  $f + 2$  rounds, whereas the best protocol uses  $n$  rounds in all runs, even for low  $f$ . Clearly, there is a fundamental gap in our understanding of early stopping protocols when  $t = n - 1$ . Motivated by this discussion, we pose the following question: *Are there early stopping broadcast protocols tolerating up to  $t = n - 1$  corruptions?*

## 1.1 Our Contribution

In this work, we answer this question in the affirmative by providing the first early stopping protocol CDC for arbitrary majority corruption, i.e.,  $t = n - 1$ . Concretely, its properties can be summarized as follows:

- CDC tolerates any number  $t < n$  of malicious corruptions.
- For any execution with  $f \leq \sqrt{t}$  faults, CDC terminates in  $\mathcal{O}(f^2)$  rounds. It always terminates within  $\mathcal{O}(t)$  rounds. Prior work achieves either  $t + 1$  rounds deterministically or requires that the maximum number of faults satisfy  $t = (1 - \epsilon) \cdot n$  for some  $\epsilon > 0$  in order to achieve early stopping.
- Our protocol is secure with respect to a *strongly adaptive adversary*. This type of adversary can observe an honest party’s messages, corrupt it, and replace its messages with its own before these messages are delivered. This sets our work apart from existing early stopping protocols for the dishonest majority regime with a strongly adaptive adversary. Namely, existing protocols require either 1) strong number theoretic hardness assumptions (i.e., time-locked puzzles) [29, 28] or 2) tolerate only  $t = n/2 + O(1)$  malicious corruptions [16, 15] in order to stop early.
- CDC is deterministic and works in the plain public key model. By comparison, existing early stopping protocols for majority corruption are all randomized and, in many cases, require strong setup assumptions.

In summary, CDC is the first early stopping protocol for  $t = n - 1$  and makes a significant improvement over the state of the art for any execution with  $f = o(\sqrt{n})$  corruptions. We give further comparison with existing literature in the related work section.

## 1.2 Technical Overview

We now give an overview of our techniques. We begin by giving a brief recap of the classical Dolev-Strong protocol [12] DSC. We then explain the difficulty of making this protocol early stopping and our key insights to overcome it.

*Recap: the Dolev-Strong protocol.* DSC achieves a round complexity of  $t+1$  against a strongly adaptive malicious adversary corrupting up to  $t < n$  parties. It works as follows for a sender  $P_s$  holding a message  $m$ :

- In round 1,  $P_s$  signs  $m$  and sends it to all parties.
- In any round  $i \leq t$ , a party  $P_j$  does as follows. If it receives a message  $m$  together with a list of valid signatures from  $i$  distinct parties for the first time, it adds  $m$  to a set of accepted messages  $\mathcal{A}$ . Then,  $P_j$  adds its own signature to the list, and forwards it to all parties so that they add it to  $\mathcal{A}$  at most one round later.
- In round  $t+1$ , a party  $P_i$  executes the above rule to update  $\mathcal{A}$ , but does not forward a new list. Instead  $P_i$  determines its output as follows. If  $\mathcal{A} = \{\}$  or  $|\mathcal{A}| > 1$ , output a default output `NoMsg`. If  $\mathcal{A} = \{m\}$ , output  $m$ .

Clearly, if an honest  $P_i$  adds  $m$  to  $\mathcal{A}$  at any point during the first  $t$  rounds of the protocol, all honest parties add it to  $\mathcal{A}$  at most one round later. If  $P_i$  adds  $m$  to  $\mathcal{A}$  in round  $t+1$ , it sees  $m$  with  $t+1$  signatures and thus knows that at least one honest party  $P_j$  has previously signed  $m$  in a previous round. Since  $P_j$  was honest in that round, it would have added its own signature to the list of  $t$  signatures it needed to add  $m$  to  $\mathcal{A}$  and passed on the resulting list of  $t+1$  signatures to all parties. Hence, all honest parties add  $m$  to  $\mathcal{A}$  by round  $t+1$ .

*Why making DSC early stopping is hard.* Stopping DSC early turns out to be very challenging. Surprisingly, however, this does not result from a fully malicious sender  $P_s$  that sends conflicting (signed) messages in DSC to break consistency, as this allows all honest parties to detect and prove that  $P_s$  is acting dishonestly. The central difficulty arises already from crash faults:  $P_s$  can simply not send any message to (some of) the parties. Note that if  $P_s$  sends no signature then DSC runs in silence for  $t+1$  rounds before outputting `NoMsg`. We would like to terminate this earlier. However, the sender might send a signature only to a single honest party in round 1. Or send it to no honest party but collude with some  $P_i$  forwarding the signature from  $P_s$  to a single honest party in round 2, *et cetera*. To detect this and stop early it would help if honest parties from round 1 could prove that they did not get a message from  $P_s$ . But how to prove this? This is a comparatively simple task in the honest majority setting where there are  $n-t \geq t+1$  honest parties. At a high-level, parties can simply collect a certificate of accusations against the sender for not sending them a message. If they can obtain  $t+1$  signed accusations, they can disqualify the sender and stop the protocol. On the other hand, if  $P_s$  cannot be disqualified, then at least one honest party must have received a message from  $P_s$  and can forward it to all other parties. Unfortunately, this strategy fails when  $t \geq n/2$ , as there are only  $n-t < t+1$  honest parties so a certificate might not be constructible. Thus, the obstacle we must overcome is to design an analogue of a certificate for the dishonest majority setting.

*Polarisers to the rescue.* Our key tool for achieving this is a novel primitive called a *polariser*. Informally, a polariser partitions the parties into two ‘polarized’ sets `Alive` and `Corrupt`. Polarisers are updated continuously throughout the protocol and maintain the following properties. First, an honest party  $P_i$  accepts a polariser from another party (and subsequently

updates its own polariser) only if it itself is in *Alive*. Moreover, for any party  $P \in \text{Corrupt}$ , a polariser contains accusation against  $P$  from *all* parties in *Alive*. And it is ensured that honest parties never accuse honest parties. As it turns out, these properties make it possible for an honest party  $P_i$  to justify any decision in our protocol and convince other honest parties to take the same decision. A crucial observation is that it follows from the properties that all honest parties are in *Alive* and therefore  $P_i$  knows it can forward the polariser and have it accepted by other *honest* parties: they too are in *Alive*. Thus, polarisers act as our replacement for certificates. The idea behind creating a polariser is surprisingly simple. As an example, suppose the sender  $P_s$  does not send a party  $P_i$  a message in the first round. Now,  $P_i$  can publicly accuse  $P_s$ . To deal with having accusations levelled against honest senders, note that honest parties will move  $P_s$  to *Corrupt* only if it was accused by all parties in *Alive*. Since  $P_i$  never accuses an honest sender  $P_s$  and is itself in *Alive*, this precludes honest parties from moving an honest  $P_s$  to *Corrupt*. However, this creates a different problem: if  $P_j$  is also dishonest, but is itself in *Alive*, it can simply not accuse  $P_s$ . In this case,  $P_s$  cannot be moved from *Alive* to *Corrupt*, since not all parties in *Alive* have accused  $P_s$ . To deal with this type of behaviour from dishonest parties, we present a recursive solution for generating a polariser. In our running example from above, we require that either  $P_j$  sends whatever it got from  $P_s$  or it accuses  $P_s$  of cheating. If  $P_i$  receives neither of those two things from  $P_j$ ,  $P_i$  knows that  $P_j$  is itself corrupt and accordingly accuses  $P_j$ . And, importantly,  $P_i$  expects every other party  $P_k$  to accuse  $P_j$  too, or send to  $P_i$  the reason that it did not accuse  $P_j$ , namely an accusation of  $P_j$  against  $P_s$  (which would resolve the issue). If  $P_k$  does not send an accusation or resolvment, then  $P_i$  will accuse  $P_k$ , *et cetera*. In each recursive step one more corrupt party is accused, and there are at most  $f$  corrupt parties, so the recursion stops in at most  $f$  steps. When it stops, then in the view of every honest party either  $P_s$  can be moved from *Alive* to *Corrupt* or a signature from  $P_s$  was received. Hence each honest party receives either a signed message from  $P_s$  or a polariser showing that  $P_s$  is corrupt.

*Justifying outputs and graded broadcast.* Polarisers can be used to justify the output (or non-output) of any subprotocol to other parties. To do so, parties will either send their entire view of the protocol transcript so far in case a subprotocol produces output, or send a polariser to justify not having output. We show a protocol **GSTM** for *graded broadcast* that is inspired by the protocol of Koo et al. [16]. Roughly, **GSTM** outputs a message  $m$  together with a grade  $g \in \{0, 1, 2\}$  and a proof  $\pi$  that justifies the combination of  $m$  and  $g$ . The grade  $g$  reflects a party's confidence in its output. Grade  $g = 2$  indicates high confidence, meaning that  $m$  should be output, and all other parties have grade at least 1 for the same message  $m$ . Grade  $g = 1$  indicates low confidence in  $m$ , meaning that some parties might not have received the message  $m$ . Finally,  $g = 0$  indicates that the message was not received, the sender was corrupt, and some dummy `NOMSG` should be output. Validity ensures that all parties output  $(m, 2)$  whenever the sender is honest. The crucial property of **GSTM**, however, is that if a dishonest party can produce a justified output then the grading rules from above apply too. So a corrupted party basically cannot justify an output unless that output could have been produced by an honest party.

*Putting things together: Protocol DC.* Using **GSTM** we are able to run a broadcast protocol **DC** that resembles the well-known phase-king approach [5] from the honest majority setting. In this style of protocol, one rotates through  $f + 1$  leaders until agreement on an output is detected. Essentially, each leader is instructed to broadcast via **GSTM** whatever it received from the previous instance. The crucial idea is that a malicious leader can never undo the progress that the protocol has made so far by making them choose a different output from one that they have already agreed on (or not choosing the output of an honest sender). This is because each time a new broadcast is initiated by a leader, its input must be justified by the entire view of the protocol so far. Therefore it must use an input which some honest party could have used, or choose to abort the protocol. More precisely, once an output  $m$  is received with a positive grade  $g$  in the  $j$ th leader iteration, parties set  $m_j = m$  and broadcast this message once it is their turn to be the leader. (If the previous king aborted then they pick the most recent such value, if there is one, and otherwise they pick a fixed default output). The consistency properties of **GSTM** and the justification that comes along with any new output ensures that a dishonest leader can never introduce a new message once parties have already agreed on a message  $m$ . Once parties see an output  $m$  with grade 2, they can detect agreement, forward the justification for this output to all parties, and terminate. After  $f + 1$  leader rotations, at least one of the leaders will be honest, giving grade 2, at which point the agreement detection is triggered and the protocol terminates.

*CDC: Ensuring  $O(t)$  round-complexity.* In the worst case, each of the leader iterations in **DC** identifies only a single party (i.e., the leader) as malicious. Since each of these steps takes roughly  $f$  rounds, we end up with an overall complexity of  $O(f^2)$  for **DC**. To avoid exceeding  $O(t)$  for the round complexity, we can simply stop the protocol after running for  $\ell = O(t)$  rounds. At this point, we can afford to run an additional instance of **DSC** to reach agreement. More precisely, any party that has not terminated by round  $\ell - 1$  will thus use **DSC** to broadcast its view of the protocol so far to all parties after completing iteration  $\ell$ . (Running for one more iteration ensures that all parties have time to be forwarded justifications from already terminated parties). Once **DSC** terminates after another  $O(t)$  rounds, all parties can locally decide on a correct output which, by the justification properties of **GSTM**, will be consistent with parties who have already terminated. This solution, however, still has an undesirable property: honest parties can terminate  $O(t)$  rounds apart, whenever one party terminates in iteration  $\ell - 1$ , but other parties keep on running. This would make the protocol very difficult to use as a subroutine in higher-level application. To have parties terminate one round apart, our actual protocol **CDC** replaces **DSC** with a new protocol **WES** (**WES** stands for *weak early stopping*) that terminates in  $O(f)$  rounds if the sender is honest and  $O(t)$  rounds in any other case—furthermore, parties terminate at most one round apart. This allows honest parties who terminate early in **DC** to always broadcast their justified outputs via **WES**. If any **WES** reports a justified output of **DC** that will be the final output, otherwise the final output will be **NO MSG**. Parties can terminate immediately upon receiving output from the first terminating instance of **WES** reporting a justified output of **DC**. If an honest party saw a justified output of **DC** it gets reported in  $O(f)$  rounds and the overall protocol terminates in  $O(f^2)$  rounds. If no honest party saw a justified output from **DC** in

$O(t)$  rounds, there are (at least)  $f = \sqrt{t}$  corruptions and then the protocol is allowed to run for  $O(t)$  rounds. We believe that WES may have other natural applications.

*Achieving Polynomial Complexity.* A final wrinkle is that sending along the protocol’s entire transcript in every step would result in exponential communication complexity. However, this turns out to be unnecessary, as  $P_i$  can simulate another party  $P_j$ ’s behaviour from its past messages. This means that  $P_j$  need only send information associated with the most recent protocol step every time it is asked to justify one of its outputs. Thus, our combined protocol CDC ends up with a communication complexity of  $O(n^4\lambda)$ , where  $\lambda$  is the length of a signature.

### 1.3 Related Work

Below, we summarize some early stopping protocols from the literature.

*Deterministic Protocols.* To the best of our knowledge, the first early stopping protocol for byzantine agreement (which implies broadcast for  $t < n/2$ ) with optimal resilience  $t < n/3$  in the information theoretic setting was due to Berman et al. [6], who gave a protocol with (optimal) round complexity  $\min\{f + 2, t + 1\}$  and exponential communication. Their work builds on earlier work of Berman and Garay [4] who achieved the same round complexity with polynomial complexity and  $n > 4t$ . Garay and Moses [17, 18] later improved the communication and computation for the corruption-optimal protocol to polynomial. However, their protocol achieves a slightly worse round complexity for the early stopping case of  $\min\{f + 5, t + 1\}$ . Much later, Abraham and Dolev [2] gave the first protocol with optimal round complexity  $\min\{f + 2, t + 1\}$  and polynomial communication/computation. In the authenticated setting with  $t < n/2$  and plain PKI, Perry and Toueg [27] showed a protocol with polynomial communication and computation complexity and a round complexity of  $\min\{2f + 4, 2t + 2\}$ .

