

# Byzantine Agreement Decomposed: Honest Majority Asynchronous Total-Order Broadcast from Reliable Broadcast

Simon Holmgaard Kamp and Jesper Buus Nielsen

Aarhus University\*  
{kamp, jbn}@cs.au.dk

**Abstract.** It is well-known that Asynchronous Total Order Broadcast (ATOB) requires randomisation and that at most  $t < n/3$  out of  $n$  players are corrupted. This is opposed to synchronous total-order broadcast (STOB) which can tolerate  $t < n/2$  corruptions and can be deterministic. We show that these requirements can be conceptually separated, by constructing an ATOB protocol which tolerates  $t < n/2$  corruptions from blackbox use of Common Coin and Reliable Broadcast. We show the power of this conceptually simple contribution by reproving, using simpler protocols, existing results on STOB with optimistic responsiveness and asynchronous fallback. We also use the framework to prove the first ATOB with sub-quadratic communication and optimal corruption threshold  $t < n/3$ , new ATOBs with covert security and mixed adversary structures, and a new STOB with asymmetric synchrony assumptions.

## 1 Introduction

MPC in asynchronous networks is complicated by the fact that a protocol can never wait to receive inputs from all parties in the presence of faults. [BCG93, Can95] introduce a model for Asynchronous MPC (AMPC) in which the ideal functionality is a trusted party who receives only the inputs from a *core set* of the parties and returns a description of the core set and the result of the computation using only these inputs. The core set is chosen by an adversarial scheduler but is restricted to contain at least  $n - t$  parties chosen independently of any inputs of honest parties. To realise this functionality they provide an Agreement on a Core Set (ACS) protocol which is resilient against  $t < n/3$  corruptions and allows the parties to agree on a set of size at least  $n - t$ . In more detail each party at the beginning of the protocol has a set of size  $n - t$  which they send to all other parties. It keeps adding new elements to the set as they arrive and in each round of the protocol waits until  $n - t$  of the sets from the previous round are included in its currently accumulated set, then sends its current set. After  $\log n$  rounds it can be shown that all the accumulated sets share a common core of size  $n - t$ . The proof goes by induction in the number of rounds and crucially relies on the fact that in each round each party takes the union of sets from  $n - t = 2t + 1$  parties to conclude that each pair of honest parties will receive at least one set from some common honest party  $P$ , and hence their set in round  $r$  will include the set sent by  $P$  in round  $r - 1$ . Finally each party runs a Byzantine Agreement (BA) protocol for each party with input 1 iff that party is in its accumulated set.

Several works stop just before running the BAs and instead follow the pattern of rounds of all-to-all communication while accumulating local sets with a large common core as a building block to achieve consensus. Following [AJM<sup>+</sup>21] we will refer to this design pattern as *Gather*. One notable recent example using Gather is the DAG-Rider protocol in [KKNS21], in which they build

---

\* Simon Holmgaard Kamp and Jesper Buus Nielsen are partially funded by The Concordium Foundation.

a DAG of blocks pointing to  $n - t$  blocks from the previous round, and for each 4 rounds, called a *wave*, rely on the fact that causal past of any  $n - t$  blocks in the last round include  $n - t$  common blocks from the first round. This allows them to conclude that a block from the first round of the wave which is randomly chosen after the wave is completed will be in the DAG of an honest party with probability at least  $2/3$ . The reason they get away with only 4 rounds of blocks in each wave and still get a large common core, is that all blocks are sent through Reliable Broadcast[Bra87] (RB). This means that each pair of blocks will have overlap on least  $n - 2t \geq t + 1$  of the blocks referenced from the previous round. This gives rise to an ATOB protocol by running multiple waves in sequence, flipping a common coin for each iteration and applying an appropriate commit rule.

In this paper we present an asynchronous Gather protocol tolerating  $t < n/2$  corruptions assuming RB. On a high level it uses the structure of the Gather in [KKNS21] and the round complexity of [BCG93,Can95] to tolerate more corruptions and a termination condition from the Gather protocol CSS in [MMNT19,DMM<sup>+</sup>20] to get constant time complexity. Interestingly, even though the protocol follows [BCG93,Can95] and runs in  $\log(n)$  rounds of communication it will terminate in time  $\mathcal{O}(\Delta_{\text{NET}})$ , where  $\Delta_{\text{NET}}$  is the actual network delivery time during the execution. It uses the RB blackbox, so one can instantiate RB in a setting where assumption  $A$  holds to get Gather for the setting where  $t < n/2$  and  $A$  holds. Using a generic transformation similar to [KKNS21] this gives a Total Order Broadcast (TOB) protocol for the same setting. We note that additional work is needed to get ATOB from Gather compared to previous works as all sub-protocols must tolerate  $t < n/2$  in the asynchronous model. Typically, if  $A$  implies RB it also implies  $t < n/2$ , so overall we get that if we can construct RB from assumption  $A$  then we can also get TOB from  $A$ . Our compiler is efficient, so this allows to construct efficient TOB in various models. In summary we get the following results:

1. In the asynchronous setting in which we assume a ground population of size  $N$  with  $T < N/3$  corruptions and a committee of  $n$  servers with honest majority we construct a protocol for RB which has communication linear in  $N$ . Using this and subset sampling to elect  $n$  servers from the ground population we get the first ATOB protocol with subquadratic communication and *optimal* resilience.
2. In the asynchronous setting we construct a RB that is covert secure against  $t < n$  corruptions. This gives an ATOB protocol with covert security against  $t < n/2$  corruptions.
3. We introduce a new weakly synchronous model, called the *asymmetric synchrony assumption* (ASA) model, where each party  $P_i$  has its own guess  $\Delta_{\text{GUESS}}^i$  at the network delay. The guess is unknown to the other parties. The only guarantee is that the actual network delay for messages sent specifically to  $P_i$  is smaller than  $\Delta_{\text{GUESS}}^i$ . A motivation for the model is *ad hoc* groups of parties which want to do TOB and might not share the same assumptions on the network, or know each other's assumptions prior to running the protocol. We give an RB for the ASA model tolerating  $t < n/2$  corruptions. We also consider the weakly ASA (WASA) model, where all honest guesses at the network delay are sound for all messages sent between all honest parties.
4. We introduce the *timeout model*, which uses time weakly. Parties do not have clocks, their only access to time is that they can set a timeout and get a callback *at least* some  $\Delta_{\text{WAIT}}$  seconds later. They are also not guaranteed to start the protocol at the same time, which makes it easier to compose the protocol with other protocols. This notion of synchrony is seemingly much easier to implement in practice as it does not need synchronised clocks nor very precise clocks, merely a bound on how fast the clock drift. We give an RB for the timeout model tolerating  $t < n/2$  corruptions.

5. In the synchronous setting we instantiate an RB which is optimistically responsive, i.e., when the actual corruption is lower than the maximal tolerated corruption then the latency of the protocol only depends on the current network delay. Our protocol tolerates  $t \leq t_s$  corruptions and is optimistic for  $t \leq t_A (\leq t_s)$  corruptions. It tolerates  $t_s + 2t_A < n$  which is optimal. This gives an optimistic responsive TOB in which the optimistic condition does not include having an honest leader. The protocol is secure in the WASA model and the timeout model. To the best of our knowledge the resulting TOB is the first optimally resilient fallback TOB for the WASA model and for the timeout model.
6. We give a RB for the ASA model with asynchronous fallback which is secure for  $2t_s + t_A < n$  and  $t_A \leq t_s$  and when *either* there are at most  $t \leq t_s$  corruptions and the network is ASA synchronous *or* there are at most  $t \leq t_A$  corruptions and the network is fully asynchronous. To the best of our knowledge the resulting TOB is the first optimally resilient fallback TOB for the ASA model and for the timeout model. Setting  $t_A = 0$  gives the contributions claimed in Item 3 and Item 4.
7. We finally note that it is an easy corollary of our framework that you can get asynchronous TOB for a model with mixed adversaries, where there are at most  $t_{\text{Byz}}$  fully Byzantine corruptions, at most  $t_{\text{CRASH}}$  additional crash-silent errors, and  $2t_{\text{CRASH}} + 3t_{\text{Byz}} < n$ .

The above results show the power of a framework where ATOB is implemented with *honest majority* given RB, which can then be implemented in different ways.

## 1.1 Related work

We discuss related work in the context of the specific contributions.

*Subquadratic ATOB with Optimal Resilience* A noteworthy property of asynchronous protocols tolerating significantly more than a third corruption, is that one can securely sample a committee to perform the task while retaining the optimal resilience against  $t < n/3$  corruptions. This was previously used in [ACKN23]. Previous solutions to subquadratic BA in the asynchronous [BKLL20,CKS20] or partially synchronous [GHM<sup>+</sup>17] setting rely on sampling committees with an honest supermajority from a ground population with  $t < (1 - \epsilon)n/3$  corruptions. As  $\epsilon$  can be arbitrarily small, the separation between optimally and “near-optimally” resilient protocols might seem purely theoretical. But the practical implications of picking a small  $\epsilon$  is that any secure protocol must have a concretely very large committee. If for instance  $\epsilon = 1/10$  (i.e. the corruption in the ground population is bounded by 30%) committees need 16037 parties to retain an honest supermajority except with probability  $2^{-60}$  (cf. table 1 of [DMM<sup>+</sup>22]). A recent work [BBK<sup>+</sup>23] on the concrete security of Algorand show that a committee of 6000 parties is needed to get 56 bits of security assuming the parties are sampled from a ground population with less than 1/5 corruption. When assuming the same corruption in the ground population our protocol would need a committee of 173 parties to get 60 bits of security, while it would need 653 parties to tolerate up to a third being corrupted (cf. Table 1).