*Randomized Protocols.* In the following, we let  $\delta$  denote the failure probability of a protocol. There are various randomized protocols for the honest majority setting with constant expected round complexity for both the  $t < n/3$  (information theoretic) setting [14, 24] and  $t < n/2$  (authenticated) setting [19, 1]. They can all be made to terminate early with worst-case failure probability  $\delta$  by running them for  $O(\log(1/\delta))$  iterations (each iteration has constant many rounds).

One can use similar ideas to turn expected constant round protocols in the dishonest majority setting with  $t < n$  corruptions into protocols that always stop early and have failure probability  $\delta$ . Here, randomized protocols were first explored by Garay et al., who showed an expected  $O(2t - n)^2$ -round protocol from plain PKI for any  $t < n$  [16]. Their approach was later improved by Fitzi and Nielsen [15], who showed a protocol with  $O(2t - n)$  complexity in the same setting. These protocols lead to early stopping protocols with round complexities  $O(\log(1/\delta) + (2t - n)^2)$  and  $O(\log(1/\delta) + 2t - n)$ , respectively, and failure probability  $\delta$ . However, since  $O(2t - n) = O(n)$  (regardless of  $f$ ) whenever  $t = (1 - \epsilon) \cdot n$  where  $\epsilon > 0$ , these protocols are not early stopping.

Chan et al. [7] presented a randomized broadcast protocol with trusted setup and tolerating any  $(1 - \epsilon)$ -fraction of adaptive corruptions (for an arbitrary constant  $\epsilon > 0$ ) by assuming no after-the-fact removal of messages. Their protocol achieves a round complexity of  $O(\log(1/\delta)) \cdot (n/(n-t))$  for failure probability  $\delta$ . Also assuming trusted setup and no after-the-fact removals, but tolerating up to  $t = n - 1$  corruptions, Wan et al. [30] give a protocol achieving expected  $O((n/(n-t))^2)$  round complexity and  $O(\log(1/\delta)/\log(n/t)) \cdot (n/(n-t))$  complexity for failure probability  $\delta$ . Protocols tolerating adaptive corruptions with after-the-fact message removals and  $t < (1 - \epsilon)n$  corruptions were studied by Wan et al. [30] and more recently by Srinivasan et al. [28], who gave protocols from time-lock-assumptions achieving round complexities of  $(n/(n-t))^2 \cdot \text{polylog}(\lambda)$  (for failure probability negligible in  $\lambda$ ) and  $O(\log(1/\delta)) \cdot (n/(n-t))$ , respectively. Most recently, Alexandru et al. [3] showed how to remove the need for trusted setup in order to obtain  $O(\log(1/\delta)) \cdot (n/(n-t))$  round complexity. Although these protocols all achieve early stopping (with failure probability  $\delta$ ) for  $t = (1 - \epsilon) \cdot n$ , they are also not early stopping when  $t = n - O(1)$ . Namely, in this case, their round complexity is at least  $O(n/(n-t)) = O(n)$ .

*Lower Bounds, communication optimizations, and weaker models.* As mentioned above [11] showed that the round complexity of an early stopping algorithm with eventual agreement in an execution with  $f$  faults is lower bounded by  $\min\{f + 2, t + 1\}$ . This was later extended by Keidar and Rajsbaum [20] to the setting of omission faults, which can fail to send or receive some of their messages. They demonstrate that early stopping broadcast/agreement algorithms require the same complexity as in the malicious setting if agreement is required to be *uniform*, i.e., omission faulty parties that output, must output consistently with the honest parties. Chandra et al. present reliable broadcast protocols achieving  $f + 2$  rounds in the crash fault model and  $2f + 3$  rounds in the omission fault model [8]. The latter result was later improved to  $f + 2$  by Parvédy and Raynal [26] to  $\min\{f + 2, t + 1\}$  round complexity and  $O(n^2 \cdot f)$  communication complexity in the omission fault model.

From the perspective of communication complexity, a result by Dolev and Lenzen [9] shows that any (deterministic) early stopping algorithm with optimal round complexity requires sending  $O(nt + t^2f)$  messages. This tightens the famous Dolev-Reischuk bound for the early stopping case [10]. On the other hand, the recent result of Lenzen and Sheikholeslami [22] demonstrate that this bound can be circumvented by presenting an early stopping protocol with (suboptimal)  $O(f)$  round complexity but significantly improved  $O(nt)$  communication complexity.

## 2 Preliminaries

We work with directed trees  $\text{Tree}$  with a single root and edges pointing towards the leafs. For a tree  $\text{Tree}$  we use  $\text{path} \in \text{Tree}$  to denote a path  $(r, \dots, l)$  from the root  $r$  to a leaf  $l$ . We let  $\text{depth}(\text{Tree}) = \max_{\text{path} \in \text{Tree}} (|\text{path}| - 1)$ . The tree with only a root thus has  $\text{depth}(\text{Tree}) = 0$  and if the tree is empty  $\text{depth}(\text{Tree}) = -1$ . For  $\text{path} = (r, \dots, l)$  we let  $\text{leaf}(\text{path}) = l$ . We assume a synchronous model with  $n$  parties  $\mathbb{P} = \{\mathbb{P}_1, \dots, \mathbb{P}_n\}$ . The computation proceeds in rounds where in each round each party can send a message to each other parties that

is guaranteed to arrive by the end of that round. We assume a rushing adversary that can adaptively corrupt parties and replace or delete any of the messages they sent for a round and which have not yet been delivered. We use  $t$  to denote the maximum number of corrupted parties and  $f \leq t$  to denote the actual number of corrupted parties. We allow  $t$  to take any value  $0 \leq t \leq n$ .

We assume a PKI. In an initial setup round each party  $P_i$  generates a key-pair  $(sk_i, vk_i)$  for a signature scheme and announces  $vk_i$  to a public bulletin board. As is standard for this line of work, we assume the Dolev-Yao model [13] and treat signatures as information theoretic objects with perfect unforgeability. Throughout, we denote the size of a signature in bits as  $\lambda$ .

**Definition 1 (Broadcast).** *Let  $\Pi$  be a protocol executed by  $n$  parties  $P_1, \dots, P_n$ , where a designated sender  $P_s$  holds input  $x$  and each party  $P_i$  terminates upon giving output  $y_i$ . We say that  $\Pi$  is a broadcast protocol if it has these properties:*

- **Validity:** *If  $P_s$  is honest, each honest party  $P_i$  outputs  $y_i = x$ .*
- **Agreement:** *For honest parties  $P_i$  and  $P_j$ ,  $y_j = y_i$ .*

*Sequential composition of protocols without simultaneous termination.* When describing our protocols we will assume that all parties get input in the same round. If a party has no input in a protocol we assume that they nonetheless get a tacit, dummy input, to ensure they know in which round to start running. We also assume that all sub-protocols give outputs in the same round. This ensures that if the parties call a sub-protocol and then proceed when it gives output, then they are still synchronized. Under these conditions we will design protocol where parties might terminate at most one round apart. This leads to problems with composition: when using a protocol as a subroutine, we assume parties give outputs in the same round. But in many of our protocols, parties terminate one round apart. This can, however, be handled using known techniques for sequential composition of protocols without simultaneous termination at a blowup in round complexity of just 2. Details can be found in [23] and Chapter 7 in [25]. Here, we sketch the main idea for completeness. Protocols which assume that the parties start in the same round will be compiled into protocols tolerating that they start one round apart. The compiler works as follows. If parties start a protocol  $\Pi$  one communication round apart, then after  $P_i$  sends its messages for protocol round  $r$  of  $\Pi$ , it will wait for two underlying communication rounds to ensure it received messages from honest parties sending their messages one underlying communication round later than  $P_i$ . Then,  $P_i$  computes its messages for the next protocol round and sends them out, waits for two communication rounds *et cetera*. This leaves the problem that the parties might terminate two communication rounds apart. This would be a problem for sequential composition as we want them to start the next protocol only one communication round apart. This can be mitigated when  $\Pi$  has justified outputs. When the first party gets an output it sends it to the other parties. The output will arrive within 1 *communication* round. When a party sees a valid incoming protocol output, they adopt this as their own output and forward it to all parties. Now all parties terminate at most one communication round apart. We will call this the *staggering compiler*.



### 3 Polarisers and Transferable Justifiers

We first put in place a tool which will allow us to write later protocols more concisely. The tool is called a *polariser* as it polarises the  $n$  parties into two disjoint sets  $S$  and  $T$  of which we know all honest parties are fully inside one of the sets. An external party might not know whether  $S$  or  $T$  contain the honest parties, but an honest party will of course know which set it is in.

Polarisers are used for proving that a corrupted party  $P_i$  did *not* send a message. To motivate their design we first discuss this issue. Consider a party  $P_i$  which is to send a message  $m$  of a particular form to  $P_k$ , say it should be signed. We would like to know when this was *not* done and have a *transferable* proof of this. If we have a bound  $t < n/2$  on how many corruptions there can be then this is easy. You can ask  $P_i$  to send  $m$  via all other parties  $P_j$  and have all  $P_j$  forward  $m$  to  $P_k$  or a signature  $\sigma_j = \text{Sig}_{\text{sk}_j}(\text{ACC}, P_i, P_j)$  which is a signed accusation of  $P_i$  that  $P_j$  did not send a message. Now  $P_k$  either gets  $m$  or a certificate of  $t + 1$  accusations which can act as a transferable proof that  $P_j$  did not send a message. In contrast, in the dishonest majority setting, simple majority voting about whether the message was sent will not solve the problem.

The solution is polarisers. The core of a polariser will be a tuple  $(\text{Alive}, \text{Corrupt}, \text{Accuse})$ , where  $\mathbb{P} = \text{Alive} \cup \text{Corrupt}$  and for all parties in  $\text{Alive}$ , we have a signed accusation for each party in  $\text{Corrupt}$ . Assuming that all honest parties send all intended messages and honest parties only accuse parties which fail to send a message, this leaves only two cases when seeing  $(\text{Alive}, \text{Corrupt}, \text{Accuse})$ . Either all the honest parties are in  $\text{Alive}$  or all the honest parties are in  $\text{Corrupt}$  (if there is an honest party in both  $\text{Alive}$  and  $\text{Corrupt}$  then an honest party accused an honest party). But it might of course be that all the parties in  $\text{Alive}$  are corrupted and falsely accusing the parties in  $\text{Corrupt}$ . This is hard to catch in the dishonest majority case where it can happen that  $|\text{Alive}| > |\text{Corrupt}|$  and all parties in  $\text{Alive}$  are corrupt. This prevents external agreement on who is corrupt.

The trick is to give up on externally valid certificates and go for a weaker type of certificate which maintains transferability only within the context of the current protocol. An honest party  $P_i$  can check whether  $P_i \in \text{Alive}$  or  $P_i \in \text{Corrupt}$ . If  $P_i \in \text{Corrupt}$ , then reject the polariser. Note that in this case *all* honest parties are in  $\text{Corrupt}$  and will therefore also reject the polariser if it is sent to them. If  $P_i \in \text{Alive}$ , then accept the polariser. Note that in this case *all* honest parties are in  $\text{Alive}$  and will therefore also accept the polariser.

**Definition 2 (Polariser).** Let  $\text{Pol} = (\text{Alive}, \text{Corrupt}, \text{Accuse})$  be a tuple where we refer to set  $\text{Alive}$  as the alive parties, to set  $\text{Corrupt}$  as the corrupt parties, and to set  $\text{Accuse}$  as the accusations. We define the following structural properties:

- **Justifiability.** For every  $P_j \in \text{Corrupt}$  and for every party  $P_i \in \text{Alive}$ , there exists  $A_{i,j} \in \text{Accuse}$ , where  $A_{i,j}$  denotes a value  $(\text{ACC}, P_i, P_j, \sigma_i)$ , where  $\text{Ver}_{\text{vk}_i}((\text{ACC}, P_i, P_j), \sigma_i) = \top$ .
- **Completeness.**  $\text{Alive} \cap \text{Corrupt} = \emptyset$  and  $\text{Alive} \cup \text{Corrupt} = \mathbb{P}$ .

We define the following contextual property:

- **Accusation Soundness.** If  $P_i$  and  $P_j$  are honest then the adversary cannot construct a valid  $A_{i,j}$  in PPT, in particular there is no such  $A_{i,j}$  in  $\text{Accuse}$ .

We call a polariser  $\text{Pol}$  a  $P_i$ -polariser if  $P_i \in \text{Pol.Corrpt}$ . We use  $\overset{i}{\not\rightarrow}$  to denote the set of  $P_i$ -polarisers. As with the Landau notation for asymptotic complexity we misuse notation and use  $\text{Pol} = \overset{i}{\not\rightarrow}$  to denote that  $\text{Pol} \in \overset{i}{\not\rightarrow}$ . We also sometimes let  $\overset{i}{\not\rightarrow}$  denote a generic element from  $\overset{i}{\not\rightarrow}$ .

In all our protocols constructing polarisers we only sign messages of the form  $(\text{ACC}, P_i, P_j, \sigma_i)$  if  $P_j$  is corrupt. Therefore:

**Lemma 1 (Polarisation Lemma).** *Let  $\text{Honest}$  be the set of honest parties and let  $\text{Pol}$  be a polariser. Then  $\text{Honest} \subset \text{Pol.Alive}$  or  $\text{Honest} \subset \text{Pol.Corrpt}$ .*

*Proof.* By Completeness,  $\text{Honest} \subset \text{Pol.Alive} \cup \text{Pol.Corrpt}$ , and it cannot be the case that  $P_i \in \text{Pol.Alive}$  and  $P_j \in \text{Pol.Corrpt}$  are honest, because then by Justifiability  $A_{i,j} \in \text{Accuse}$ , contradicting Accusation Soundness.  $\square$

### 3.1 Transferable Justifiers

Our second general tool is the concept of a transferable justifier for a protocol output. Recall that the purpose of polarisers is to get a transferable proof that some party did not send a message. From these, we build increasingly complex messages and eventually a justified output. It is helpful to have some general machinery for talking about transferable justifiers.