To instantiate our protocol we need to use a RB, which usually would have all-to-all communication in a committee with an honest supermajority. However, we construct a RB protocol which only depends linearly on the size of the ground population that needs honest supermajority while the remaining communication is done in a committee with honest majority. As a rough sketch: a sender sends a message to the ground population who threshold signs it and sends their shares

back to the sender who then combines the shares and sends it to the committee that will gossip it among themselves. Our actual implementation is a bit more involved in order to get not only subquadratic but concretely very efficient communication which can be amortized optimal when either the messages or the ground population is large. This approach is not secure against an adaptive adversary and it is an open problem if one can get an adaptively secure subquadratic asynchronous agreement with optimal resilience.

The DAG rider protocol [KKNS21] can achieve amortized communication overhead of  $\mathcal{O}(n)$  when blocks include  $\mathcal{O}(n \log n)$  values, i.e.,  $\mathcal{O}(n \log n \kappa)$  bits. This is done by instantiating it with the RB protocol by Cachin and Tessaro [CT05] which communicates  $\mathcal{O}(n|m| + n^2 \log n \kappa)$  bits to broadcast a message of length  $|m|$ . An important caveat of the amortized complexity of [KKNS21] is that it only goes through in their model because there is no distinction between clients and servers. This means that each value sent from a server counts as throughput. If the protocol would be deployed in a setting in which clients send their messages to servers who broadcast them, then communication complexity increases by a factor  $n$ , as all servers could in the worst case include the same messages in each round. Our RB protocols uses similar techniques as in [CT05] to get the communication cost of broadcasting  $\beta$  bits from and to the ground population down to the amortized optimal  $\mathcal{O}(N\beta)$  when the protocol consumes on average at least  $n \log(n)$  messages per epoch and *either* the size of the ground population is at least  $n^2 \log n$  or the average messages size is at least  $n \log(n) \kappa$ . The time complexity is expected  $\mathcal{O}(1)$  network delays as in [KKNS21].

*Honest Majority Asynchronous Total-Order Broadcast with Covert Security* [AL07] introduces a notion of *covert adversaries* in which some parties are Byzantine corrupted but do not want to be “caught”. Protocols are said to be covert secure if any breach of security will result in the honest parties detecting a specific party who did not follow the protocol. [AO12] introduce a strengthened model in which cheating parties are not only detected by an honest party with some probability, but that party also receives a certificate demonstrating that a party cheated. We informally prove that we can instantiate asynchronous RB with dishonest majority and covert security with public verifiability, which in turn gives ATOB for honest majority secure against covert adversaries. The idea is simple: senders sign their messages which are then flooded. If different messages were signed, they will eventually meet and the sender is detected as corrupted.

*Optimistic Responsiveness* Optimistic responsive protocols [PS17] have safety in synchronous networks when the number of corruptions is bounded by  $t_s$  but additionally a latency that only depends on the actual (unknown) network delay when at most  $t_A < t_s$  parties are corrupted. In some cases additional optimistic conditions, e.g., a leader being honest, must be met. We provide an optimistic responsive RB protocol matching the optimal bounds for  $t_A$  and  $t_s$ , which in turn gives an optimistic responsive TOB protocol. As the resulting protocol only retroactively elects a leader, and any valid leader points to a core set of  $n - t$  inputs, the resulting protocol does not impose any restrictions on the optimistic case besides having at most  $t_A$  corruptions. The paper [HKL20] considers asynchronous RB protocols with different corruption levels for the different security properties validity, agreement and termination. This is related, but different as we consider synchronous protocols and different running times for different thresholds.

*Synchronous Consensus with Asynchronous Fallback* The concept of synchronous protocols tolerating  $t_s$  corruptions and additionally tolerating asynchrony when  $t_A < t_s$  parties are corrupted were introduced in [BKL19]. It turns out that the protocol used to achieve optimistic responsiveness has

synchronous security with asynchronous fallback when adjusting the timeout used and instantiating it with the optimal bounds for  $t_A$  and  $t_S$  in this model. Once more this immediately gives a TOB protocol with optimal thresholds in this model.

*Asymmetric Synchrony Assumptions* To the best of our knowledge it is new that one can do TOB for  $t < n/2$  in the ASA model. The closest related work seems to be [DDFN07] where the authors consider asymmetric trust in MPC. Each party has its own assumptions on which subsets of the parties may be corrupted. However, [DDFN07] considers a fully synchronous network proceeding in rounds, i.e., all parties share assumptions on the network delay, start the protocol at the same time and have access to clocks to implement a round-based abstraction. Here we relax the model further and assume that the parties might not even share assumptions on the network delay. We show how to get TOB in this model. Our protocol has the interesting property that it is responsive in the largest *honest*  $\Delta_{\text{GUESS}}^i$ , i.e., the corrupted parties cannot do a denial of service attack by proposing artificially large  $\Delta_{\text{GUESS}}^i$ . Asymmetric trust in distributed systems and specifically total-order broadcast was also studied in [CT19,CZ21,Cac21], but again only the assumptions on corruptions are asymmetric, not assumptions on network delays.

*Mixed Adversaries* Mixed adversaries were studied in consensus and MPC before. In [GP92,MP91] the authors give a consensus protocols for  $t_{\text{CRASH}} + 3t_{\text{BYZ}} < n$ , but for the synchronous model. In [FHM98] the authors give information theoretic MPC protocols which tolerates mixed adversaries. For the asynchronous model  $2t_{\text{CRASH}} + 3t_{\text{BYZ}} < n$  is trivially optimal. An honest party can only wait to hear from  $n - t$  parties without deadlocking, where  $t = t_{\text{BYZ}} + t_{\text{CRASH}}$ . If  $2t_{\text{CRASH}} + 3t_{\text{BYZ}} = n$  then two honest parties each hearing from  $n - t$  parties may only have an overlap of  $n - 2t = t_{\text{BYZ}}$ . Therefore honest parties can be partitioned and never receive information from each other and still have to give an output. We are not aware of any concrete ATOB protocol for a mixed model with  $2t_{\text{CRASH}} + 3t_{\text{BYZ}} < n$ . We do, however, not claim that the feasibility is new. We are highlighting the construction to show how easy it is to derive an ATOB for  $2t_{\text{CRASH}} + 3t_{\text{BYZ}} < n$  in our framework. Getting a RB from  $2t_{\text{CRASH}} + 3t_{\text{BYZ}} < n$  is trivial, see Section 6.3, and the rest of the construction does not consider mixed adversaries at all and works for  $t < n/2$ , i.e.,  $2t_{\text{CRASH}} + 2t_{\text{BYZ}} < n$ .

*Subquadratic AMPC with Optimal Resilience* The final result we present is the fairly simple observation that using our subquadratic ATOB protocol as the broadcast channel in the AMPC protocol in [Coh16] which only requires honest majority assuming such a broadcast channel, we get AMPC which can be run by the honest majority committee and hence has subquadratic complexity with optimal resilience. It has previously been shown that AMPC can be linear in the circuit size[CP15]. [CHL21] focuses on adaptive security and give a protocol with  $\mathcal{O}(n \text{poly}(\kappa))$  communication and near-optimal resilience. But using our approach the communication complexity can be sublinear in the size of the inputs and subquadratic in the size of the output for computations on many inputs. Specifically if at least  $n^2 \log n$  members of the ground population give input of combined size  $\beta$ , and the output has size  $\gamma$  then the communication complexity is  $\mathcal{O}(n\beta + nN\gamma)$ .

## 2 Network Model and Technical Preliminaries

We use  $\kappa$  to denote the security parameter. We use  $\sigma$  to denote a statistical security parameter, i.e., for all fixed  $\sigma$  it holds that the security of a protocol is  $2^{-\sigma} + \text{negl}(\kappa)$  as  $\kappa$  goes to  $\infty$ . In the

main analysis we assume that  $\sigma = \Theta(\kappa)$  and only use  $\kappa$ . We revisit the distinction between the parameters in Section 4 when discussing concrete parameters.

We consider protocols for a fixed set of parties  $\mathbb{G} = \{G_1, \dots, G_N\}$ , which we call the *ground population*. We assume that at most  $T < N/3$  or these are corrupted. We also consider a small set of parties  $\mathbb{S} = \{S_1, \dots, S_n\}$ , which we call the servers. For notational simplicity we assume that  $S_i = G_i$ . We assume that at most  $t < n/2$  or these are corrupted. We offer the reader the intuition that  $\mathbb{S}$  has been sampled from the set  $\mathbb{G}$  large enough that  $T < N/3$  implies that  $t < n/2$ , except with negligible probability. We therefore assume that  $n = \mathcal{O}(\sigma)$ . For simplicity we also assume that  $n = \Theta(\kappa)$  as it allows a more concise analysis in places. In Section 4 we discuss how to set  $n$  to ensure  $t < n/2$  except with probability  $2^{-\sigma}$  for different security levels  $\sigma$ .

We assume that the parties all have pairwise authenticated channels. We assume an asynchronous communication model where the adversary schedules the delivery of the messages without any restrictions. We use a simple model of eventuality. We say that event  $E$  (like termination of a protocol) *eventually* happens (in a protocol with session identifier  $\text{sid}$ ) if it holds that at any point in time if the event  $E$  did not happen for session  $\text{sid}$ , then there is still a message in transit from an honest party to an honest party with session identifier  $\text{sid}$ . Note that a protocol could hack this definition by having two honest parties  $P_1$  and  $P_2$  sending back and forth a PING message forever. Then by definition all events eventually happen. However, all our protocols generate an expected finite number of messages and the simple notion of *eventually* is meaningful for such protocols.

We also consider synchronous protocols. We use the above system model but assume a global clock  $c \in \mathbb{N}$  incremented by the adversary. We say that the network is  $\Delta_{\text{NET}}$ -synchronous if all message sent by time  $c_0$  by an honest party to a honest party is delivered at time  $c_1 \leq c_0 + \Delta_{\text{NET}}$ . We do not give the parties access to the clock. Instead, when describing synchronous protocols we will use time via an explicit timeout mechanism. A party can create a timeout of some duration  $\Delta$ . We say that the party calls  $\text{Timeout}(\text{name}, \Delta)$ . It is then guaranteed that the party at some point in time at least  $\Delta$  time units after the call  $\text{Timeout}(\text{name}, \Delta)$  will be activated with input  $\text{name}$ . It is not guaranteed to happen after exactly duration  $\Delta$ . We do not assume that parties in a protocol start at the same time. We call this the *timeout model*.

We let  $\text{Lg}(n) = \lceil \log_2(n) \rceil$ . We use an erasure code  $\text{EC} = (\text{Enc}, \text{Dec})$ . The encoding of  $m \in \{0, 1\}^{(t+1)\text{Lg}(n)}$  as  $(m_1, \dots, m_n)$  proceeds as follows. Encode  $m$  as  $(\alpha_0, \alpha_1, \dots, \alpha_t) \in \text{GF}(2^{\text{Lg}(n)})^{t+1}$ , let  $f(\mathbf{x}) = \sum_{i=0}^t \alpha_i \mathbf{x}^i$ , and let  $m_j = f(j)$  for  $j = 1, \dots, n$ . We can decode from  $t + 1$  values  $(j, m_j)$  by interpolating  $f(\mathbf{x})$  and reading off  $(\alpha_0, \alpha_1, \dots, \alpha_t)$ . When  $n = 2t + 1$  then clearly  $|m_j| \leq |m|/n$ .

We assume a collision resistant hash function  $\text{Hash}$ . Among other uses we use it to compute a Merkle-tree  $h = \text{HashTree}(m_1, \dots, m_n)$  for messages  $m_i$ . In this context we use  $\text{Path}_i = \text{Path}(m_1, \dots, m_n, i)$  to denote the usual path used to prove that  $m_i$  is consistent with being message number  $i$  in the tree with root  $h$ . We let  $\text{VerPath}$  be the algorithm verifying that  $\text{Path}_i$  proves that  $m_i$  is in  $h$ .

We assume a signature scheme  $(\text{Gen}, \text{Sig}, \text{Ver})$  which is EUF-CMA secure [GMR88] and we assume a PKI. In some initial synchronous round all parties  $G_i \in \mathbb{G}$  sample  $(\text{vk}_i, \text{sk}_i) \leftarrow \text{Gen}(1^\kappa)$ , sends  $\text{vk}_i$  to a trusted third party which makes public  $(\text{vk}_1, \dots, \text{vk}_N)$ . The adversary gets to see  $\text{vk}_i$  for all honest  $G_i$  before picking its own keys, and it does not have to pick its own keys at random, it can use any  $\text{vk}_j$  for a corrupted  $G_j$ .

All our definitions and protocols tacitly assumes that a session identifier  $\text{sid}$  is given to distinguish different runs of a protocol. When signing a message in a session with session identifier  $\text{sid}$  we tacitly add  $\text{sid}$  to the signed message to avoid replay attacks.

We assume some common setup among  $\mathbb{G}$ . Initial some PPT algorithm  $(\text{pv}, \text{sv}_1, \dots, \text{sv}_N) \leftarrow \text{Setup}(1^\kappa)$  is run and  $(\text{pv}, \text{sv}_i)$  is given to  $\mathbb{G}_i$ . The secret values can be correlated, for instance a sharing of a secret key. We do not assume special setup among the servers  $\mathbb{S}$ . The reason is that we think of  $\mathbb{S}$  as having been sampled from  $\mathbb{G}$ , possibly on the fly. Therefore any specialized setup for  $\mathbb{S}$  would reasonably have to be computed on the fly too, and we do not want to hide the cost of this.

We use a threshold signature scheme with unique signatures (**Setup**, **Sig**, **Ver**, **Combine**), where  $(\text{vk}, \text{sk}_1, \dots, \text{sk}_N) \leftarrow \text{Setup}(1^\kappa)$  generates a verification key  $\text{vk}$  and a signing share  $\text{sk}_i$  for  $\mathbb{G}_i$ , **Sig** <sub>$\text{sk}_i$</sub>  partially signs  $m$ , **Ver** can verify a partial signature, and **Combine** computes  $\sigma = \text{Sig}_{\text{sk}}(m)$  from  $T+1$  verified shares, where  $T < N/3$ . Given only  $T$  of the signing keys  $\text{sk}_i$  the scheme is still EUF-CMA. A detailed definition can be found in for instance [CKS00].

For one result we assume a threshold fully homomorphic encryption scheme (**Setup**, **Enc**, **Eval**, **Dec**, **Ver**, **Combine**), where  $(\text{ek}, \text{dk}_1, \dots, \text{dk}_N) \leftarrow \text{Setup}(1^\kappa)$  generates an encryption key  $\text{ek}$  and a decryption share  $\text{dk}_i$  for  $\mathbb{G}_i$ , **Enc** <sub>$\text{ek}$</sub>  encrypts, **Eval** <sub>$\text{ek}$</sub>  $(f, \cdot)$  applies  $f$  to the messages in ciphertexts, **Dec** <sub>$\text{dk}_j$</sub>  $(c) = y_j$  produces a partial decryption, **Ver** verifies a decryption share, and **Combine** compute **Dec** <sub>$\text{dk}$</sub>  $(c)$  from  $T+1$  verified shares, where  $T < N/3$ . Given only  $T$  decryption keys the scheme is IND-CPA. A detailed security definition can be found in [Coh16].

## 2.1 Eventual Justifiers

In our protocols all messages will have a message identifier  $\text{mid}$  specifying which protocol it belongs to, what round of the protocol it comes from, sent by whom and so on. Each message identifier  $\text{mid}$  specifies a party  $\mathbb{P}^{\text{mid}}$ , which we think of as the party which is to send the message identified by  $\text{mid}$ . Here  $\mathbb{P}$  is either a server  $\mathbb{S}$  or a ground member  $\mathbb{G}$ . Each  $\text{mid}$  also specifies a so-called justifier  $J^{\text{mid}}$ , which is a predicate depending on the message  $m$  and the local state of a party. When we write pseudo-code then we write  $J^{\text{mid}}(m)$  to denote that the party  $\mathbb{P}$  executing the code computes  $J^{\text{mid}}$  on  $m$  using its current state. In definitions and proofs we write  $J^{\text{mid}}(m, \mathbb{P}, t)$  to denote that we apply  $J^{\text{mid}}$  to  $m$  and the local state of  $\mathbb{P}$  at time  $t$ . The following definition is adopted from a similar definition in [DMM<sup>+</sup>20].

**Definition 1 (Justifier).** *For a message identifier  $\text{mid}$  we say that  $J^{\text{mid}}$  is a justifier if the following properties hold.*

**Monotone:** *If for an honest  $\mathbb{P}$  and some time  $t$  it holds that  $J^{\text{mid}}(m, \mathbb{P}, t) = \top$  then at all  $t' \geq t$  it holds that  $J^{\text{mid}}(m, \mathbb{P}, t') = \top$ .*

**Propagating:** *If for honest  $\mathbb{P}$  and some point in time  $t$  it holds that  $J^{\text{mid}}(m, \mathbb{P}, t) = \top$ , then eventually the execution will reach a time  $t'$  such that  $J^{\text{mid}}(m, \mathbb{P}', t') = \top$  for all honest parties  $\mathbb{P}'$ .*

**Definition 2 (Justified Protocol).** *We will work with justified protocols II. If a protocol has an input justifier  $J_{\text{IN}}$  it means that a message identifier  $\text{mid}$  is associated with the input,  $J_{\text{IN}} = J^{\text{mid}}$ , and it is  $\mathbb{P}^{\text{mid}}$  which gets the input. Furthermore, if  $\mathbb{P}^{\text{mid}}$  is honest then it is guaranteed that  $J_{\text{IN}}(m) = \top$  whenever  $m$  is input to  $\mathbb{P}^{\text{mid}}$ . If a protocol has an output justifier  $J_{\text{OUT}}$  it means that a message identifier  $\text{mid}$  is associated with the output,  $J_{\text{OUT}} = J^{\text{mid}}$ , it is  $\mathbb{P}^{\text{mid}}$  which gives the output, and when it gives the output it sends it to all parties with message identifier  $\text{mid}$ . Furthermore, if  $\mathbb{P}^{\text{mid}}$  is honest then it is a security property of the protocol that  $J_{\text{OUT}}(m) = \top$  whenever  $m$  is output by  $\mathbb{P}^{\text{mid}}$ .*

We will sometimes talk about a property holding for all *possible justified outputs* of a justified protocol, by which we mean that it holds also for adversarial outputs, as long as these are justified. These outputs might not have been the output of any party in the protocol, the adversary is free to cook them up. However, all outputs will be sent, and an output is only considered an output if it meets  $J_{\text{OUT}}$ . Therefore we can identify adversarial outputs by messages that honest parties would accept as justified.

**Definition 3 (Possible Justified Outputs).** *Let  $\Pi$  be a protocol with output justifier  $J$ . When we say that an  $\ell$ -ary predicate  $P$  holds for all possible justified outputs we mean: Run the protocol  $\Pi$  under attack by the adversary. At some point the adversary may output a sequence  $(P^1, \text{mid}^1, m^1), \dots, (P^\ell, \text{mid}^\ell, m^\ell)$ . We say that the adversary wins if  $\text{mid}^1, \dots, \text{mid}^\ell$  are identified with outputs of  $\Pi$ ,  $P^1, \dots, P^\ell$  are honest, for  $j = 1, \dots, \ell$  it holds that  $J^{\text{mid}^j}(m^j) = \top$  at  $P^j$ , and  $P(m^1, \dots, m^\ell) = \perp$ . Otherwise the adversary loses the game. Any PPT adversary should win with negligible probability.*

Note that since the outputs of a justified protocol are sent to all parties they will become known by the adversary, and we assume the output  $m_i$  of all honest parties  $P_i$  are justified, i.e.,  $J^{\text{mid}^i}(m_i) = \top$  at  $P^{\text{mid}^i}$  at the time of output, we have as a special case that  $P(m^1, \dots, m^\ell) = \top$  for all honest outputs  $m^1, \dots, m^\ell$ , as the adversary can let  $P_j = P^{\text{mid}^j}$  in the game. However, *for all possible justified outputs* implies in addition that the adversary cannot even cook up outputs which look justified to some honest parties and which does not have the property  $P$ .