**Definition 3 (Justifier).** *We call a PPT predicate  $J : \mathbb{P} \times \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{\top, \perp\}$  a justifier predicate. If for a party  $P_i$ , a message  $m$  and a proof  $\pi$  we have that  $J(P_i, m, \pi) = \top$  then we say that  $P_i$  accepts the message  $m$  with justifier  $\pi$ . We require that justifiers are transferable, i.e., if  $P_i$  and  $P_j$  are honest then  $J(P_i, m, \pi) = \top$  implies that  $J(P_j, m, \pi) = \top$ . We use  $J(m, \pi) = \top$  to denote that  $J(P_i, m, \pi) = \top$  for all honest  $P_i$ . By transferability this is the implied if it holds for a single honest  $P_i$ .*

As an example consider a protocol where  $P_j$  was to send a message and let  $\text{NoMsg}^{(j)}$  be a special symbol denoting that  $P_j$  sent no message. Then a justifier predicate for this could be  $J(P_i, \text{NoMsg}^{(j)}, \text{Pol}) \equiv P_i \in \text{Pol.Honest} \wedge \text{Pol} = \overset{j}{\not\rightarrow}$ , i.e.,  $P_i$  accepts that  $P_j$  sent nothing if  $\text{Pol}$  proves that  $P_j$  is corrupt and that  $P_i$  is honest. This is a transferable justification qua Lemma 1.

**Definition 4 (Justified Inputs/Outputs).** *We say that a protocol  $\Pi$  has justified inputs if it takes an input justifier  $J_{\text{In}}$  as parameter and works for any justifier predicate  $J_{\text{In}}$ . We write  $\Pi_{J_{\text{In}}}$  to specify the value of  $J_{\text{In}}$  being used in a given run. When a protocol  $\Pi_{J_{\text{In}}}$  with justified inputs is being called by an honest party  $P_i$  with input  $x_i$  then  $x_i$  must be of the form  $x_i = (m_i, \pi_i)$  such that  $J_{\text{In}}(P_i, m_i, \pi_i) = \top$ . We say that a protocol  $\Pi$  has an output justifier if the protocol, as part of its code, specifies a justifier predicate  $J_{\text{Out}}$ . We denote the output justifier of  $\Pi$  by  $\Pi.J_{\text{Out}}$ . We say that a protocol  $P_i$  has justified outputs if it has an output justifier and the outputs  $y_i$  of honest  $P_i$  are of the form  $y_i = (m_i, \pi_i)$  and it always holds that  $\Pi.J_{\text{Out}}(P_i, m_i, \pi_i) = \top$  after a run of the protocol with a PPT adversary.*

An important tool in our protocols is that justified outputs can be passed on to other parties. Therefore, not even adversarial parties should be able to claim wrong outputs. The following notion is later used to phrase this.

**Definition 5 (Adversarial Justified Output (AJO)).** *Let  $\Pi$  be a protocol with an output justifier and let  $\mathcal{A}$  be a PPT adversary. Consider the following experiment: Execute  $\Pi$  with  $\mathcal{A}$  in the role of the adversary. When all honest parties  $P_i$  have produced an output  $y_i = (m_i, \pi_i)$ , give all  $y_i$  to  $\mathcal{A}$ . We say that  $\mathcal{A}$  generates  $\ell$  adversarial justified outputs (AJOs)  $(m^1, \dots, m^\ell)$  if it outputs  $(P^1, m^1, \pi^1), \dots, (P^\ell, m^\ell, \pi^\ell)$  such that for all  $j = 1, \dots, \ell$  such that  $P^j$  is honest and  $\Pi.J_{\text{Out}}(P^j, m^j, \pi^j) = \top$ . Otherwise, we say that no outputs were generated.*

Note that the triple  $(P^j, m^j, \pi^j)$  with  $\Pi.J_{\text{Out}}(P^j, m^j, \pi^j) = \top$  does not mean that  $P^j$  produced the output  $(m^j, \pi^j)$ . It merely means that  $P^j$  would *accept* the output  $(m^j, \pi^j)$  given its current state and the predicate  $\Pi.J_{\text{Out}}$ . Note also that if a property holds for all AJOs it also holds for honest outputs as the adversary are given the honest outputs and can reuse them as a triple in  $(P^1, m^1, \pi^1), \dots, (P^\ell, m^\ell, \pi^\ell)$ .

## 4 Send Transferable Messages

We now present a protocol which forces a potentially corrupt sender to send a message, solving the *missing message problem* discussed earlier. Throughput, we let  $\text{NoMsg}$  be a designated symbol where  $\text{NoMsg} \notin \{0, 1\}^*$ . We use it to signal that a sender was corrupt and did not send a normal message. Ultimately,  $\text{NoMsg}$  could be mapped to a normal message, like the empty string, but it helps the exposition to assume  $\text{NoMsg} \notin \{0, 1\}^*$ . We also use another such symbol  $\perp$  and assume that  $\perp \notin \{0, 1\}^*$  and  $\perp \neq \text{NoMsg}$ .

**Definition 6 (Send Transferable Message Protocol).** *Let  $\Pi_{J_{\text{Msg}}}$  be a protocol run among  $n$  parties  $P_1, \dots, P_n$  where  $J_{\text{Msg}}$  is the parametrisable input justifier predicate. Assume that  $\Pi_{J_{\text{Msg}}}$  specifies a designated sender  $P_s$  holding an input  $m \in \{0, 1\}^*$  along with a justifier  $\pi$  such that  $J_{\text{Msg}}(P_s, m, \pi) = \top$  and parties terminate upon generating output  $y_i$  in  $P_i$ . The protocol specifies an output justifier predicate  $\Pi.J_{\text{Out}}$ .*

- **Correctness:** Honest  $P_i$  outputs  $y_i = (m_i, \pi_i)$ , where  $m_i \in \{0, 1\}^* \cup \{\text{NoMsg}\}$  and  $\pi_i \in \{0, 1\}^*$ .
- **Justified Output:** Outputs are justified. When honest  $P_i$  outputs  $y_i = (m_i, \pi_i)$  then  $J_{\text{Out}}(P_i, m_i, \pi_i) = \top$ .
- **Justified Message:** Only justified inputs can appear in justified outputs. For all AJOs  $y_i = (m_i, \pi_i)$  either  $m_i = \text{NoMsg}$  or the justifier  $\pi_i$  is of the form  $\pi_i = (\pi_i^1, \pi_i^2)$  and  $J_{\text{Msg}}(P_i, m_i, \pi_i^1) = \top$  for all honest  $P_i$ .
- **Validity:** Honest senders manage to send their intended message and only that message. If  $P_s$  is honest and has input  $(m, \pi)$ , then all honest parties  $P_j$  have output  $y_j = (m_j, \pi_j)$  with  $m_j = m$ . Furthermore, for all AJOs  $m'$  it holds that  $m' = m$ .
- **Agreement:** All outputs are the same or  $\text{NoMsg}$ . For all AJOs  $m^1$  and  $m^2$  it holds that  $m^1 = \text{NoMsg}$  or  $m^2 = \text{NoMsg}$  or  $m^1 = m^2$ .

We say that  $\Pi$  is a send transferable message with agreement (STMA) protocol with  $J_{\text{Out}}$ -justified output if it has the above properties. If it lacks agreement we call it an STM protocol. If an STMA protocol has the additional property that  $m \neq \text{NoMsg}$  for all AJOs, then we call it a justifiable broadcast. We call an output a legal STM output if it has the correctness and justified output properties.

*Remark 1.* Note that justifiable broadcast ensures that all outputs are the same, so it implies the notion of broadcast in Definition 1. It additionally has input and output justifiers, which will be convenient when using it as sub-protocol. It is straight forward to see that we can create a justified broadcast protocol  $\text{DSC}_{P_s, J_{\text{Msg}}}$  by using Dolev-Strong with sender  $P_s$  and  $t = n - 1$  corruptions and where parties only accept signatures from  $P_s$  on  $(m, \pi_{\text{Msg}})$  where  $J_{\text{Msg}}(m, \pi_{\text{Msg}}) = \top$ . The round complexity is  $\mathcal{O}(n)$ . We use this protocol later.

*Remark 2.* Note that honest parties have input  $m \neq \text{NoMsg}$ , so by Validity, if output  $m = \text{NoMsg}$  can be justified, then  $P_s$  is corrupt.

*Remark 3.* We have required that for all AJOs  $y_i = (m_i, \pi_i)$  either  $m_i = \text{NoMsg}$  or the justifier  $\pi_i$  is of the form  $\pi_i = (\pi_i^1, \pi_i^2)$  and  $J_{\text{Msg}}(P_i, m_i, \pi_i^1) = \top$  for all honest  $P_i$ . Note that the definition does not restrict what  $\pi_i^2$  is or how it affects the output of  $J_{\text{Out}}$ . Typically  $\pi_i^2$  will be a protocol dependent value proving that the message  $m_i$  resulted from running  $\Pi$  and  $J_{\text{Out}}$  will check that this is the case. Typically  $\pi_i^2$  also contains a signature on  $m_i$  from  $P_s$  to ensure validity.

## 4.1 Polariser Cast

We present a STM protocol PC called *polariser cast*. The protocol will proceed roughly as follows.

1. The sender signs its justified input and sends it to all parties.
2. If the sender  $P_s$  did not send a justified signed input in round 0 then each  $P_j$  accuses  $P_s$ .
3. If  $P_s$  should have been accused but a party  $P_j$  did not accuse  $P_s$  in round 1, then all parties  $P_k$  will accuse  $P_j$  in round 2, unless  $P_j = P_s$  such that it was already accused, *et cetera*.
4. Since only corrupted parties are accused, at some point there are no more parties to accuse, and at this point an output can be computed. Either  $P_s$  sent a signed message or  $P_s$  can be moved to **Corrupt** along with all parties with enough accusations.
5. This gives each party an output candidate, but different honest parties might hold different signed messages  $m$ .
6. The parties exchange their output candidates, and if some  $P_j$  has different signed values in any of them, then this is used as a transferable proof that  $P_s$  is corrupt.

Before describing the protocol in detail we give some helping definitions. During the protocol each  $P_i$  will keep a set  $S$  of received well-formed *elements*.

**Definition 7 (Well-formed elements).** We call an element  $e$  well-formed if it is of one of the following two forms.

**Inputs:**  $e = (\text{IN}, m, \pi, \sigma)$ , where  $\text{Ver}_{\text{vk}_s}((\text{IN}, m), \sigma) = \top$  and  $J_{\text{Msg}}(\text{P}_s, m, \pi) = \top$ .

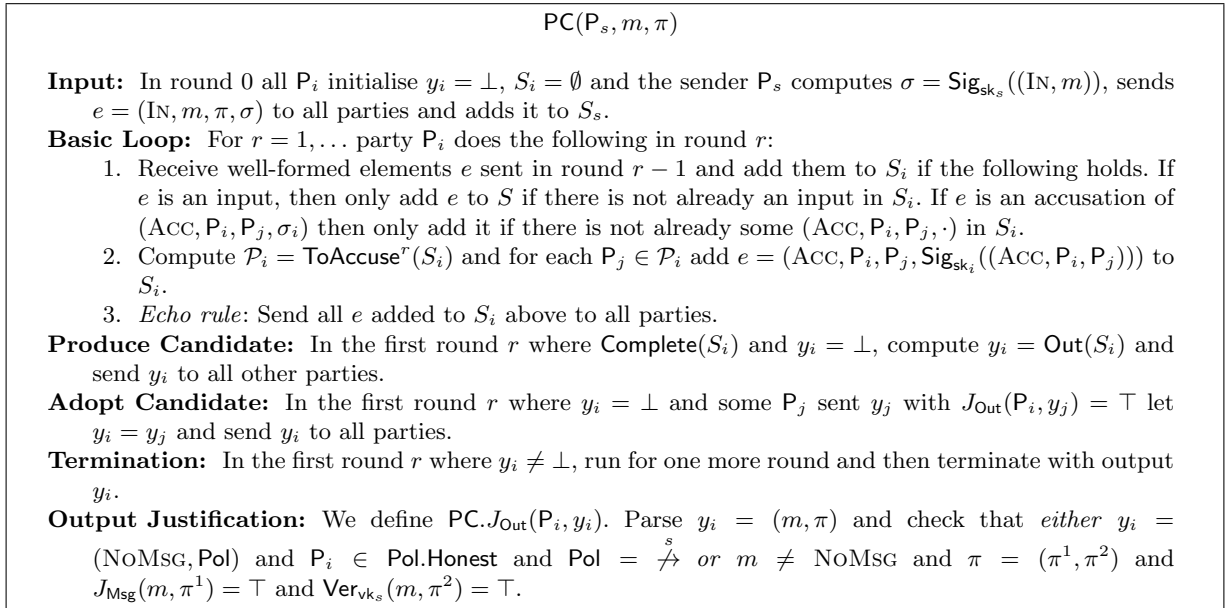
**Accusations:**  $e = (\text{ACC}, \text{P}_i, \text{P}_j, \sigma_i)$ , where  $\text{Ver}_{\text{vk}_i}((\text{ACC}, \text{P}_i, \text{P}_j), \sigma_i) = \top$ .

Each  $\text{P}_i$  has its own version of  $S$ . When we need to distinguish it we denote it by  $S_i$ . We use  $S_i^r$  to denote the value of  $S_i$  at  $\text{P}_i$  in round  $r$ . We define some helper functions used in the protocol.

**Detect corruption:** The function  $\text{ToAccuse}^r(S) \subset \mathbb{P}$  takes as input a set of well-formed elements and a round number  $r$  and computes a set of parties  $\mathcal{P} \in \mathbb{P}$ , which we think of as being corrupt.

**Complete:** We call a set of well-formed elements  $S$  *complete* if there are no more parties to accuse, i.e.,  $\text{Complete}(S) \equiv \exists r > 0 (\text{ToAccuse}^r(S) = \emptyset)$ .

**Output:** The function  $\text{Out}(S)$  takes as input a complete set of well-formed elements and computes a possible output of PC, i.e., if  $\text{Complete}(S)$  then  $\text{Out}(S) = (m, \pi)$  where  $\pi$  is a signature on  $m$  under  $\text{vk}_s$  or  $\text{Out}(S) = (\text{NoMsg}, \text{Pol})$  where  $\text{Pol} = \not\rightarrow^s$ .