## 2.2 Justified Reliable Subcast

We use the notion of reliable broadcast (RB) (cf. [Bra87]). For a message identifier  $\text{mid}$  we can force a possible corrupt sender  $G_s$  to send a message  $m$  associated to  $\text{mid}$  such that all servers which receive  $m$  will receive the same  $m$ . Furthermore, if any honest server receives  $m$  then they all receive  $m$ . We look at a version where all parties in  $\mathbb{G}$  can broadcast but only the subset of parties  $\mathbb{S} \subseteq \mathbb{G}$  learn the output, so we call it a reliable subcast. Allowing only  $\mathbb{S}$  to broadcast gives a normal RB among  $\mathbb{S}$ . Setting  $\mathbb{S} = \mathbb{G}$  gives a normal reliable broadcast among  $\mathbb{G}$ .

**Definition 4 (Justified Reliable Subcast).** *A protocol  $\Pi$  for  $n$  servers  $S_1, \dots, S_n$  and  $N$  ground members  $G_1, \dots, G_N$ , where all parties have input  $\text{mid}$ . The message identifier  $\text{mid}$  contains the identity of a sender  $G_s$  along with the description of a justifier  $J_{\text{mid}}$ . The sender additionally has input  $m \in \{0, 1\}^*$  for which  $J_{\text{mid}}(m) = \top$  at  $G_s$  at the time of input.*

**Validity:** *If honest  $G_s$  has input  $(\text{mid}, m)$  and an honest  $S_i$  has output  $(\text{mid}, m')$  then  $m' = m$ .*

**Agreement:** *For all possible justified outputs  $(\text{mid}, m)$  and  $(\text{mid}, m')$  it holds that  $m = m'$ .*

**Eventual Output 1:** *If  $G_s$  is honest and has input  $(\text{mid}, \cdot)$ , and all honest  $G_j$  start running the protocol, then eventually all honest  $S_i$  have output  $(\text{mid}, \cdot)$ .*

**Eventual Output 2:** *If an honest  $S_j$  has output  $(\text{mid}, \cdot)$ , and all honest servers start running the protocol then eventually all honest  $S_i$  have output  $(\text{mid}, \cdot)$ .*

*If only servers  $S_i$  are allowed to give inputs then we call  $\Pi$  a RB (for the servers).*

We will later give many different implementations of reliable subcast. We will also show how to get total-order broadcast among  $\mathbb{S}$  from reliable subcast secure against  $t < n/2$  servers. We will also need that the servers can reliably broadcast a message “out” to the ground population. In this case we assume that all servers agree on the message  $m$  to be sent. We call this primitive outcast.



**Definition 5 (Reliable Outcast).** *A protocol for  $n$  servers  $S_1, \dots, S_n$  and  $N$  ground members  $G_1, \dots, G_N$ , where all parties have input  $\text{mid}$ . The servers additionally have an already agreed upon input  $m \in \{0, 1\}^*$ .*

**Validity:** *If all honest  $S_i$  have input  $(\text{mid}, m)$  and an honest  $G_i$  has output  $(\text{mid}, m')$  then  $m' = m$ .*

**Eventual Output:** *If all honest  $S_i$  have input  $(\text{mid}, m)$  and start running the protocol, then eventually all honest  $G_i$  which start running the protocol have an output  $(\text{mid}, \cdot)$ .*

Outcasting is trivial when at most  $t < n/2$  servers are corrupt. First each server takes  $m$  and uses an erasure code to encode it as  $(m_1, \dots, m_n)$  such that it can be decoded using  $n - t$  of the values  $m_i$ . Using standard techniques and assuming  $|m| \geq n \log n$ , this can be done with  $\sum_{i=1}^n |m_i| = \mathcal{O}(|m|)$ . Then each server computes a Merkle-Damgård hash tree on  $(m_1, \dots, m_n)$  with root  $h$ . Server  $S_i$  sends to each  $G_j$  the root  $h$  and  $m_i$  and a path  $h_i$  showing that  $m_i$  is in leaf  $i$  of the tree with root  $h$ . Each ground member  $G_j$  waits for  $n - t$  identical reports of  $h$  and adopts this value. It was sent by at least one honest server, so it must be correct. Then it waits for  $n - t$  messages  $m_i$  along with proofs that they are in  $h$ . From these  $n - t$  values it computes  $m$ . We call this protocol  $\Pi_{\text{OUTCAST}}$ . This costs communication  $\mathcal{O}(nN \log(n)\kappa + N|m|)$ . For  $|m| \geq n \log(n)\kappa$  and  $n = \Theta(\kappa)$  this is  $\mathcal{O}(N|m|)$ , which is asymptotically optimal as each ground member has to receive  $m$ . More details are given in Section 4.

### 2.3 Justified Leader Election

In the following definitions we will often drop the explicit mentioning of message identifiers. They are tacitly identified by variable names and adding explicit message identifiers gives no more insight. In addition, in each of the following protocols each party has an additional input  $\text{sid}$ , the session identifier, allowing to identify different runs. We often let it be tacit. All security properties are defined relative to the same  $\text{sid}$  and all protocol messages tacitly contains  $\text{sid}$ .

**Definition 6 (Justified Leader Election).** *A protocol for  $n$  servers  $S_1, \dots, S_n$  where the input of the parties is a fixed symbol  $\text{ELECT}$ . The output is  $j \in [n]$ . There is an output justifier  $J_{\text{OUT}}$  specified by the protocol.*

**Liveness:** *If all honest servers start running the protocol with input  $\text{ELECT}$ , then eventually all honest parties  $S_i$  have an output.*

**Validity:** *For all possible justified outputs  $j$  it holds that  $j \in \{1, \dots, n\}$ .*

**Agreement:** *For all possible justified outputs  $j$  and  $j'$  it holds that  $j = j'$ .*

**Unpredictable:** *If no honest party had input  $\text{ELECT}$  yet, then the adversary cannot guess  $j$  negligibly better than at random. In particular, consider the game where the adversary can run the protocol and at any point in time where no honest party has input  $\text{ELECT}$  yet—and at most once—can output  $j'$ . It wins the game if it can continue the execution of the protocol and make an honest party output  $j'$ . No PPT adversary should win this game with probability better than  $1/n + \text{negl}$ .*

We note that the leader election protocol from [CKS00] works for the asynchronous model with  $t < n/2$  corruptions if a threshold signature scheme with unique signatures has been set up among the servers. On  $(\text{ELECT}, \text{sid})$  server  $S_i$  sends its signature share on  $\text{sid}$ . Given  $n - t$  signature shares a server reconstructs  $\sigma_{\text{sid}} = \text{Sig}_{\text{sk}}(\text{sid})$ , computes  $c_{\text{sid}} = \text{Hash}(\sigma_{\text{sid}})$ . Modelling  $\text{Hash}$  as a random

oracle this gives a uniformly random  $c_{\text{sid}}$  unknown until  $\sigma_{\text{sid}}$  is known, in particular until the first honest parties gets input `ELECT`. This costs communication  $\mathcal{O}(n^2\kappa)$ .

In a setting where we do not want to assume special setup among the servers we can share  $\text{sk}$  among the ground population with reconstruction threshold  $T + 1$ . On input  $(\text{ELECT}, \text{sid})$  a server  $S_i$  sends an authenticated  $(\text{ELECT}, \text{sid})$  to the ground population. On seeing  $t + 1$  such messages a ground member sends its signature share to each of the  $n$  servers. Given  $T + 1$  of the signature shares the server can compute the coin as above. This costs communication  $\mathcal{O}(Nn\kappa)$ . Below we let  $\text{ELECT} = Nn\kappa$ .

## 2.4 Causal Cast

We now present a framework for describing protocols for DAG-style protocols (cf. [KKNS21]) in a modular way and in combination with non-DAG style protocols. In a DAG-style protocol all messages  $m$  are reliably broadcast and they point to the other reliably broadcast messages they were computed from. Therefore the receiver can recompute and check the message  $m$ . In fact, the message  $m$  never has to be sent, the receiver can compute it itself. This allows to compress the communication complexity. All that needs to be reliably broadcast are the pointers to the messages used to compute  $m$ . This implements reliable, causal communication against a Byzantine adversary. We will therefore call our system below causal cast (CauCast). Each message  $m$  to be CauCast will have a message identifier  $\text{mid}$ . The set of message identifiers is divided into four disjoint sets of free-choice identifiers, computed-message identifiers, leader-election identifiers, and constant identifiers. We assume that the type of  $\text{mid}$  can be determined efficiently. Each free-choice, computed message, and constant identifier  $\text{mid}$  specifies a sender  $S^{\text{mid}}$ . Each free-choice identifier  $\text{mid}$  specifies an input justifier  $J_{\text{IN}}^{\text{mid}}$ . Each computed message identifier specifies a next message function  $\text{NextMessage}^{\text{mid}}$ . This function is PPT and takes as input a set of pairs  $M = \{(\text{mid}_j, m_j)\}_{j=1}^{\ell}$  and outputs  $m = \text{NextMessage}^{\text{mid}}(M)$ , where  $m = \perp$  indicates that  $M$  is not a valid set of inputs for computing the message for  $\text{mid}$ . For leader-election identifiers  $\text{mid}$  we have that  $\text{mid} = \text{sid}$  for a session identifier  $\text{sid}$  for a justified leader-election protocol  $\Pi_{\text{ELECT}}$  with output justifier  $\Pi_{\text{ELECT}} \cdot J_{\text{OUT}}$ . The constant identifiers are just a convenient tool to define that some values on which the servers already agree have been “delivered”, for instance hardwired values of the protocol.