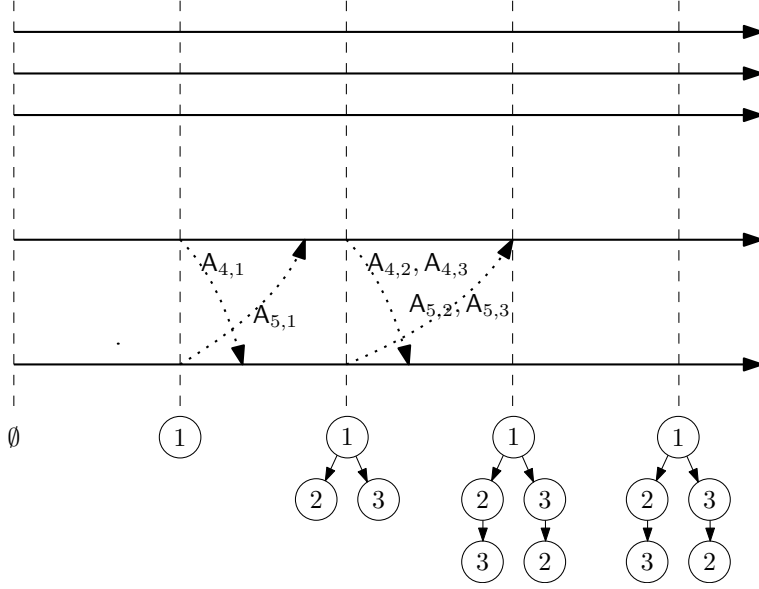


**Fig. 1.** A Sending Transferable Message Protocol.

The protocol is given in Fig. 1. We now proceed to define  $\text{ToAccuse}$  and  $\text{Out}$ . We use a tree-based definition where  $\text{Tree}$  is a tree of missing accusations. What messages are missing depends on what round we are in, so the function  $\text{Tree}^r$  also depends on  $r$ . The nodes of the tree will be elements  $(\text{P}, \rho) \in \mathbb{P} \times \mathbb{N}$ .

**Definition 8 (Tree Function).** *The output of  $\text{Tree}^r(S)$  is computed as:*

1. Let  $T$  be the empty tree.



**Fig. 2.** Party  $P_1$  is the sender. Parties  $P_1$ ,  $P_2$ , and  $P_3$  are corrupted. Their timelines are shown as the top three. Parties  $P_4$  and  $P_5$  are honest and their timelines shown at the bottom. Time runs from left to right and vertical dashed lines are round separators with the first one showing the beginning of round 0. To not clutter the figure, we do not show messages sent *to* corrupted parties. Below the timelines we show the tree built by  $P_5$ . In round 0 it is empty. In round 0 party  $P_1$  does not send its signature  $\sigma$ . Therefore  $P_5$  adds the root  $(P_1, 0)$ . We are then in round 1, so all parties in leafs of paths of length 1 should be accused, i.e.,  $P_1$ . In round 1 party  $P_4$  therefore accuses  $P_1$  and  $P_5$  also accuses  $P_1$ . The corrupted parties do not accuse  $P_1$ . Therefore  $P_5$  adds edges from  $(P_1, 0)$  to  $(P_2, 1)$  and  $(P_3, 1)$ . We are then in round 2, so all parties in leafs of paths of length 2 should be accused, so both honest parties accuse  $P_2$  and  $P_3$ . The corrupted parties do not accuse. Therefore  $P_3$ 's missing accusation of  $P_2$  is added as an edge and  $P_2$ 's missing accusation of  $P_3$  is added as an edge. Note that for instance  $P_1$ 's missing accusation of  $P_2$  is not added as an edge as we only add parties not already on a path. We are then in round 3 so parties in leafs on paths of length 3 should be accused, i.e.,  $P_3$  and  $P_2$ . However, no accusations are actually sent as equivalent accusations were sent already. We are then in round 4. No new nodes are added. Parties in leafs of paths of length 4 should be accused. There are no such paths, so the accusation is considered complete, and the protocol ends. We have  $\text{Alive} = \{P_4, P_5\}$ ,  $\text{Corrupt} = \{P_1, P_2, P_3\}$ , and  $\text{Accuse} = \{A_{4,1}, A_{5,1}, A_{4,2}, A_{4,3}, A_{5,2}, A_{5,3}\}$ , so we have a legal  $P_1$ -polariser.

2. If  $\nexists(\text{IN}, m, \cdot, \cdot) \in S$  then add  $(P_s, 0)$  to  $T$  as the root.
3. For  $\rho = 1, \dots, r$ :
  - (a) For all  $\text{path} \in T$  with  $|\text{path}| = \rho$ :
    - i. Let  $(P_j, \rho - 1) = \text{leaf}(\text{path})$ .
    - ii. For all  $P_k \in \mathbb{P}$  where  $(P_k, \cdot) \notin \text{path}$  and  $\nexists(\text{ACC}, P_k, P_j, \cdot) \in S$ , add the edge  $((P_j, \rho - 1), (P_k, \rho))$  to  $T$ .
4. Output  $T$ .

We think of  $\text{Tree}^r(S)$  as the tree of missing elements relative to an honest run of  $\text{PC}(P_s, \dots)$ . As an example, if  $P_s$  is honest it should send  $(\text{IN}, m, \cdot, \cdot)$  in round 0, so if  $\nexists(\text{IN}, m, \cdot, \cdot) \in S$  then we add  $(P_s, 0)$  to signify that  $P_s$  omitted a message in round 0. Therefore, if the tree has the path  $((P_s, 0))$  then all  $P_k$  should have accused  $P_s$  in round 1. If  $P_k$  does not do this, then in the iteration of the loop with  $\rho = 1$  the path  $\text{path} = ((P_s, 0))$  of length 1 will get considered and so will  $(P_s, 0) = \text{leaf}(\text{path})$ . So if  $P_k \neq P_s$  did not accuse  $P_s$  then we will have  $\nexists(\text{ACC}, P_k, P_s, \cdot) \in S$  and hence  $((P_s, 0), (P_k, 1))$  gets added to the tree  $T$ . So we add an edge to  $(P_k, 1)$  to signify that in round 1 party  $P_k$  failed an obligation and then points from  $(P_s, 0)$  to say that the obligation was to accuse  $P_s$  because  $P_s$  failed an obligation in the previous round. The reason that  $\text{Tree}^r$  depends on  $r$  is that some accusation might be missing simply because the parties did not have a chance to send them yet.

We now define  $\text{ToAccuse}^r(S)$ . To motivate the definition, consider a tree  $\text{Tree}^{r-1}(S_i)$  at  $P_i$  at the beginning of round  $r$ , i.e., after receiving the accusations  $(\text{ACC}, \cdot, \cdot)$  sent out in round  $r - 1$ . If  $\text{path} = (\dots, (P_j, r - 1)) \in \text{Tree}^{r-1}(S_i)$  this is because  $P_j$  missed an obligation in round  $r - 1$ . Therefore  $P_i$  must accuse  $P_j$ . This motivate the following definition

$$\text{ToAccuse}^r(S) = \{ P_j \mid \exists(\dots, (P_j, r - 1)) \in \text{Tree}^{r-1}(S) \} .$$

If a set  $S$  is complete then it allows to compute an output as follows.


**Definition 9 (Output).** *The function  $\text{Out}(S)$  is defined as follows.*

1. The input is a complete set  $S$ , so we can find the smallest  $r$  such that  $\text{ToAccuse}^r(S) = \emptyset$ .
2. If  $r = 1$  then pick  $(\text{IN}, m, \sigma, \pi) \in S$  and output  $(m, (\sigma, \pi))$ .
3. If  $r > 1$  then output  $(\text{NoMsg}, \text{Pol} = (\text{Alive}, \text{Corrupt}, \text{Accuse}))$ , where

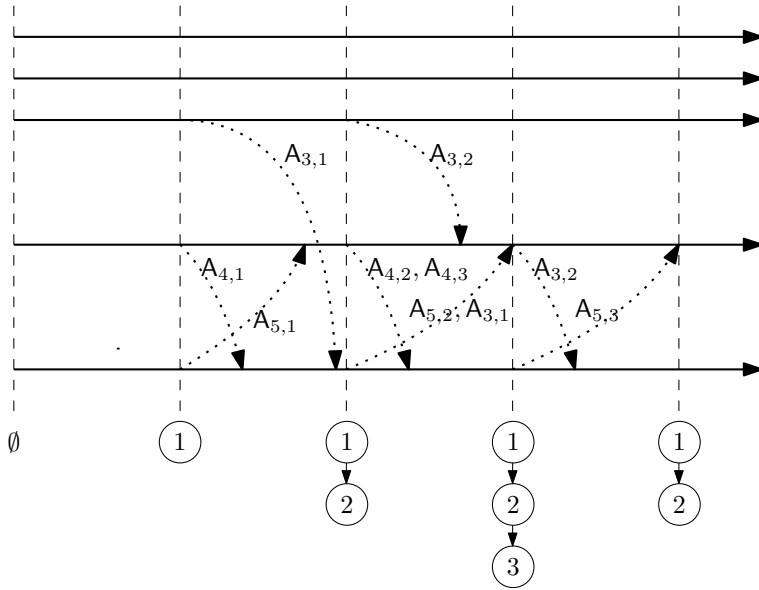
$$\begin{aligned} \text{Corrupt} &= \{ P_j \mid \exists(P_j, \cdot) \in \text{nodes}(\text{Tree}^r(S)) \} , \\ \text{Accuse} &= S , \\ \text{Alive} &= \mathbb{P} \setminus \text{Corrupt} . \end{aligned}$$

We proceed to prove PC secure when using the above definitions of  $\text{ToAccuse}$ ,  $\text{Complete}$ , and  $\text{Out}$ . Before the proof it may be instructive to consider the example runs of the protocol in Figs. 2 to 4.

**Definition 10 (equivalent sets).** *Let  $S$  and  $T$  be two sets of well-formed elements. We say that  $S \sqsubseteq T$  if  $\exists(\text{IN}, m, \pi, \sigma_s) \in S$  implies that  $\exists(\text{IN}, m', \pi', \sigma'_s) \in T$  and  $\exists(\text{ACC}, P_i, P_j, \sigma_i) \in S$  implies that  $\exists(\text{ACC}, P_i, P_j, \sigma'_i) \in T$ . We call two sets equivalent if  $S \sqsubseteq T$  and  $T \sqsubseteq S$ . We call two elements equivalent if  $\{e_1\}$  and  $\{e_2\}$  are equivalent.*



**Fig. 3.** For notation see Fig. 2. In round 0 party  $P_1$  does not send its signature  $\sigma$ . Therefore  $P_5$  adds the root  $(P_1, 0)$ . We are then in round 1 and parties  $P_4$  and  $P_5$  accuse  $P_1$ . The corrupted parties do not accuse  $P_1$ , but  $P_3$  forwards a signature  $\sigma$  by  $P_1$  to  $P_5$ . Therefore the tree computed by  $P_5$  in round 2 is again empty. Therefore the accusation is over for  $P_5$ . It outputs  $\sigma$  (it terminates one round later, not shown). Note that  $P_4$  is in the same situation as in Fig. 2. Its tree will look like that of  $P_5$  in round 2 in Fig. 2. So, it accuses  $P_2$  and  $P_3$ . By the echo rule  $P_5$  forwards  $\sigma$  to  $P_4$  which will then have an empty tree by round 3 and outputs  $\sigma$ .



**Fig. 4.** For notation see Fig. 2. In round 0 party  $P_1$  does not send its signature  $\sigma$ . Therefore  $P_5$  adds the root  $(P_1, 0)$ . We are then in round 1 and parties  $P_4$  and  $P_5$  accuse  $P_1$ . Party  $P_3$  accuses  $P_1$  but sends the accusation only to  $P_5$ . Party  $P_2$  does not accuse  $P_1$ . Party  $P_5$  adds an edge representing the missing accusation of  $P_1$  by  $P_2$ . Party  $P_4$  being in the same situation as in Fig. 2 accuses  $P_2$  and  $P_3$ . Party  $P_5$  accuses all parties which are leafs on paths of length 2, i.e., party  $P_2$ . It also forwards  $A_{3,1}$  because of the echo rule. Now  $P_3$  accuses  $P_2$  but only towards  $P_4$ . Therefore  $P_5$  is missing the accusation of  $P_2$  by  $P_3$  and adds an edge to represent it. It then has a path of length 3 in round 3 and thus accuses  $P_3$ . But in the same round  $P_4$  forwards  $A_{3,2}$  because of the echo rule. Therefore by round 4 the tree computed by party  $P_5$  is back to height 1 and the protocol ends for  $P_5$ . It outputs  $\text{Alive} = \{P_3, P_4, P_5\}$ ,  $\text{Corrupt} = \{P_1, P_2\}$  and  $\text{Accuse} = \{A_{4,1}, A_{5,1}, A_{3,1}, A_{5,2}, A_{4,2}, A_{5,3}, A_{3,2}\}$ , which is a legal  $P_1$ -polariser.



**Lemma 2 (propagation lemma).** *For all honest  $P_i$  and  $P_j$  and rounds  $r > 0$  reached in the protocol it holds that  $S_i^{r-1} \sqsubseteq S_j^r$ .*

*Proof.* This follows from the fact that all elements  $s$  added to  $S_i$  get forwarded to  $P_j$  and if  $s$  is considered well-formed by  $P_i$  then it is also considered well-formed by  $P_j$ . Therefore  $s$  is added to  $S_j^r$ , unless  $S_j^r$  contains an equivalent element.  $\square$

**Lemma 3 (tree monotonicity lemma).** *For all sets of well-formed elements  $S, T$  and all  $r \geq 0$  it holds that*

$$S \sqsubseteq T \implies \text{Tree}^r(T) \subseteq \text{Tree}^r(S) .$$

*Proof.* Note that for an object  $o$ , root or edge, to be included in  $\text{Tree}^r(T)$  it is required that some element is *missing* in  $T$ , i.e.,  $\#(\text{IN}, m, \cdot, \cdot) \in T$  or  $\#(\text{ACC}, P_k, P_j, \cdot) \in T$ . Since  $S \sqsubseteq T$  these conditions imply that  $\#(\text{IN}, m, \cdot, \cdot) \in S$  and  $\#(\text{ACC}, P_k, P_j, \cdot) \in S$ , so the same object  $o$  gets included in  $\text{Tree}^r(S)$ .  $\square$

The following corollary is important in showing that honest do not accuse honest parties. The tree  $\text{Tree}^{r-1}(S_i^{r-1})$  is the tree that  $P_i$  used in round  $r - 1$  to calculate who it should accuse. The tree  $\text{Tree}^{r-1}(S_j^r)$  is the tree that  $S_j$  uses in round  $r$  to calculate who  $S_i$  ought to have sent an accusation against—it uses the same function  $\text{Tree}^{r-1}$ , but its own set  $S_j^r$ . If  $\text{Tree}^{r-1}(S_j^r) \sqsubseteq \text{Tree}^{r-1}(S_i^{r-1})$  then  $P_j$  does not expect to receive any accusations which are not sent.