The system guarantees liveness, agreement on all messages, and that all messages are valid inputs, valid outputs of a leader election, or computed correctly from other valid messages.

**Definition 7 (Causal Cast).** *A protocol for  $n$  servers  $S_1, \dots, S_n$  is called a causal cast (CauCast) if it has the following properties.*

**Free-Choice Send:** *A server  $S_i$  can have input  $(\text{CAUCAST-SEND}, \text{mid}, m)$  where  $\text{mid}$  is a free choice identifier  $S_i = S^{\text{mid}}$  and  $J_{\text{IN}}^{\text{mid}}(m) = \top$  at  $S^{\text{mid}}$  at the time of input.*

**Computed-Message Send:** *A server  $S_i$  can have input  $(\text{CAUCAST-SEND}, \text{mid}, m, \text{mid}_1, \dots, \text{mid}_\ell)$ , where  $\text{mid}$  is a computed-message identifier,  $S_i = S^{\text{mid}}$ ,  $S_i$  earlier gave outputs  $(\text{CAUCAST-DEL}, \text{mid}_j, m_j)$  for  $j = 1, \dots, \ell$ , and  $\perp \neq m = \text{NextMessage}^{\text{mid}}((\text{mid}_1, m_1), \dots, (\text{mid}_\ell, m_\ell))$ .*

**Constant Send:** *A server  $S_i$  can have input  $(\text{CAUCAST-SEND}, \text{mid}, m)$  where  $\text{mid}$  is a constant identifier. In that case it is guaranteed that all servers eventually have the same input  $(\text{CAUCAST-SEND}, \text{mid}, m)$ .*

**Free-Choice Validity:** A server  $S_i$  can have output  $(\text{CAUCAST-DEL}, \text{mid}, m)$ , where  $\text{mid}$  is a free-choice identifier. It then holds that  $J_{\text{IN}}^{\text{mid}}(m) = \top$  at  $S_i$  at the time of output. Furthermore, if  $S_j = S^{\text{mid}}$  is honest, then  $S_j$  had input  $(\text{CAUCAST-SEND}, \text{mid}, m)$ .

**Leader Election Validity:** A server  $S_i$  may output  $(\text{CAUCAST-DEL}, \text{mid}, m)$  where  $\text{mid}$  is a leader-election identifier. In that case  $\text{mid} = \text{sid}$  is a session identifier for a justified leader election and  $\Pi_{\text{ELECT}}.J_{\text{OUT}}^{\text{sid}}(m)$  at  $S_i$  at the time of output.

**Computed-Message Validity:** A server  $S_i$  can have output  $(\text{CAUCAST-DEL}, \text{mid}, m, \text{mid}_1, \dots, \text{mid}_\ell)$ , where  $\text{mid}$  is a computed-message identifier. In that case  $S_i$  earlier gave outputs  $(\text{CAUCAST-DEL}, \text{mid}_j, m_j, \dots)$  for  $j = 1, \dots, \ell$ , and  $\perp \neq m = \text{NextMessage}^{\text{mid}}((\text{mid}_1, m_1), \dots, (\text{mid}_\ell, m_\ell))$ .

**Constant Validity:** A server  $S_i$  can have output  $(\text{CAUCAST-DEL}, \text{mid}, m)$ . In that case it immediately before had input  $(\text{CAUCAST-SEND}, \text{mid}, m)$ .

**Liveness:** If an honest server  $S_i$  had input  $(\text{CAUCAST-SEND}, \text{mid}, \dots)$  or some honest server had output  $(\text{CAUCAST-DEL}, \text{mid}, \dots)$  and all honest servers are running the system, then eventually all honest servers have output  $(\text{CAUCAST-DEL}, \text{mid}, \dots)$ .

**Agreement:** For all possible justified outputs  $(\text{CAUCAST-DEL}, \text{mid}, m, \dots)$  and  $(\text{CAUCAST-DEL}, \text{mid}, m', \dots)$  it holds that  $m' = m$ .

Below is a protocol implementing CauCast. It uses a sub-protocol for reliable broadcast (RB) among the servers and a sub-protocol  $\Pi_{\text{ELECT}}$  for justified leader election. The protocol  $\Pi_{\text{ELECT}}$  may be initialised by other protocols. The CauCast merely reports its outputs—this allows the elections to be inputs for computed messages in the causal cast system.

**Free-Choice Send:** On input  $(\text{CAUCAST-SEND}, \text{mid}, m)$  at  $S_i$  where  $\text{mid}$  is a free choice  $S_i$  will RB  $m$  with message identifier  $\text{mid}$ .

**Computed-Message Send:** On input  $(\text{CAUCAST-SEND}, \text{mid}, m, \text{mid}_1, \dots, \text{mid}_\ell)$  at  $S_i$ , where  $\text{mid}$  is a computed-message identifier,  $S_i$  will RB  $(\text{mid}_1, \dots, \text{mid}_\ell)$  with message identifier  $\text{mid}$ .

**Free-Choice Deliver:** On RB  $m$  from with message identifier  $\text{mid}$  from  $S^{\text{mid}}$  where  $\text{mid}$  is a free-choice identifier, wait until  $J_{\text{IN}}^{\text{mid}}(m) = \top$  (possibly waiting forever if this never happens) and then output  $(\text{CAUCAST-DEL}, \text{mid}, m)$ .

**Computed-Message Deliver:** On RB  $(\text{mid}_1, \dots, \text{mid}_\ell)$  with message identifier  $\text{mid}$  from  $S^{\text{mid}}$  wait until outputs  $(\text{CAUCAST-DEL}, \text{mid}_j, m_j)$  were given for  $j = 1, \dots, \ell$  (possibly waiting forever if this never happens), compute  $m = \text{NextMessage}^{\text{mid}}((\text{mid}_1, m_1), \dots, (\text{mid}_\ell, m_\ell))$ , and output  $(\text{CAUCAST-DEL}, \text{mid}, m, \text{mid}_1, \dots, \text{mid}_\ell)$  if  $m \neq \perp$ .

**Leader-Election Deliver:** On output  $(\text{ELECT}, \text{sid}, K)$  from  $\Pi_{\text{ELECT}}$  output  $(\text{CAUCAST-DEL}, \text{sid}, K)$ .

**Constant Deliver:** On input  $(\text{CAUCAST-SEND}, \text{mid}, m)$  where  $\text{mid}$  is a constant identifier, output  $(\text{CAUCAST-DEL}, \text{mid}, m)$ .

**Definition 8 (CauCast Conventions).** When using CauCast we let the justifier  $J^{\text{mid}}$  of a message  $m$  at  $S_i$  be that  $(\text{CAUCAST-DEL}, \text{mid}, m, \dots)$  was output by  $S_i$ . When describing protocols using CauCast we assume that all outputs were CauCast and this is how they are justified. Specifically, the session identifier  $\text{sid}$  of the protocol specifies a message identifier  $\text{mid}$  and  $J_{\text{OUT}}^{\text{sid}}(m) = J^{\text{mid}}(m)$ . When describing protocols using CauCast we will always as the first instruction CauCast the inputs of the protocol as a free-choice message. In the typical case where the input  $x$  is the output  $y$  of some previous protocol we will by convention have that  $y = m$  for some message with message identifier  $\text{mid}$ , where  $m$  was by convention already CauCast. In that case we do not reCauCast the

input  $x$ . All servers will by convention know  $\text{mid}$  and that  $x = m$ , so getting input for the protocol just means waiting for  $(\text{CAUCAST-DEL}, \text{mid}, m)$  and then letting  $x = m$ .

We address communication complexity. We can represent a session identifier  $\text{sid}$  with  $\kappa$  bits as we can always hash the session identifiers. We can let all servers  $G_i$  number the messages they send using a counter  $c_i = 1, 2, \dots$ . Then message identifiers can be of the form  $(\text{sid}, i, c_i)$ . We do not need to send  $\text{sid}$  along with all  $\text{mid}$  as it is the same for all  $\text{mid}$  in the protocol. Using that  $n, N \in \text{poly}(\kappa)$  and that  $c_i$  can become at most polynomially large in a poly-time run of the system each  $\text{mid}$  can therefore be represented in  $\log(\kappa)$  extra bits. We can therefore represent  $(\text{mid}, \text{mid}_1, \dots, \text{mid}_\ell)$  as  $\kappa + \ell \log(\kappa)$  bits. Therefore the communication complexity is that of reliably broadcasting all free-choice messages  $m$ , the leader elections, plus the complexity of reliably broadcasting  $\kappa + \ell \log(\kappa)$  per computed message. In many cases a computed message is computed from  $n - t$  outputs of  $n$  possible messages with session identifiers  $\text{sid}_1, \dots, \text{sid}_n$  known by all servers. In these cases we can send an  $n$ -bit vector indicating the  $n - t$  session identifiers to use. Then the communication is only  $\mathcal{O}(\kappa + n) = \mathcal{O}(\kappa)$  bits per computed message. We return to this when analysing the complexity of concrete protocols below.

**Definition 9 (Complexity).** *We say that a protocol  $\Pi$  using  $\text{CauCast}$  has (expected communication) complexity*

$$\mathcal{O}(c_1 \text{IN} + c_2 \text{RS} + c_3 \text{RS}_\# + c_4 \text{ELECT} + c_5)$$

*if the following holds in expectation, using the above methods for compression, and under  $\mathcal{O}$ :  $c_1$  is the total number of bits that the protocol needs to reliably subcast<sup>1</sup> its inputs,  $c_2$  is the total number of bits that the protocol needs to reliably subcast intermediary values and outputs,  $c_3$  is the number of reliable subcast instances run,  $c_4$  is the number of calls to the reliable leader election primitive, and  $c_5$  is the total number of bits sent otherwise. Notice that we count all messages sent by all servers. By time complexity of a protocol we mean how long wall clock time it takes to run, measured in units of the longest it takes to deliver a message during the execution, called  $\Delta_{\text{NET}}$ .*

The reason why we single out the complexity of inputs is that when we compose protocols,  $c_1 \text{IN}$  falls away if the inputs are outputs of a previous protocol, as their cost is included in  $\text{RS}$  of that protocol.