**Corollary 1.** *For all honest  $P_i$  and  $P_j$  and rounds  $r > 0$  reached in the protocol it holds that*

$$\text{Tree}^{r-1}(S_j^r) \sqsubseteq \text{Tree}^{r-1}(S_i^{r-1}) .$$

*Proof.* By the preceding lemmas we have that  $S_i^{r-1} \sqsubseteq S_j^r$  and that  $S \sqsubseteq T \implies \text{Tree}^\rho(T) \subseteq \text{Tree}^\rho(S)$  for all  $S, T$  and  $\rho$ . Set  $S = S_i^{r-1}$ ,  $T = S_j^r$  and  $\rho = r - 1$ .  $\square$

**Lemma 4.** *If  $P_i$  is honest and accuses  $P_j$  then  $P_j$  is corrupted.*

*Proof.* By construction, if  $P_i$  accuses  $P_j$  in round  $r$  then  $P_j \in \text{ToAccuse}^r(S_i^r)$ . By definition this means that  $\exists(\dots, (P_j, r - 1)) \in \text{Tree}^{r-1}(S_i^r)$ . If  $r = 1$  then this implies that  $\exists((P_j, 0)) \in \text{Tree}^0(S_i^r)$ , and hence  $P_j = P_s$ , as only  $P_s$  can occur in the root. Therefore  $\#(\text{IN}, m, \cdot, \cdot) \in S_i$ . Hence,  $P_j = P_s$  did not send its signed input to  $P_i$  in round 0. Therefore  $P_j = P_s$  is corrupted. If  $r > 1$  then we have that  $\exists(\dots, (P_k, r - 2), (P_j, r - 1)) \in \text{Tree}^{r-1}(S_i^r)$ . By construction (see Item 3(a)ii in Definition 8) this implies that  $\exists(\dots, (P_k, r - 2)) \in \text{Tree}^{r-1}(S_i^r)$  and  $\#(\text{ACC}, P_j, P_k, \cdot) \in S_i^r$ . If  $P_j$  is honest this implies that  $\exists(\dots, (P_k, r - 2)) \in \text{Tree}^{r-1}(S_j^{r-1})$ , as  $\text{Tree}^{r-1}(S_i^r) \sqsubseteq \text{Tree}^{r-1}(S_j^{r-1})$ . Therefore  $P_k \in \text{ToAccuse}^{r-1}(S_j^{r-1})$ . So if  $P_j$  is honest it sent  $(\text{ACC}, P_j, P_k, \cdot)$  to  $P_i$  in round  $r - 1$ . This contradicts  $\#(\text{ACC}, P_j, P_k, \cdot) \in S_i^r$ .  $\square$

The following lemma shows that when the set  $S$  is complete at an honest party such that it terminates, then  $\text{Out}(S)$  produces a correct output, i.e., if there is no signature in  $S$ , then a polariser is produced proving that  $P_s$  is corrupt.

**Lemma 5 (justified output).** *If  $S$  is held by an honest party  $P_i$  and  $\text{Complete}(S)$ , such that  $P_i$  produces output  $\text{Out}(S)$ , then  $\text{PC}.J_{\text{Out}}(P_i, \text{Out}(S))$ .*

*Proof.* We want to prove that  $\text{PC}.J_{\text{Out}}(P_i, \text{Out}(S)) = \top$ . This means that if we let  $y_i = \text{Out}(S)$ , as defined in Definition 9, then we have to make sure that  $\text{PC}.J_{\text{Out}}(P_i, y_i)$  where  $\text{PC}.J_{\text{Out}}$  is defined in **Output Justification** in Fig. 1. We write this out. First parse  $y_i = (m, \pi)$ . We then have to prove that *either* 1)  $y_i = (\text{NoMsg}, \text{Pol})$  and  $P_i \in \text{Pol.Honest}$  and  $\text{Pol} = \overset{s}{\neq}$  or 2)  $m \neq \text{NoMsg}$  and  $\pi = (\pi^1, \pi^2)$  and  $J_{\text{Msg}}(m, \pi^1) = \top$  and  $\text{Ver}_{\text{vks}}(m, \pi^2) = \top$ .

By  $\text{Complete}(S)$  there exists  $r$  such that  $\text{ToAccuse}^r(S) = \emptyset$ . Assume that  $r = 1$ . From  $\text{Complete}(S)$  we get that  $\text{ToAccuse}^1(S) = \emptyset$ . This implies that  $\nexists(\dots, (P_j, r-1)) \in \text{Tree}^{r-1}(S)$  which for  $r = 1$  means that  $\nexists((P_j, 0)) \in \text{Tree}^0(S)$ , which by the construction of the tree  $T$  in the algorithm  $\text{Tree}^0$  defined in Definition 8 means that it is *not* the case that  $\nexists(\text{IN}, m, \cdot, \cdot) \in S$ . So, there is some well-formed  $(\text{IN}, m, \cdot, \cdot) \in S$ . Therefore the output is of the form in case 2 above.

Assume then that  $r > 1$ . From  $\text{Complete}(S)$  we get that  $\text{ToAccuse}^r(S) = \emptyset$  and by  $r$  being minimal we have that  $\text{ToAccuse}^{r-1}(S) \neq \emptyset$ . From  $\text{ToAccuse}^{r-1}(S) \neq \emptyset$ , it follows that there is at least one path in  $\text{Tree}^{r-1}(S)$  and hence also a root. Therefore  $(P_s, 0) \in \text{nodes}(\text{Tree}^{r-1}(S))$  and hence  $P_s \in \text{Pol.Corrupt}$ . We therefore just have to show that  $\text{Pol}$  is a legal polariser. Completeness follows from  $\text{Alive} = \mathbb{P} \setminus \text{Corrupt}$ . Since  $S$  contains only well-formed elements and  $\text{Accuse} = S$  by Definition 9, for justifiability it is sufficient to prove that for all  $(P_i, P_j) \in \text{Alive} \times \text{Corrupt}$  it holds that  $(\text{Accuse}, P_i, P_j, \cdot) \in S$ . So, assume that  $(P_i, P_j) \in \text{Alive} \times \text{Corrupt}$ . This implies that  $(P_j, \rho) \in \text{nodes}(\text{Tree}^\rho(S))$  for some  $\rho < r$ . We argue that this implies that  $S$  contains  $(\text{ACC}, P_i, P_j, \cdot)$ . Assume to the contrary that  $S$  does not contain  $(\text{ACC}, P_i, P_j, \cdot)$ . Then it would be the case that  $((P_j, \rho), (P_i, \rho+1)) \in \text{Tree}^r(S)$  by Item 3(a)ii in Definition 8 unless  $(P_i, \cdot)$  was already on the path in question (but  $(P_i, \cdot)$  being on the path contradicts  $P_i \in \text{Alive}$  as we added to  $\text{Corrupt}$  all parties on all paths by construction of Definition 9). But if  $((P_j, \rho), (P_i, \rho+1)) \in \text{Tree}^r(S)$  then from  $\rho+1 \leq r$  and because we assume that  $(\text{ACC}, P_i, P_j, \cdot) \notin S$  it is not the case that  $\text{ToAccuse}^r(S) \neq \emptyset$ , as we would have  $P_i \in \text{ToAccuse}^r(S)$  by construction. Since we have as premise that  $\text{ToAccuse}^r(S) = \emptyset$  it follows that  $(\text{ACC}, P_i, P_j, \cdot) \in S$ , as desired.  $\square$

**Theorem 1.** *The protocol PC is an STM protocol for  $t < n$ . Furthermore, assume that PC has inputs  $(m, \pi)$  with  $|(m, \pi)| \leq \ell$ . Let  $\lambda$  be the length of a signature. Then the protocol has communication complexity  $\mathcal{O}(n^2\ell + n^4\lambda)$  and the size of the justified output is at most  $\mathcal{O}(\ell + n^2\lambda)$ . The protocol uses at most  $f + 2$  rounds.*

*Proof.* Correctness follows by construction of  $\text{Out}$ . Justified output follows from Lemma 5. Justified message follows by construction of  $J_{\text{Out}}$ . Validity follows by the fact that if  $P_s$  is honest then in any polariser  $\text{Pol}$  accepted by an honest party it holds that  $P_s \in \text{Alive}$  by Lemma 4. We then count communication complexity. We ignore constant factors in the counting. In round 0 party  $P_s$  sends to all parties its input of length  $\ell$  and a signature. This is the sending of  $n(\ell + \lambda)$  bits. During the basic loop each party forwards at most one well-formed input from  $P_s$  to other parties, so this is at most  $n^2(\ell + \lambda)$  bits. Besides this, each  $P_i$  might send an accusation of each  $P_j$  which will then be relayed by each other party. This

is the flooding of at most  $n^4\lambda$  bits. The output consists of at most one well-formed input of  $P_s$  so is at most  $\ell + n^2\lambda$  bits. In Produce Candidate and Adopt Candidate each party sends it to at most  $n$  other parties, yielding communication at most  $n^2\ell + n^4\lambda$ . We consider round complexity. For illustration consider the easy case for  $f = 0$ . In the first round  $P_s$  sends its signed input to all parties and in round 2 all parties send their adopted candidate  $y_i$ . This is  $2 = f + 2$  rounds. Note then that if the protocol in rounds  $r \geq 1$  do not send an adopted candidate it is because  $\neg\text{Complete}(S)$ , which implies that  $\text{ToAccuse}^r(S) \neq \emptyset$  which by Item 3(a)ii in Definition 8 implies that  $P_k$  for which  $(P_k, \cdot) \notin \text{path}$  is added to  $T$  in `Tree`, extending `path` by length 1. After this `path` contains one more corrupted party. There are at most  $f$  corrupted parties. So this adds at most  $f$  extra rounds, for a total of at most  $2 + f$ .  $\square$

The bulk for the communication of PC is the flooding of up to  $n^2$  accusations. It turns out this can be compressed accross multiple runs of PC as each accusation needs only be sent once.

**Lemma 6 (Amortized Communication Complexity).** *Assume that in the life time of the system  $\iota$  instances  $\text{PC}^1, \dots, \text{PC}^\iota$  are run and that  $\text{PC}^i$  has inputs  $(m^i, \pi^i)$  with  $|(m^i, \hat{\pi}^i)| \leq \ell_i$ , where  $\hat{\pi}^i$  is  $\pi^i$  with all accusations removed. Then the communication complexity of running all  $\iota$  copies can be compressed to  $\mathcal{O}(n^2 \sum_i \ell_i + n^4\lambda)$  without affecting the security of the protocol.*

*Proof.* If in a given system an accusation  $e = (\text{ACC}, P_i, P_j, \sigma_i)$  was sent as part of running one PC, then do not send it again. Also, add it to all sets  $S_i$  in all copies. Also, do not send it as part of the justifications after it was sent once. If  $e = (\text{ACC}, P_i, P_j, \sigma_i)$  was received once then add it to all incoming justifications. This can be seen to not affect the execution of the protocol. If  $\pi$  justified  $m$  then it also does so after adding one more accusation. This way, overall, each of the  $n^2$  possible  $e = (\text{ACC}, P_i, P_j, \sigma_i)$  will be sent at most once per pair of parties for a total of  $n^4\lambda$  communication.  $\square$

## 5 Generic Transferable Justifiers

Our second to last general tool is generic transferable justifiers. This is a short hand capturing the idea that a message can be justified by sending along all messages used to compute it and let the receiver recompute the message. In all protocols  $\Pi$  which follow, the protocols proceed in super rounds, where in each super round  $s = 1, 2, \dots$  the parties invoke a sub-protocol  $\Pi^s$ , wait for its outputs, and then run the next super round. In the first super round, we assume that each  $P_i$  has an input  $(x_i, \pi_i)$ , where  $J_{\text{In}}(P_i, x_i, \pi_i) = \top$  and that this is the input to the first sub-protocol  $\Pi^1$ . As a sentinel, let  $x_i^0 = x_i$  and  $\pi_i^0 = \pi_i$  and define  $\Pi^0.J_{\text{Out}} := \Pi.J_{\text{In}}$ . Consider then a super round  $s$  where in the previous  $s - 1$  super rounds the parties ran protocols  $\Pi^1, \dots, \Pi^{s-1}$  and in super round  $s$  are to run  $\Pi^s$ . For  $k = 1, \dots, s - 1$ , let  $y_i^k$  be the output of  $P_i$  from  $\Pi^k$  and let  $\pi_i^k$  be the justifier. Then the message to be input to  $\Pi^s$  by  $P_i$  will be computed using a function

$$x_i^s = \text{NxtInp}(\{y_i^k\}_{k=0}^{s-1}) \quad (1)$$

and the accompanying justifier computed using a function

$$\pi_i^s = \text{NxtJst}(\{(y_i^k, \pi_i^k)\}_{k=0}^{s-1}) . \quad (2)$$

**Definition 11 (Generic Justifier).** *When we say that we use a generic justifier predicate in a setting as described above then we mean that*

$$\text{NxtJst}(\{(y_i^k, \pi_i^k)\}_{k=0}^{s-1}) = \{(y_i^k, \pi_i^k)\}_{k=0}^{s-1} .$$

Furthermore, the input justifier predicate for  $\Pi^s$  will be

$$\begin{aligned} \Pi^s . J_{\text{In}}(\text{P}_j, x_i^s, \{(y_i^k, \pi_i^k)\}_{k=0}^{s-1}) &\equiv \\ x_i^s = \text{NxtInp}(\{(y_i^k)\}_{k=0}^{s-1}) \wedge \bigwedge_{k=0}^{s-1} \Pi^k . J_{\text{Out}}(\text{P}_j, y_i^k, \pi_i^k) . \end{aligned}$$

We also use generic justifiers for the outputs. We simply show that the output can be computed from justified outputs of the sub-protocols.

**Definition 12 (Generic Justifier for Output).** *In a protocol with  $\sigma$  super-rounds, we give generic justified outputs by computing  $x_i^{\sigma+1}$  as if it was an input for a virtual round  $\sigma + 1$  (which we can also think of as the first round of the next protocol where  $x_i^{\sigma+1}$  is input) and then we let the output of  $\text{P}_i$  be  $y_i = x_i^{\sigma+1}$  and  $\pi_i = \pi_i^{\sigma+1}$  where  $x_i^{\sigma+1}$  is computed as in Eq. (1) and  $\pi_i$  as in Eq. (2).*

When an output is generically justified, then from the justifier one can extract a view of the protocol execution of the party producing it. We capture this in the following definition.

**Definition 13 (Unfolded View).** *Consider an AJO  $y$  in some  $\sigma$ -super-round protocol  $\Pi$  as described above using generic justifiers.<sup>4</sup> Let  $\pi'$  be the justifier. By the unfolded view of  $y$  we mean*

$$\text{unfold}(y, \pi') := ((x, \pi), (y^1, \pi^1), \dots, (y^\sigma, \pi^\sigma), y) ,$$

where by construction  $\pi'$  contains  $x = x^0$  and  $\pi = \pi^0$  such that  $\Pi . J_{\text{In}}(x, \pi) = \top$  and outputs  $y^1, \dots, y^\sigma$  of its sub-protocols  $\Pi^1, \dots, \Pi^\sigma$  along with justifiers  $\pi^1, \dots, \pi^\sigma$  such that  $\Pi^s . J_{\text{Out}}(y^s, \pi^s) = \top$  and such that  $y = \text{NxtInp}(\{y^k\}_{k=0}^\sigma)$ <sup>5</sup> is a correctly constructed output.

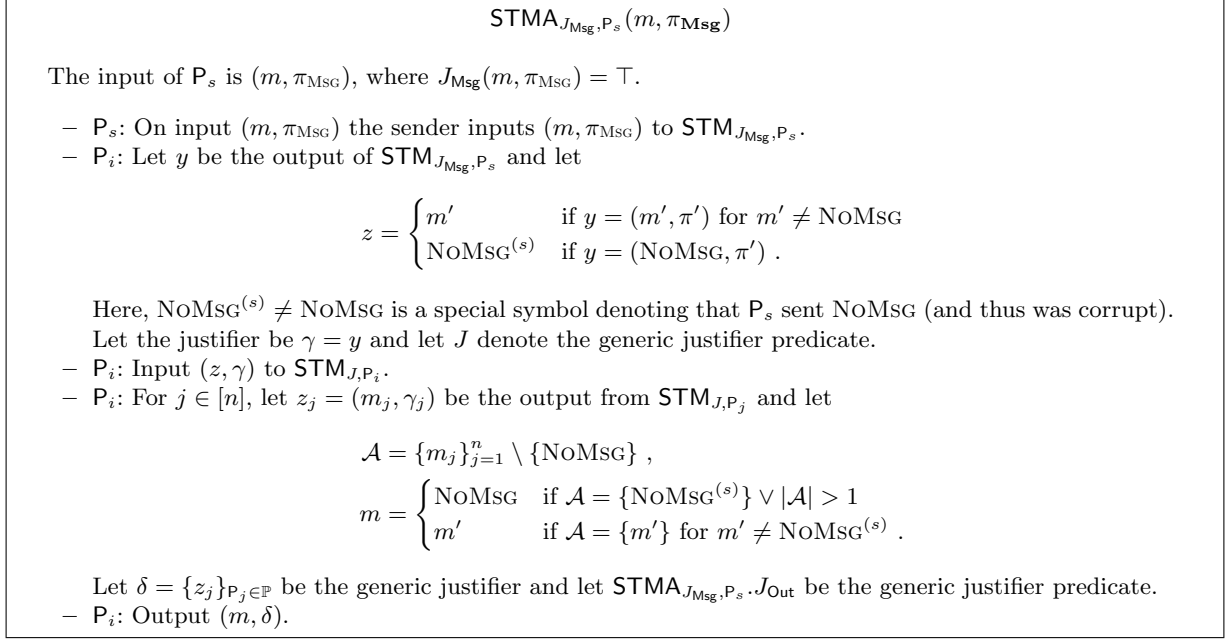
*Remark 4 (Inconsistent Unfolded Views).* Note that the unfolded view does not demonstrate that the input  $x^s$  to  $\Pi^s$  was computed according to the protocol from  $(x^0, \dots, x^{s-1})$ . It only shows that the output  $y^s$  which is included in the justifier  $\pi'$  for  $\Pi^s$  was justified and the final  $y$  was computed from  $x$  and these justified  $y^s$ . In particular, if we were to unfold  $y^s$  it might give an input  $x^{s'}$  where  $x^{s'} \neq x^s$ . This is intended and seems crucial in controlling that the size of generic justifiers does not grow exponentially. It is also a desired feature that  $y^s$  might not be consistent with  $x^s = \text{NxtInp}(\{y^k\}_{k=0}^{s-1})$  in the view of  $\text{P}_i$ . This will soon allow us that  $\text{P}_i$  takes over a justified output  $y_j^s$  from another party  $\text{P}_j$  for some sub-protocol, i.e., it lets  $y_i^s = y_j^s$  without having to recursively check consistency of the output it takes over.

<sup>4</sup> Recall that this means that some honest party would accept  $y$ .

<sup>5</sup> Recall that we compute the output as if it was the input for a next virtual round.

## 6 Send Transferable Messages with Agreement and Justified Grade Cast

We now show how to add agreement to any STM protocol by giving an STMA protocol using STM as sub-protocol. The protocol is given in Fig. 5.



**Fig. 5.** An STMA protocol.

**Theorem 2.** *Protocol  $\text{STMA}_{J_{\text{Msg}}, P_s}$  is an STMA protocol. Assuming that STM has inputs  $(m, \pi)$  with  $|m| \leq \ell$ , it has communication complexity  $\mathcal{O}(n^2\ell + n^4\lambda)$  and the size of the justified output is at most  $\mathcal{O}(\ell + n^2\lambda)$ . The amortized complexity can be optimised to be as in Lemma 6. If  $P_s$  is honest, it uses 2 rounds.*

Validity and agreement are proven below. The remaining properties are trivial. For communication, use that we run STM twice on inputs of length  $\mathcal{O}(\ell)$ .

**Lemma 7 (Validity).** *Protocol  $\text{STMA}_{J_{\text{Msg}}, P_s}$  in Fig. 5 is valid.*

*Proof.* Suppose  $P_s$  is honest and has input  $(m, \pi_{\text{Msg}})$ . In this case,  $P_s$  inputs  $(m, \pi_{\text{Msg}})$  to  $\text{STM}_{J_{\text{Msg}}, P_s}$  such that  $J_{\text{Msg}}(P_s, m, \pi_{\text{Msg}}) = \top$  and hence  $J_{\text{Msg}}(m, \pi_{\text{Msg}}) = \top$  as  $P_s$  is honest.<sup>6</sup> By validity of  $\text{STM}_{J_{\text{Msg}}, P_s}$ , every AJO is of the form  $y = (m, \pi')$  with  $J_{\text{Msg}}(m, \pi') = \top$  and all honest parties get an output. Hence, all honest parties  $P_i$  set  $\gamma_i = (m, \pi')$ ,  $z = m$  and input  $(z, \gamma)$  to  $\text{STM}_{J, P_i}$ . By validity of  $\text{STM}_{J, P_j}$ , every AJO  $(m, \gamma_j)$  has  $J(m, \gamma_j) = \top$  when  $P_j$  is

<sup>6</sup> Recall that  $J(x, \pi) = \top$  denotes that  $J(P_i, x, \pi) = \top$  for all honest  $P_i$ , which holds if it holds for a single  $P_i$  (cf. Definition 3).

honest. For all dishonest  $P_j$ , the properties of the generic justifier and the justified messages property of  $\text{STM}_{J,P_j}$  implies that for all AJOs  $z_j$  either  $z_j = (m, \gamma_j)$  s.t.  $J_{\text{Msg}}(m, \gamma_j) = \top$  or  $z_j = (\text{NoMsg}, \gamma_j)$  s.t.  $\text{STM}_{J,P_j} \cdot J_{\text{Out}}(m, \gamma_j) = \top$ . Hence, for all honest  $P_i$ ,  $m_j = m$  for all honest  $P_j$  and  $m_j \in \{m, \text{NoMsg}\}$  for all dishonest  $P_j$ . Note that all AJOs must include  $\delta = \{z_j\}_{P_j \in \mathbb{P}}$  and be valid with respect to the generic justifier. The properties of the generic justifier therefore imply that the unfolded view of any AJO must have  $\mathcal{A} = \{m\}$  and hence be of the form  $(m, \delta)$ .  $\square$

**Lemma 8 (Agreement).** *Protocol  $\text{STMA}_{J_{\text{Msg}}, P_s}$  in Fig. 5 has agreement.*

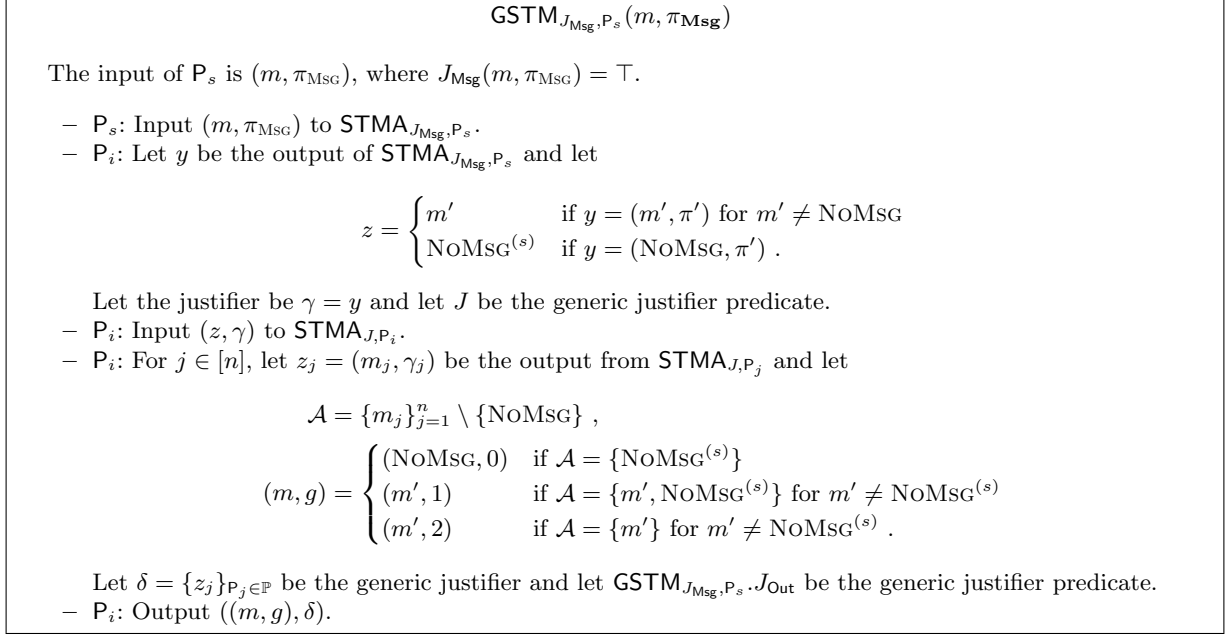
*Proof.* By the properties of STM, any AJO of  $\text{STM}_{J_{\text{Msg}}, P_s}$  is justified by  $\text{STM}_{J_{\text{Msg}}, P_s} \cdot J_{\text{Out}}$ . We first show that if there exists an AJO  $(m' \neq \text{NoMsg}, \pi')$  from  $\text{STM}_{J_{\text{Msg}}, P_s}$ , then any AJO must be of the form  $(m', \delta)$  or  $(\text{NoMsg}, \delta)$ , where  $\delta = \{z_j\}_{P_j \in \mathbb{P}}$  s.t. the output is valid w.r.t. the generic justifier. To see this, note that by validity of  $\text{STM}_{J, P_i}$ , all its AJOs must be of the form  $z_i = (m', \gamma_i)$ . By the properties of the generic justifier, any party  $P_j$  must include  $z_i$  in the justifier of its output  $(m, \delta)$  in order for it to be justifiable toward any honest party. Hence, the unfolded view of any AJO must form its output based on  $\mathcal{A}$  which includes  $m'$ . This shows that only  $(m', \delta)$  and  $(\text{NoMsg}, \delta)$  can be justifiable outputs for  $P_j$  in this case. The case where there exist a justified output  $(\text{NoMsg}, \pi)$  from  $\text{STM}_{J_{\text{Msg}}, P_s}$  is similar. Here, validity of  $\text{STM}_{J, P_i}$  and the properties of the generic justifier together instead imply that any AJO must include  $\text{NoMsg}^{(s)}$  in  $\mathcal{A}$  in order for its output to be generically justifiable. Hence, all AJOs will have  $m = \text{NoMsg}$ .  $\square$

## 6.1 Justified Grade Cast

In STMA, some parties might have output  $\text{NoMsg}$  while some have  $m \neq \text{NoMsg}$ . We now show how to upgrade an STMA to a graded STM where the output contains a grade  $g \in \{0, 1, 2\}$  which indicates the confidence in this output. When  $g = 2$ , no AJO can have  $m \neq \text{NoMsg}$ . Furthermore, grades are at most 1 apart and honest senders always produce grade 2. A detailed definition is given in Definition 14. Our protocol is given in Fig. 6.

**Definition 14 (Graded Send Transferable Message).** *Let  $\Pi$  be a protocol run among  $n$  parties  $P_1, \dots, P_n$  and let  $J_{\text{Msg}}$  be a parametrizable input justifier predicate. Assume that  $\Pi$  specifies a designated sender  $P_s$  holding input  $(m, \pi_{\text{Msg}})$  such that  $J_{\text{Msg}}(P_s, m, \pi_{\text{Msg}}) = \top$  and parties terminate upon generating output in  $\Pi$ . As part of its code the protocol specifies an output justifier predicate  $\Pi \cdot J_{\text{Out}}$ . We say that  $\Pi$  is a graded send transferable message (GSTM) protocol if it has the following properties:*

- **Correctness:** *Honest  $P_i$  outputs  $y_i = ((m_i, g_i), \pi_i)$  for  $g_i \in \{1, 2\}$  and  $m_i \neq \text{NoMsg}$  or  $y_i = ((m_i = \text{NoMsg}, g_i = 0), \pi_i)$ .*
- **Justified Output:** *When honest  $P_i$  outputs  $y_i = ((m_i, g_i), \pi_i)$  then  $J_{\text{Out}}(P_i, (m_i, g_i), \pi_i) = \top$ .*
- **Justified Message:** *For all AJOs  $y_i = ((m_i, g_i), \pi_i)$  either  $m_i = \text{NoMsg}$  and  $g_i = 0$  or  $\pi_i$  is of the form  $\pi_i = (\pi_i^1, \pi_i^2)$  and  $J_{\text{Msg}}(m_i, \pi_i^1) = \top$ .*



**Fig. 6.** A justified gradecast protocol.

- **Validity:** If  $P_s$  is honest and has input  $(m, \pi_{\text{Msg}})$ , then all honest parties  $P_j$  have output  $m_j = m$  and  $g_j = 2$ . Furthermore, for all AJOs  $(m', g')$  it holds that  $m' = m$  and  $g' = 2$ .
- **Graded Agreement:** For all AJOs  $(m^1, g^1)$  and  $(m^2, g^2)$  it holds that  $|g^1 - g^2| \leq 1$  and if  $g^1, g^2 > 0$  then  $m^1 = m^2$ .