### 3 Total Order Broadcast from RB and $t < n/2$

We now present a protocol for total order broadcast. We will here make a distinction between ground population which can broadcast and the servers holding the ledger.

**Definition 10 (Total Order Subcast).** *A protocol for  $n$  servers  $S_1, \dots, S_n$  and ground population  $G_1, \dots, G_N$ . There is an input justifier  $J_{\text{IN}}$ . Each  $S_i$  holds a list  $\text{Ledger}_i$ .*

**Input:**  $G_i$  get inputs  $m$  where  $J_{\text{IN}}(m) = \top$ . We use  $\text{Scheduled}$  to denote the set of  $(i, m)$  such that  $m$  was input at an honest  $G_i$ . We use  $\text{Scheduled}^t$  to denote the value of this set at time  $t$ .

**Liveness:** For all honest  $S_i$  and all times  $t$  it holds eventually that  $\text{Scheduled}^t \subseteq \text{Ledger}_i$ .

**Monotone:** For all  $S_i$  and  $t' \geq t$  it holds that  $\text{Ledger}_i^t \sqsubseteq \text{Ledger}_i^{t'}$ .

**Agreement:** For all  $S_i$  and  $S_j$  it holds that  $\text{Ledger}_i^t \sqsubseteq \text{Ledger}_j^t$  or  $\text{Ledger}_j^t \sqsubseteq \text{Ledger}_i^t$ .

<sup>1</sup> I.e., the sum of the length of messages input to a reliable subcast.

If the above holds but with each  $G_i$  holding  $\text{Ledger}_i$  and Liveness, Monotone and Agreement holding for all honest  $G_i$  and  $G_j$  instead of  $S_i$  and  $S_j$ , then we call the protocol an *Total Order Broadcast* (for the ground population).

In this section we build a Total Order Subcast which has communication complexity  $\mathcal{O}(\beta \cdot \text{RS})$  when there is enough traffic, as discussed in detail later. By letting the servers outcast the ledger to the ground population the complexity becomes  $\mathcal{O}(\beta \cdot \text{RS} + \beta N)$  as we can outcast  $\beta$  bits with complexity  $\beta N$  when outcasting in large enough blocks. The term  $\beta N$  is clearly optimal as all  $N$  ground members need to receive all  $\beta$  bits in Total Order Broadcast. In the following we therefore focus on presenting and proving the Total Order Subcast. It will be build up via a sequence of increasingly powerful primitives.

### 3.1 Justified Gather

We describe and analyse our Justified Gather protocol. We consider protocols where each server has an input  $B_i \in \{0, 1\}^*$ , which we will call a block below.

**Definition 11 (block set).** A block set is a set of pairs  $U = \{(S_j, B_j)\}_{S_j \in P}$ , where  $P \subset \mathbb{S}$  and  $|P| \geq n - t$ .

We think of  $(S_j, B_j) \in U$  as  $S_j$  having input  $B_j$ . The Gather primitive says that each server  $S_i$  has a block set  $U_i$  as output and there is a large common core, i.e., a set  $U$  of size  $n - t$  which is a subset of each  $U_i$ . So all servers to some extent agree on a large set  $U$ , but some  $S_i$  might have extra elements in  $U_i$  and they do not know which are the extra ones.

**Definition 12 (Justified Gather).** A protocol for  $n$  servers  $S_1, \dots, S_n$ . There is an input justifier  $J_{\text{IN}}$  and an output justifier  $J_{\text{OUT}}$  specified by the protocol. All honest  $S_i$  have an input  $B_i$  for which  $J_{\text{IN}}(B_i) = \top$  at  $S_i$  at the time the input is given.

**Liveness:** If all honest parties start running the protocol with a  $J_{\text{IN}}$ -justified input then eventually all honest servers have a  $J_{\text{OUT}}$ -justified output.

**Justified Blocks:** For all possible justified outputs  $U$  and all (potentially corrupt)  $S_i$  and all  $(S_i, B_i) \in U$  it holds that  $J_{\text{IN}}(B_i) = \top$ .

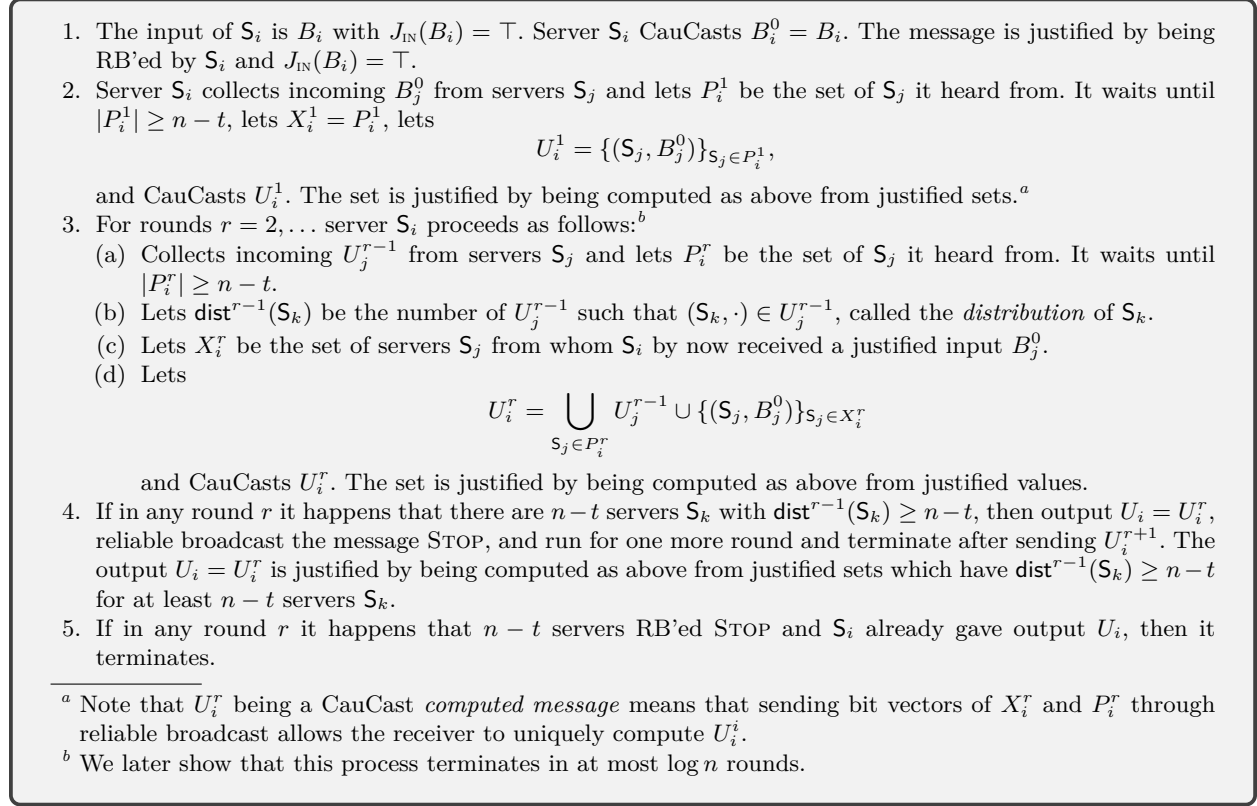
**Validity:** For all possible justified outputs  $U$  and all honest  $S_i$  and all  $(S_i, B_i) \in U$  it holds that  $S_i$  had input  $B_i$ .

**Agreement:** For all possible justified outputs  $U$  and  $U'$  and all  $(S_i, B_i) \in U$  and  $(S_i, B'_i) \in U'$  it holds that  $B_i = B'_i$ .

**Large Core:** For all possible justified outputs  $(U^1, \dots, U^m)$  it holds that  $|\bigcap_{k=1}^m U^k| \geq n - t$ .

The protocol works for  $t < n/2$  corruptions and is inspired by the ACS protocol in [Can95] which worked for  $t < n/3$  corruptions. Making a variant of this protocol work for  $t < n/2$  is central to the results in this paper. The main observation is that if the sending of sets in [Can95] is replaced by CauCast then it works for  $t < n/2$ . The protocol from [Can95] proceeds via a  $\log(n)$ -round gather structure. We add to this gather structure that late original blocks are added when they arrive, this is the  $X_i^r$  sets below. This means that in constant time all honest parties have their blocks enter the gather structure, which will build a large core in constant time. We then use a termination condition from the Gather protocol CSS in [MMNT19, DMM<sup>+</sup>20]. This termination condition allows detecting that a large core will be build in the ensuing round. This allows terminating no later than

when all honest inputs entered the gather structure. This makes the overall protocol constant time. Note that the protocol might still run in  $\log(n)$  rounds of communication even though it runs in constant time. But this requires that the messages of the fastest honest parties are delivered with delay  $\Delta_{\text{NET}}/\log(n)$  and the slowest honest parties have their messages delivered with delay  $\Delta_{\text{NET}}$ , so there must be an asymptotic gap in the delays. The protocol is presented in Fig. 1.