**Theorem 3.** Protocol  $\text{GSTM}_{J_{\text{Msg}}, P_s}$  is a GSTM for  $t < n$ . Assume that GSTM has inputs  $(m, \pi)$  with  $|m| \leq \ell$ . Then the protocol has communication complexity  $\mathcal{O}(n^2\ell + n^4\lambda)$  and the size of the justified output is at most  $\mathcal{O}(\ell + n^2\lambda)$ . The amortized complexity can be optimised to be as in Lemma 6.

As for STMA, all properties but validity and agreement are trivial and so is the communication complexity.

**Lemma 9 (Validity).** Protocol  $\text{GSTM}_{J_{\text{Msg}}, P_s}$  in Figure 6 is valid.

*Proof.* Suppose that  $P_s$  is honest and has input  $m$ . By validity of  $\text{STMA}_{J_{\text{Msg}}, P_s}$ , every AJO  $y$  for that protocol is  $y = (m', \pi')$  for  $m' = m$ . Now let  $((m, g), \delta)$  be any AJO of  $\text{GSTM}_{J_{\text{Msg}}, P_s}(m, \pi_{\text{Msg}})$  and assume for the sake of contradiction that  $g \in \{0, 1\}$ . Then the unfolded view of  $((m, g), \delta)$  contains an AJO for  $(\text{NoMsg}, 0)$  or  $(m', 1)$  from  $\text{STMA}_{J, P_j}$ . By definition of  $\mathcal{A}$  the unfolded view of either of these AJOs will contain an AJO for  $y = (\text{NoMsg}, \pi')$  from  $\text{STMA}_{J_{\text{Msg}}, P_s}$ , a contradiction.  $\square$

**Lemma 10 (Graded Agreement).** Protocol  $\text{GSTM}_{J_{\text{Msg}}, P_s}$  in Figure 6 has graded agreement.

*Proof.* Suppose we have AJOs  $((m_i, g_i), \delta_i)$  and  $((m_j, g_j), \delta_j)$  with  $g_i, g_j > 0$ . By the properties of the generic justifier, it must be the case that the unfolded view of  $((m_i, g_i), \delta_i)$  holds an AJO for  $m_i \neq \text{NoMsg}$  from some  $\text{STMA}_{J, P_k}$ . By the justified message property of  $\text{STMA}$  it follows that  $m_i$  is justified using the input justifier  $J$  of  $\text{STMA}_{J, P_k}$ , which was the genetic justifier checking that  $m_i$  was an AJO of  $\text{STMA}_{J_{\text{Msg}}, P_s}$ . Symmetrically, we can conclude that  $m_j$  was an AJO of  $\text{STMA}_{J_{\text{Msg}}, P_s}$ . By the agreement property of  $\text{STMA}_{J_{\text{Msg}}, P_s}$  it follows that  $m_i = m_j$ . It remains to show that it cannot be the case for two AJOs that  $g_i = 2$  and  $g_j = 0$ . Assume that  $g_i = 0$ . Suppose then for the sake of contradiction that we have AJOs  $((m_i, g_i), \delta_i)$  and  $((m_j, g_j), \delta_j)$  with  $g_i = 0$  and  $g_j = 2$ . The unfolded view of  $((m_i, g_i), \delta_i)$  holds an AJO for  $m_i = \text{NoMsg}^{(s)}$  from all  $\text{STMA}_{J, P_k}$  which did not output  $\text{NoMsg}$ , in particular for all honest  $P_k$ . Symmetrically, we can conclude that  $m_j$  was an AJO from  $\text{STMA}_{J, P_k}$  for all honest  $P_j$ . Since  $g_j = 2$  we have  $m_j \neq \text{NoMsg}^{(s)}$ . Since there is at least one honest  $P_j$  we have found one  $\text{STMA}_{J, P_k}$  where  $P_k$  is honest and it has AJO  $\text{NoMsg}^{(s)}$  and AJO  $m_j \neq \text{NoMsg}^{(s)}$ . This breaks validity of  $\text{STMA}_{J, P_k}$  as the honest  $P_k$  cannot have inputted both  $\text{NoMsg}^{(s)}$  and  $\neq \text{NoMsg}^{(s)}$ .  $\square$

## 7 Early Stopping Broadcast

We now present our main result, early stopping broadcast.

### 7.1 Diagonal Cast

We begin by building an early stopping protocol, called DC (for *diagonal cast*), which is early stopping, but may run up to  $O(t^2)$  rounds. The protocol will be a justifiable broadcast in the sense of Definition 6, which implies that it is also a broadcast protocol. The protocol DC uses **GSTM** as (blackbox) sub-protocol. Since we use DC as a building block in our final protocol, we also make its output justified. During the protocol we use a helper function computing a party's next vote. In each round the parties run one **GSTM** to get output  $(m, g)$ . Consider a party which in the previous rounds saw outputs  $(m^1, g^1), \dots, (m^r, g^r)$ . Then it should use the  $m$  from the latest **GSTM** with  $g^\rho > 0$ . If no such  $\rho$  exists it should use  $\text{NoMsg}^{(s)}$  to indicate the sender  $P_s$  was corrupt. More formally we use this function:

$$\text{NxtInp}((m^1, g^1), \dots, (m^r, g^r)) = \begin{cases} \text{NoMsg}^{(s)} & \text{if } g^1 = \dots = g^r = 0 \\ m_{\max\{i \in [r] \mid g_i > 0\}} & \text{otherwise.} \end{cases}$$

Below we let  $i^* = \max\{i \in [r] \mid g_i > 0\}$  when we are not in the case  $g^1 = \dots = g^r = 0$ . At the end of the protocol we want to map  $\text{NoMsg}^{(s)}$  to  $\text{NoMsg}$ . For this we use a simple helper:  $\text{Out}(m) = \text{NoMsg}$  if  $m = \text{NoMsg}^{(s)}$  and  $\text{Out}(m) = m$  otherwise. The protocol is given in Fig. 7.

**Lemma 11 (Validity).** *DC satisfies validity.*

*Proof.* Suppose that  $P_1$  is an honest sender holding input  $(m, \pi_{\text{Msg}})$ . Since  $J_{\text{Msg}}(P_1, m, \pi_{\text{Msg}}) = \top$  validity of  $\text{GSTM}_{J_{\text{Msg}}, P_1}$  implies that all honest parties get output  $(m, 2, \pi)$ . Hence, all parties terminate with output  $m$  after running for one more iteration of the loop.  $\square$



**Diagonal Cast**  $DC_{J_{\text{Msg}}, P_s}$

Without loss of generality, assume that  $P_1$  is the sender, i.e.,  $s = 1$ .

- In round 1:
  - $P_1$ : Has input  $(m, \pi_{\text{Msg}})$  with  $J_{\text{Msg}}(P_1, m, \pi_{\text{Msg}}) = \top$  and inputs it to  $\text{GSTM}_{J_{\text{Msg}}, P_1}$ .
  - $P_i$ : Let  $y = ((m_i^1, g_i^1), \pi_i^1)$  be the output from  $\text{GSTM}_{J_{\text{Msg}}, P_1}$ .
  - $P_i$ : If  $g_i^1 = 2$  then send  $(1, (m_i^1, g_i^1), \pi_i^1)$  to all other parties, output  $m = \text{Out}(m_i^1)$ , run for one more round (to ensure that no honest party terminates until all honest gave output), and then terminate.
- For  $j \in \{2, \dots, n\}$  do as follows:
  - $P_j$ : Compute  $m_j = \text{NxtInp}((m_j^1, g_j^1), \dots, (m_j^{j-1}, g_j^{j-1}))$ .
  - $P_j$ : Compute the generic justifier  $\gamma_j = \{((m_j^k, g_j^k), \pi_j^k)\}_{k=1}^{j-1}$  and let  $J^j$  be the generic justifier predicate for this round.
  - $P_j$ : Input  $(m_j, \gamma_j)$  to  $\text{GSTM}_{J^j, P_j}$ .
  - $P_i$ : Let  $y = ((m_i^j, g_i^j), \pi_i^j)$  be the output from  $\text{GSTM}_{J^j, P_j}$ .
  - $P_i$ : If  $g_i^j = 2$  then send  $(j, (m_i^j, g_i^j), \pi_i^j)$  to all other parties, output  $m = \text{Out}(m_i^j)$ , run for one more round, and then terminate.
- $P_i$ : Upon receiving a tuple  $(j, (m', g'), \pi')$  (for any  $j \in [n]$  and in any round) such that  $\text{GSTM}_{J^j, P_j} \cdot J_{\text{Out}}(P_i, (m', g'), \pi') = \top$  and  $g' = 2$ , forward it to all parties, output  $m = \text{Out}(m')$ , run for one more round, and then terminate.
- The output justifier for the protocol is  $DC_{J_{\text{Msg}}, P_s} \cdot J_{\text{Out}}(m, \pi)$ , which is true if  $\pi = (j, (m', g'), \pi')$  and  $j \in [n]$  and  $\text{GSTM}_{J^j, P_j} \cdot J_{\text{Out}}(P_i, (m', g'), \pi') = \top$  and  $m = \text{Out}(m')$ .

**Fig. 7.** The broadcast protocol.

**Lemma 12 (Stabilisation).** *Let  $j$  be the minimal iteration such that for some honest  $P_i$  an AJO  $((m, 2), \pi)$  satisfying  $\text{GSTM}_{J^j, P_j} \cdot J_{\text{Out}}(P_i, (m, 2), \pi) = \top$  can be produced. Then for any  $j' \geq j$  and any honest  $P_{i'}$  and any AJO  $((m', 2), \pi')$  satisfying  $\text{GSTM}_{J^{j'}, P_{j'}} \cdot J_{\text{Out}}(P_{i'}, (m', g'), \pi') = \top$  it holds that  $g' = 0$  or  $m' = m$ . Furthermore, if  $j' = j$  then  $g' > 0$  and thus  $m' = m$ .*

*Proof.* We do induction in  $|j' - j| = 0, 1, \dots$ . For the base case  $|j' - j| = 0$  we have  $j' = j$  and that  $g' \in \{1, 2\}$  and  $m' = m$  by graded agreement of  $\text{GSTM}_{J^j, P_j}$ . Assume then the induction hypothesis for all  $|j' - j| < \ell$ . We prove it for  $|j' - j| = \ell$ . We have to prove that  $g' = 0$  or  $m = m'$ . So it is enough to prove that if  $g' > 0$  then  $m' = m$ . When  $g' > 0$  then by the justified message property of  $\text{GSTM}_{J^{j'}, P_{j'}}$  we have that  $m'$  is justified by  $J^{j'}$ , which was the generic justifier for iteration  $j'$ . Therefore  $m'$  was computed as

$$m' = \text{NxtInp}((m^1, g^1), \dots, (m^{j'-1}, g^{j'-1}))$$

from AJOs. By induction hypothesis  $g^j \geq 1$  and  $m^j = m$ , and for  $j \leq k \leq j' - 1$  it holds that  $g^k = 0$  or  $m^k = m$ . This by construction gives

$$\text{NxtInp}((m^1, g^1), \dots, (m^{j'-1}, g^{j'-1})) = m ,$$

as desired. □

**Corollary 2.** *DC has agreement.*

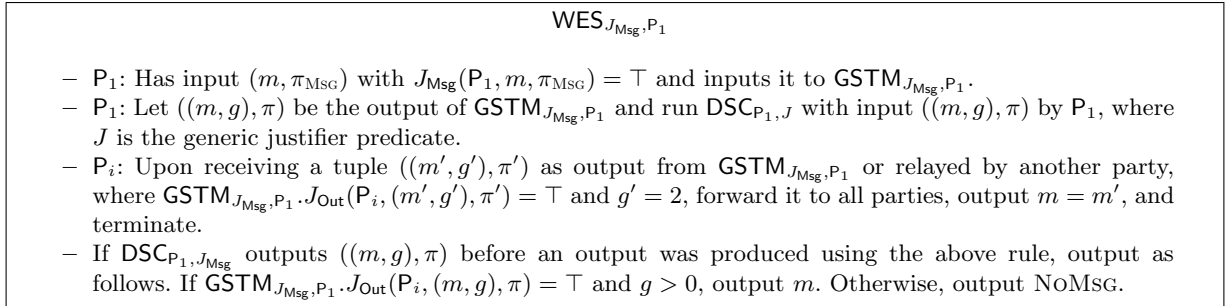
*Proof.* When an honest party  $P_i$  produces output  $m'$  then it by construction produces or receives an AJO  $(m', g = 2)$  for some  $\text{GSTM}_{J_{j'}, P_{j'}}$ . Clearly there is a smallest  $j$  for which an AJO  $(m, g = 2)$  can be produced for  $\text{GSTM}_{J_j, P_j}$ . By Lemma 12 it holds that  $m' = m$ . This holds for all honest outputs  $m'$ .  $\square$

**Lemma 13.** *DC terminates in  $8(f + 1)(f + 2)$  rounds. If  $P_1$  is honest then it runs in at most  $8(f + 2)$  rounds.*

*Proof.* The protocol terminates at the latest once the first honest party acts as the sender in some iteration  $j$  as this gives  $g_j = 2$ . This is worstcase in iteration  $(f + 1)$ . Each iteration runs one  $\text{GSTM}$ , which uses two  $\text{STMA}$ , which each uses two  $\text{STM}$ , which each uses  $f + 2$  rounds. This gives  $4(f + 2)$  rounds per  $\text{GSTM}$ , and applying the staggering compiler contributes a factor of 2.  $\square$

## 7.2 Weak Early Stopping $\mathcal{O}(f)$ and Worstcase $\mathcal{O}(t)$

The protocol DC can do early stopping, but if there are  $f = \omega(\sqrt{t})$  corruptions it runs for more than  $\mathcal{O}(t)$  rounds, which is asymptotically sub-optimal. We solve this by capping the running time at  $\mathcal{O}(t)$ . Doing this safely is subtle, and we now present a protocol with weak early stopping which helps doing it safely. Weak early stopping means that the protocol achieves early stopping when the sender is honest. If the sender is corrupt it may run for  $\mathcal{O}(n)$  rounds. We describe the protocol with  $P_1$  as sender, but it can be adopted to any  $P_s$ .



**Fig. 8.** The weak early stopping broadcast protocol.

**Theorem 4.** *The protocol  $\text{WES}_{J_{\text{Msg}}, P_1}$  is a broadcast protocol. Parties terminate one round apart and the round complexity is  $\mathcal{O}(t)$ . If  $P_1$  is honest then the round complexity is  $\mathcal{O}(f)$ .*

*Proof.* Validity is trivial: if  $P_1$  is honest then  $\text{GSTM}_{J_{\text{Msg}}, P_1}$  outputs  $((m, g), \pi)$  with  $g = 2$  and all honest parties output  $m$ . This happens within one run of  $\text{GSTM}$ , so in  $\mathcal{O}(f)$  rounds. We argue agreement. If all honest parties give output by receiving a tuple  $((m', g'), \pi')$  which is an AJO for  $\text{GSTM}_{J_{\text{Msg}}, P_1}$  and with  $g' = 2$  then agreement follows from agreement of  $\text{GSTM}_{J_{\text{Msg}}, P_1}$ . If all honest parties give output by receiving  $((m, g), \pi)$  from DSC then agreement follows from agreement of DSC. Assume then that some honest  $P_i$  gives output

using  $((m' = m_i, g' = 2), \pi')$  which is an AJO for  $\text{GSTM}_{J_{\text{Msg}}, P_1}$  and some honest  $P_j$  gives output  $m_j$  by receiving  $((m, g), \pi)$  from  $\text{DSC}_{P_1, J}$ . Since  $g' = 2$  and  $m' = m_i$  it follows from graded agreement for  $\text{GSTM}_{J_{\text{Msg}}, P_1}$  that  $m = m_i$  and  $g > 0$  for all AJOs for  $\text{GSTM}_{J_{\text{Msg}}, P_1}$ . Since  $J$  is the generic justifier it follows that all AJOs  $((m, g), \pi)$  for  $\text{DSC}_{P_1, J}$  has  $m'' = m_i$  and  $g'' > 0$ . Therefore  $P_j$  has output  $m_j = m = m_i$ . Parties terminate one round apart as they terminate one round apart in  $\text{DSC}$  and if they terminate by the  $g' = 2$  rule then they forward the AJO and then all honest parties terminate in the next round.  $\square$

### 7.3 Early Stopping $\mathcal{O}(f^2)$ and Worstcase $\mathcal{O}(t)$

We now give a broadcast protocol Capped Diagonal Cast which basically runs Diagonal Cast for a capped number of rounds and uses WES to have all parties report whether they saw an output from DC before the time cap. Since the reports are sent with WES parties will agree on the reports and make the same decision. Furthermore, if an honest party saw an output it will be reported with early stopping. We again describe the protocol with  $P_1$  as sender, but can trivially adopt to any  $P_s$ .

**Capped Diagonal Cast**  $\text{CDC}_{P_1, J_{\text{Msg}}}$

We describe the protocol from the view of party  $P_i$ .

- $P_1$ : Has input  $(m, \pi_{\text{MSG}})$  with  $J_{\text{Msg}}(m, \pi_{\text{MSG}}) = \top$  and inputs it to  $\text{DC}_{P_1, J_{\text{Msg}}}$ .
- $P_i$ : Participate in  $\text{DC}_{P_1, J_{\text{Msg}}}$  for at most  $8(t+1)$  rounds.<sup>a</sup>
- $P_i$ : If  $\text{DC}_{P_1, J_{\text{Msg}}}$  produced output  $(m, \pi)$  in or before round  $8n$  then run  $\text{WES}_{P_i, \top}$  with input  $(m, \pi)$  in the round where  $\text{DC}_{P_1, J_{\text{Msg}}}$  produced output. If by round  $8(t+1)$  protocol  $\text{DC}_{P_1, J_{\text{Msg}}}$  did not produce output run  $\text{WES}_{P_i, \top}$  with input  $\text{NoMSG}$  in round  $8(t+1)$ .
- $P_i$ : If and when the first  $\text{WES}_{P_j, \top}$  outputs  $(m, \pi)$  such that  $\text{DC}_{P_1, J_{\text{Msg}}}.J_{\text{Out}}(m, \pi) = \top$ , output  $m$ .
- $P_i$ : If  $\text{WES}_{P_1, \top}, \dots, \text{WES}_{P_n, \top}$  all terminated and none had an output  $(m, \pi)$  such that  $\text{DC}_{P_1, J_{\text{Msg}}}.J_{\text{Out}}(m, \pi) = \top$ , then output  $\text{NoMSG}$ .

---

<sup>a</sup> Here we count base communication rounds, not iterations of  $\text{GSTM}$ .

**Fig. 9.** An Early Stopping Broadcast protocol with  $\mathcal{O}(\min((f+1)^2, n))$  rounds.

**Theorem 5.** *The protocol  $\text{CDC}_{J_{\text{Msg}}, P_1}$  is a broadcast protocol. Parties terminate one round apart and the round complexity is  $\mathcal{O}(\min(f^2, t))$ . If the sender is honest output is given in  $\mathcal{O}(f)$  rounds. From  $\text{CDC}_{J_{\text{Msg}}, P_1}$  we can get a broadcast protocol  $\text{EBC}$  for definition Definition 1 with the same communication complexity simply by dropping the input justifier predicate  $J_{\text{Msg}}$  and the justifier  $\pi_{\text{MSG}}$ .*

*Proof.* First note that all  $\text{WES}_{P_i, \top}$  are started at most one round apart. Namely, no party starts them later than by round  $8(t+1)$ . So if they are to be started 2 rounds apart the first is started in round  $8(t+1) - 2$  or earlier. But then  $\text{DC}_{P_1, J_{\text{Msg}}}$  terminated by round  $8(t+1) - 2$  at the first honest party. But then it terminated by round  $8(t+1) - 1$  at all honest. Hence all honest started  $\text{WES}_{P_i, \top}$  exactly when  $\text{DC}_{P_1, J_{\text{Msg}}}$  terminated, which is at most one round apart. We can then apply the staggering compiler to ensure that  $\text{WES}_{P_i, \top}$  tolerates being started

one round apart. This gives a  $\mathcal{O}(1)$  blowup in round complexity. We then argue validity: if  $P_1$  is honest then  $DC_{J_{\text{Msg}}, P_1}$  outputs  $(m, \pi)$  within  $8(f + 2) \leq 8(t + 2)$  rounds. Therefore an honest party  $P_j$  will send  $(m, \pi)$  on  $WES_{P_j, \tau}$  which will output  $(m, \pi)$  within  $\mathcal{O}(f)$  rounds and then all honest output  $m$ . This all happens within  $\mathcal{O}(f)$  rounds. Agreement is trivial from agreement of  $WES$ , as the output is computed deterministically from the outputs of  $WES_{P_1, \tau}, \dots, WES_{P_n, \tau}$ . Consider then the round complexity. We have that all copies of  $WES$  are started after  $\mathcal{O}(\min(f^2, t))$  rounds as  $DC_{P_1, J_{\text{Msg}}}$  stops after  $\mathcal{O}(f^2)$  rounds and we cap after  $8(t + 1)$  rounds. If  $DC_{P_1, J_{\text{Msg}}}$  did give an output at all honest parties before round  $8(t + 1)$  then it will be input to  $WES_{P_i, \tau}$  which will output after  $\mathcal{O}(f)$  rounds, and hence output is produced in  $\mathcal{O}(f^2)$  rounds as  $DC_{P_1, J_{\text{Msg}}}$  terminates in  $\mathcal{O}(f^2)$  rounds. If  $DC_{P_1, J_{\text{Msg}}}$  did not give output within  $8(t + 1)$  rounds then  $8(t + 1) = \mathcal{O}(f^2)$ , so  $\mathcal{O}(\min(f^2, t)) = \mathcal{O}(t)$ . And in this case the overall protocol runs for at most  $\mathcal{O}(t)$  rounds as each  $WES$  terminates in  $\mathcal{O}(t)$  rounds.  $\square$

## 8 Acknowledgements

We would like to thank Juan Garay for fruitful discussions in the early stages of this work. We would also like to thank Fatima Elsheimy for many detailed technical and editorial comments.

## References

1. Ittai Abraham, T.-H. Hubert Chan, Danny Dolev, Kartik Nayak, Rafael Pass, Ling Ren, and Elaine Shi. Communication complexity of byzantine agreement, revisited. In Peter Robinson and Faith Ellen, editors, *38th ACM PODC*, pages 317–326. ACM, July / August 2019.
2. Ittai Abraham and Danny Dolev. Byzantine agreement with optimal early stopping, optimal resilience and polynomial complexity. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 605–614. ACM Press, June 2015.
3. Andreea B. Alexandru, Julian Loss, Charalampos Papamanthou, and Giorgos Tsimos. Sublinear-round broadcast without trusted setup against dishonest majority. Cryptology ePrint Archive, Report 2022/1383, 2022. <https://eprint.iacr.org/2022/1383>.
4. Piotr Berman and Juan A. Garay.  $n/4$ -resilient distributed consensus in  $t + 1$  rounds. *Mathematical Systems Theory*, 26(1):3–19, 1993.
5. Piotr Berman, Juan A. Garay, and Kenneth J. Perry. Towards optimal distributed consensus (extended abstract). In *30th FOCS*, pages 410–415. IEEE Computer Society Press, October / November 1989.
6. Piotr Berman, Juan A. Garay, and Kenneth J. Perry. Optimal early stopping in distributed consensus (extended abstract). In *WDAG*, pages 221–237, 1992.
7. T.-H. Hubert Chan, Rafael Pass, and Elaine Shi. Sublinear-round byzantine agreement under corrupt majority. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part II*, volume 12111 of *LNCS*, pages 246–265. Springer, Heidelberg, May 2020.
8. Tushar Deepak Chandra and Sam Toueg. Time and message efficient reliable broadcasts. In *WDAG*, pages 289–303, 1990.
9. Danny Dolev and Christoph Lenzen. Early-deciding consensus is expensive. In Panagiota Fatourou and Gadi Taubenfeld, editors, *32nd ACM PODC*, pages 270–279. ACM, July 2013.
10. Danny Dolev and Rüdiger Reischuk. Bounds on information exchange for byzantine agreement. In Robert L. Probert, Michael J. Fischer, and Nicola Santoro, editors, *1st ACM PODC*, pages 132–140. ACM, August 1982.
11. Danny Dolev, Rüdiger Reischuk, and H. Raymond Strong. Early stopping in byzantine agreement. *J. ACM*, 37(4):720–741, 1990.

12. Danny Dolev and H. Raymond Strong. Authenticated algorithms for Byzantine agreement. *SIAM Journal on Computing*, 12(4):656–666, 1983.
13. Danny Dolev and Andrew Chi-Chih Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–207, 1983.
14. Paul Feldman and Silvio Micali. Optimal algorithms for byzantine agreement. In *20th ACM STOC*, pages 148–161. ACM Press, May 1988.
15. Matthias Fitzi and Jesper Buus Nielsen. On the number of synchronous rounds sufficient for authenticated byzantine agreement. In *DISC*, pages 449–463, 2009.
16. Juan A. Garay, Jonathan Katz, Chiu-Yuen Koo, and Rafail Ostrovsky. Round complexity of authenticated broadcast with a dishonest majority. In *48th FOCS*, pages 658–668. IEEE Computer Society Press, October 2007.
17. Juan A. Garay and Yoram Moses. Fully polynomial byzantine agreement in  $t+1$  rounds. In *25th ACM STOC*, pages 31–41. ACM Press, May 1993.
18. Juan A. Garay and Yoram Moses. Fully polynomial byzantine agreement for  $n \geq 3t$  processors in  $t + 1$  rounds. *SIAM Journal on Computing*, 27(1):247–290, 1998.
19. Jonathan Katz and Chiu-Yuen Koo. On expected constant-round protocols for byzantine agreement. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 445–462. Springer, Heidelberg, August 2006.
20. Idit Keidar and Sergio Rajsbaum. A simple proof of the uniform consensus synchronous lower bound. *Information Processing Letters*, 85(1):47–52, 2003.
21. Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *ACM Trans. Programming Languages and Systems*, 4(3):382–401, 1982.
22. Christoph Lenzen and Sahar Sheikholeslami. A recursive early-stopping phase king protocol. In *ACM PODC*, pages 60–69. ACM, 2022.
23. Yehuda Lindell, Anna Lysyanskaya, and Tal Rabin. Sequential composition of protocols without simultaneous termination. In Aletta Ricciardi, editor, *21st ACM PODC*, pages 203–212. ACM, July 2002.
24. Silvio Micali. Very simple and efficient byzantine agreement. In Christos H. Papadimitriou, editor, *ITCS 2017*, volume 4266, pages 6:1–6:1, 67, January 2017. LIPIcs.
25. Jesper Buus Nielsen. *On Protocol Security in the Cryptographic Model*. PhD thesis, Aarhus University, 2003.
26. Philippe Raipin Parvédy and Michel Raynal. Optimal early stopping uniform consensus in synchronous systems with process omission failures. In *SPAA*, volume 302–310, 2004.
27. Kenneth J. Perry and Sam Toueg. An authenticated byzantine generals algorithm with early stopping. Technical report, Cornell University, 1984.
28. Shravan Srinivasan, Julian Loss, Giulio Malavolta, Kartik Nayak, Charalampos Papamanthou, and Sri Aravinda Krishnan Thyagarajan. Transparent batchable time-lock puzzles and applications to byzantine consensus. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023, Part I*, volume 13940 of *LNCS*, pages 554–584. Springer, Heidelberg, May 2023.
29. Jun Wan, Hanshen Xiao, Srinivas Devadas, and Elaine Shi. Round-efficient byzantine broadcast under strongly adaptive and majority corruptions. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part I*, volume 12550 of *LNCS*, pages 412–456. Springer, Heidelberg, November 2020.
30. Jun Wan, Hanshen Xiao, Elaine Shi, and Srinivas Devadas. Expected constant round byzantine broadcast under dishonest majority. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part I*, volume 12550 of *LNCS*, pages 381–411. Springer, Heidelberg, November 2020.