**Fig. 1.** Protocol  $\Pi_{\text{GATHER}}$ .

**Theorem 1.** *If  $t < n/2$  then  $\Pi_{\text{GATHER}}$  is a Justified Gather. If  $\beta = \sum_{i=1}^n |B_i|$  then it has complexity*

$$\mathcal{O}(\beta \text{IN} + n \log(n) \text{RS}_{\#} + n^2 \log(n) \text{RS}) .$$

*Let  $\Delta_{\text{RB}}$  be the actual delivery time of an RB while the protocol is run. Then all honest servers give output within time  $4\Delta_{\text{RB}}$ . In particular, if the RB has constant time complexity then  $\Pi_{\text{GATHER}}$  has constant time complexity.*

*Proof.* We argue Justified Blocks. The message  $B_i$  is justified by being RB'ed by  $S_i$  and  $J_{\text{IN}}(B_i) = \top$ . All later messages are justified by justified messages from the previous rounds, and therefore ultimately each  $B_i$  is justified by being RB'ed by  $S_i$  and  $J_{\text{IN}}(B_i) = \top$ . This also shows Validity. Agreement follows similarly. For  $B_i$  and  $B'_i$  to appear as output blocks, they must be justified, which implies that they were both reliably broadcast by  $S_i$ . And then they are identical. To argue liveness and large core we need a helper lemma.

We need to analyse the development of the support of the core. A core  $C^r$  for round  $r$  would be a set of size at least  $n - t$  which is a subset of all justified  $U_i^r$ . This in particular means it is a subset of  $\bigcap_{S_j \in P_i^r} U_j^r$  for all justified  $P_i^r$ , where we call  $P_i^r$  justified if all  $S_j \in P_i^r$  sent a justified  $U_j^r$ . We say that the core has support  $\sigma$  in round  $r$  if it holds for all  $P_i^r$  of *size at most*  $\sigma$  that  $\bigcap_{S_j \in P_i^r} U_j^r$  has size at least  $n - t$ . We let the support of a round  $r$  be the largest such  $\sigma$ , i.e.,

$$\Sigma^r = \max \left\{ \sigma \geq 1 \left| \left( \min_{\text{justified } P_i^r : |P_i^r| \leq \sigma} \left| \bigcap_{S_j \in P_i^r} U_j^r \right| \right) \geq n - t \right. \right\},$$

where we let  $\Sigma^r = 0$  if the set is empty.

**Lemma 1 (growing core support).**  $\Sigma^r \geq 2^{r-1}$ .

*Proof.* We prove the lemma by induction. The basis is  $r = 1$  where  $U_i^1 = \{(S_j, B_j^0)\}_{S_j \in P_i^1}$ . Note that it has size  $n - t$  by construction and we only consider intersections between sets of size  $\sigma = 2^{r-1} = 1$ , so the basis case is trivial. We then assume the induction hypothesis for  $r$  and prove it for  $r + 1$ . We have to prove support  $\sigma = 2^r$ . Consider a justified set  $P_i^{r+1}$  with  $|P_i^{r+1}| \leq \sigma$ . It is enough to show that

$$\left| \bigcap_{S_j \in P_i^{r+1}} U_j^{r+1} \right| \geq n - t.$$

We can pair the  $2^r$  servers<sup>2</sup> in  $P_i^{r+1}$  into  $(S_{j_{k,1}}, S_{j_{k,2}})$  for  $k = 1, \dots, 2^{r-1}$ . For each of these pairs there exists a server  $S_{j_k}$  such that  $U_{j_k}^r \subseteq U_{j_{k,1}}^{r+1}$  and  $U_{j_k}^r \subseteq U_{j_{k,2}}^{r+1}$ . This is because both  $U_{j_{k,1}}^{r+1}$  and  $U_{j_{k,2}}^{r+1}$  to be justified contain the union of  $n - t \geq t + 1$  justified sets  $U_j^r$ , and these sets being justified means that they are reliably broadcast, so since  $(n - t) + (t + 1) > n$  one identical  $U_j^r$  will be a subset of both  $U_{j_{k,1}}^{r+1}$  and  $U_{j_{k,2}}^{r+1}$ . This implies that  $U_{j_k}^r \subseteq U_{j_{k,1}}^{r+1} \cap U_{j_{k,2}}^{r+1}$ , which implies that

$$\bigcap_{k=1}^{2^{r-1}} U_{j_k}^r \subseteq \bigcap_{k=1}^{2^{r-1}} (U_{j_{k,1}}^{r+1} \cap U_{j_{k,2}}^{r+1}).$$

By the induction hypothesis that round  $r$  had support  $2^{r-1}$  we have that

$$\left| \bigcap_{k=1}^{2^{r-1}} U_{j_k}^r \right| \geq n - t.$$

By construction

$$\bigcap_{k=1}^{2^{r-1}} (U_{j_{k,1}}^{r+1} \cap U_{j_{k,2}}^{r+1}) = \bigcap_{S_j \in P_i^{r+1}} U_j^{r+1}.$$

This combines to give the desired conclusion. □

Now let  $\ell = \lceil \log_2(n) \rceil + 1$ . Then  $\Sigma^\ell \geq 2^{\ell-1} \geq n$ , which implies that

$$\left| \bigcap_{j=1}^n U_j^\ell \right| \geq n - t.$$

<sup>2</sup> In case  $|P_i^{r+1}| < \sigma$  one can consider inserting copies of one server to pad up to  $\sigma$  for the argument to go through.

We first argue Liveness. We do this in two steps. We argue that if no honest server terminates, then at least one honest server terminates. We then argue that if one honest server terminates, then all honest servers terminate. Assume for the sake of contradiction that no honest server terminates. So they all reach round  $\ell + 1$ . Now, clearly for all  $S_k \in \bigcap_{j=1}^n U_j^\ell$  we will have that  $\text{dist}^\ell(S_k) \geq n - t$  as it will be in all received  $U_j^\ell$  in round  $\ell + 1$ . Therefore, in round  $\ell + 1$  all honest servers will have  $n - t$  servers  $S_k$  with  $\text{dist}(S_k) \geq n - t$  and will output  $U_i^{\ell+1}$  if they did not already give output. And then they all run the next round, and then terminate, a contradiction. Assume then that some honest  $S_i$  terminates. Let  $S_i$  be the first honest server to terminate and let  $r_i$  be the round in which it happened. Then in round  $r_i - 1$  server  $S_i$  had  $n - t$  servers  $S_k$  with  $\text{dist}^{r_i-2}(S_k) \geq n - t$ . Each such  $S_k$  was in  $n - t \geq t + 1$  of the sets  $U_j^{r_i-2}$ . Since all servers in the next round collect  $n - t$  sets  $U_j^{r_i-2}$  all servers (corrupted and honest) must collect at least one  $U_j^{r_i-2}$  which contains  $S_k$ . Therefore  $S_k$  will be in *all* justified  $U_j^{r_i-1}$ . Note that all honest servers run at least until round  $r_i$ . In that round they will then all have  $\text{dist}^{r_i-1}(S_k) \geq n - t$  for all  $n - t$  servers  $S_k$  for which  $S_i$  had  $\text{dist}^{r_i-2}(S_k) \geq n - t$  in round  $r_i - 1$ . Therefore they give output and send STOP. Therefore they will all eventually see  $n - t$  STOP messages and then they terminate.

We then argue Large Core. We argued above that all honest servers terminate. It is clear by inspection that if an honest server terminates it also gave an output  $U_i$ . We have that  $U_i = U_i^{r_i}$  for some  $r_i$ . We argued as part of termination that when the first  $S_i$  terminated with output  $U_i = U_i^{r_i}$  then in round  $r_i - 1$  the server  $S_i$  had  $n - t$  servers  $S_k$  with  $\text{dist}^{r_i-2}(S_k) \geq n - t$  and all these servers ended up in all  $U_j^{r_i}$ . Therefore these  $S_k$  servers will be part of all justified outputs  $U_j$ .

We then address the communication complexity. In the first round all the blocks are CauCast, for a total of  $\beta$  bits. In the next  $\mathcal{O}(\log(n))$  rounds there are no further free choice inputs. The servers only need to CauCast  $P_i^r$  and  $X_i^r$ , which can be represented as two  $n$ -bit vectors. Sending the stop signal is asymptotically negligible. We then consider the time complexity. Let  $\Delta_{\text{RB}}$  be the actual maximum delivery time of an RB while the protocol is run. Within time  $2\Delta_{\text{RB}}$  all honest  $B_i$  will via some  $X_j^r$  have entered  $U_j^r$  at all honest  $S_j$ . Let  $r$  be the round in which this happens. In the next round ( $r + 1$ ),  $B_i$  will then be in all justified  $U_k^{r+1}$  as all  $S_k$  must add in their union one of the  $n - t \geq t + 1$  sets  $U_j^r$  containing  $B_i$ . Therefore, in the next round ( $r + 2$ ) we will have  $\text{dist}^{r+1}(S_i) \geq n - t$  for all honest  $S_i$ . And then all honest servers give output. Each round consists of one round of reliable broadcast, so it takes time at most  $\Delta_{\text{RB}}$ . Therefore, all in all, within  $4\Delta_{\text{RB}}$  time all honest servers give output.  $\square$

### 3.2 Justified Graded Gather

We describe and analyse our Justified Graded Gather protocol. This is just a Justified Gather, where each server also has knowledge about the common core. Each  $S_i$  will output a set  $T_i$  of size at least  $n - t$  which is guaranteed to be a subset of the common core  $U$ .

**Definition 13 (Justified Graded Gather).** *A protocol for  $n$  servers  $S_1, \dots, S_n$ . There is an input justifier  $J_{\text{IN}}$  and an output justifier  $J_{\text{OUT}}$  specified by the protocol. All honest  $S_i$  have an input  $B_i$  for which  $J_{\text{IN}}(B_i) = \top$  at  $S_i$  at the time the input is given.*

**Liveness:** *If all honest servers start running the protocol with a  $J_{\text{IN}}$ -justified input then eventually all honest servers have a  $J_{\text{OUT}}$ -justified output.*

**Justified Blocks:** *For all possible justified outputs  $(U, T)$  and all (potentially corrupt)  $S_i$  and all  $(S_i, B_i) \in U$  it holds that  $J_{\text{IN}}(B_i) = \top$ .*



**Sub Core:** For all possible justified outputs  $((U^1, T^1), \dots, (U^m, T^m))$  it holds that  $T^1 \subseteq \bigcap_{k=1}^m U^k$ .

**Validity:** For all possible justified outputs  $(U, T)$  and all honest  $S_i$  and all  $(S_i, B_i) \in U$  it holds that  $S_i$  had input  $B_i$ .

**Agreement:** For all possible justified outputs  $(U, T)$  and  $(U', T')$  and all  $(S_i, B_i) \in S$  and  $(S_i, B'_i) \in U'$  it holds that  $B_i = B'_i$ .

**Large Sub Core:** For all possible justified outputs  $((U^1, T^1), \dots, (U^m, T^m))$  it holds that  $|\bigcap_{k=1}^m T^k| \geq n - t$ .

1. The input of  $S_i$  is  $B_i$  with  $J_{\text{IN}}(B_i) = \top$ . All servers run  $\Pi_{\text{GATHER}}$  with  $S_i$  inputting  $B_i$  justified by  $J_{\text{IN}}$ .
2. Server  $S_i$  waits for output  $U'_i$  and CauCasts  $U'_i$ .
3. Server  $S_i$  collects  $U'_j$  from servers  $S_j$ , lets  $P_i$  be the set of  $S_j$  it heard from and waits until  $|P_i| \geq n - t$ .
4. Server  $S_i$  lets

$$U_i = \bigcup_{S_j \in P_i} U'_j, \quad T_i = \bigcap_{S_j \in P_i} U'_j$$

CauCasts  $(U_i, T_i)$  and outputs  $(U_i, T_i)$ . The outputs are justified by being computed as above from justified sets.

**Fig. 2.**  $\Pi_{\text{GRADEDGATHER}}$

**Theorem 2.** If  $t < n/2$  then  $\Pi_{\text{GRADEDGATHER}}$  is a Justified Graded Gather. If  $\beta = \sum_{i=1}^n |B_i|$  and  $\Pi_{\text{GRADEDGATHER}}$  uses  $\Pi_{\text{GATHER}}$  from Fig. 1 as sub-protocol, then it has complexity

$$\mathcal{O}(\beta \text{IN} + n^2 \log(n) \text{RS} + n \log(n) \text{RS}_{\#}) .$$

The time complexity is  $\mathcal{O}(1)$  times that of the RB used.

*Proof.* Complexities, Liveness, Justified Blocks, Validity and Agreement follow from the same properties of Justified Gather. We have that  $\bigcap_{k=1}^n T_k = \bigcap_{k=1}^n \bigcap_{S_j \in P_i} U'_j$ , so Large Sub Core follows from Large Core of Justified Gather. Sub Core follows from the below lemma.

**Lemma 2 (Sub Core).** Let  $(\cdot, T_i)$  and  $(U_\iota, \cdot)$  be any possible justified outputs. Then  $T_i \subseteq U_\iota$ .

*Proof.* It is enough to argue that if  $(S_k, B_k) \in T_i$  then  $(S_k, B_k) \in U_\iota$ . If  $(S_k, B_k) \in T_i$  then  $S_k \in U'_j$  for all  $S_j \in P_i$ . Since  $|P_i| \geq n - t \geq t + 1$  it follows that any server receiving  $n - t$  justified sets  $U'_j$  will also receive a set  $U'_j$  with  $S_k \in U'_j$ . Namely, the sets are reliably broadcast so if two servers receive justified  $U'_j$  and  $\hat{U}'_j$  then  $U'_j = \hat{U}'_j$ . Since all servers collect  $n - t$  justified sets  $U^\ell$  to justify  $U_\iota = \bigcup_{S_j \in P_i^\ell} U_j^\ell$  it follows that  $(S_j, B_k) \in U_\iota$ .  $\square$

### 3.3 Justified Graded Block Selection

We now present our graded block selection protocol. Here each server has as input a block  $B_i$  and as output a block  $C_i$ . The goal is to let  $C_i$  be one of the inputs and to agree on  $C_i$ . Since corrupted servers can pick their own input and we allow that  $C_i = B_i$  for a corrupt  $S_i$  we simply define validity by saying that the output should be some justified input. Note that this implies that if there is only

1. The input of  $S_i$  is  $B_i$  with  $J_{IN}(B_i) = \top$ .
2. The servers run  $\Pi_{\text{GRADEDGATHER}}$  with input  $B_i$  and input justifier  $J_{IN}$ . Let the output of  $S_i$  be  $(U_i, T_i)$ .
3. The servers run a justified leader election to elect a justified king  $S_k$ . Each server starts the leader election only after getting output from the Graded Gather.
4. Server  $S_i$  outputs

$$(C_i, g_i) = \begin{cases} (B_k, 2) & \text{if } \exists(S_k, B_k) \in T_i \\ (B_k, 1) & \text{if } \exists(S_k, B_k) \in U_i \setminus T_i \\ (B_i, 0) & \text{if } \nexists(S_k, \cdot) \in U_i . \end{cases}$$

and CauCasts  $(C_i, g_i)$  and outputs  $(C_i, g_i)$ . The output is justified by being computed as above from justified values.

**Fig. 3.**  $\Pi_{\text{GRADEDSELECTBLOCK}}$

one possible justified input, then that will become the only justifiable output. We will not always be able to perfectly agree on the output, instead the output will have a grade  $g \in \{0, 1, 2\}$ . The grades are never more than 1 apart and if the grade is 2 then there was agreement on  $C_i$ . Finally, we want that with some non-zero probability the grade will be 2.

**Definition 14 (Justified Graded Block Selection).** *A protocol for  $n$  servers  $S_1, \dots, S_n$ . There is an input justifier  $J_{IN}$  and an output justifier  $J_{OUT}$  specified by the protocol. All honest  $S_i$  have an input  $B_i$  for which  $J_{IN}(B_i) = \top$  at the time the input  $B_i$  is given. The output of the protocol is a block  $C_i$  justified by  $J_{OUT}$ .*

**Liveness:** *If all honest servers start running the protocol with a  $J_{IN}$ -justified input then eventually all honest servers have a  $J_{OUT}$ -justified output.*

**Justified Output:**  *$J_{IN}(C_i) = \top$  holds for all possible  $J_{OUT}$ -justified outputs  $C_i$ .*

**Graded Agreement:** *For all possible justified outputs  $(C_i, g_i)$  and  $(C_j, g_j)$  it holds that  $|g_i - g_j| \leq 1$ . Furthermore, if both  $g_i, g_j > 0$  then  $C_i = C_j$ .*

**Positive Agreement:** *There exists  $\alpha > 0$  such that with probability at least  $\alpha - \text{negl}$  some honest  $S_i$  will have output  $(C_i, g_i)$  with  $g_i = 2$ .*

**Stability:** *If there are possible justified outputs  $C_i$  and  $C_j$  with  $C_i \neq C_j$  then there exist two justified inputs  $B_i$  and  $B_j$  with  $B_i \neq B_j$ .*

**Theorem 3.** *If  $t < n/2$  then  $\Pi_{\text{SELECTBLOCK}}$  is a Justified Graded Block Selection. When  $\beta = \sum_{i=1}^n |B_i|$  and when using  $\Pi_{\text{GRADEDGATHER}}$  from Fig. 2 as sub-protocol the complexity is*

$$\mathcal{O}(\beta \text{ IN} + n^2 \log(n) \text{ RS} + n \log(n) \text{ RS}_{\#} + \text{ELECT}) .$$

*Proof.* We start with the complexity. The protocol in Fig. 2 has complexity

$$\mathcal{O}(\beta \text{ IN} + n^2 \log(n) \text{ RS} + n \log(n) \text{ RS}_{\#}) .$$

In addition to this  $\Pi_{\text{SELECTBLOCK}}$  only does one leader election. It has to send no more PUJM information as the justifier for  $(C_i, g_i)$  is the justified  $B_i$ , the justified  $(U_i, T_i)$ , and the justified  $S_k$ , which have all been CauCast already. Liveness is straight forward. We argue Justified Output. Let  $(C_i, g_i)$  be any justified output. If  $g_i = 0$  then by definition  $C_i = B_i$  is a justified input. If  $g_i > 0$  then  $B_i = B_k$  for  $(S_k, C_k) \in U_i$  and therefore  $C_k$  is a justified input to the Graded Gather which

also used  $J_{\text{IN}}$  as input justifier. Then use the Justified Blocks property. To argue Graded Agreement let  $(C_i, g_i)$  and  $(C_j, g_j)$  be any justified outputs. To argue that  $|g_i - g_j| \leq 1$  it is sufficient to prove that if  $g_i = 2$  then  $g_j \neq 0$ . So assume that  $g_i = 2$ . Then  $(S_k, C_i) \in T_k$  for some justified  $T_k$ . Therefore, by Sub Core,  $(S_k, C_i) \in U_j$ , and therefore  $g_j \geq 1$ . Assume then that  $g_i, g_j > 0$ . In that case  $(S_k, C_i) \in U_i$  and  $(S_k, C_j) \in U_j$ , so by Agreement of the Graded Gather it follows that  $C_i = C_j$ . We then argue Positive Agreement for  $\alpha = 1/2$ . It is sufficient to argue that with probability  $1/2$  it holds for some honest  $S_i$  that  $S_k \in T_i$ . Consider the first honest  $S_i$  to start running the leader election. When this happens  $T_i$  is already defined, and  $S_k$  is unpredictable. Since  $|T_i| = n - t > n/2$  it follows that  $S_k \in T_i$  with probability at least  $1/2 - \text{negl}$ . To argue Stability just note that it holds for both  $C_i$  and  $C_j$  that they are justified inputs of  $\Pi_{\text{SELECTBLOCK}}$ .  $\square$

### 3.4 Justified Block Selection

We now present our (ungraded) block selection protocol. The difference from graded block selection is that all possible justified outputs  $C_i$  should be identical.

**Definition 15 (Justified Block Selection).** *A protocol for  $n$  servers  $S_1, \dots, S_n$ . There is an input justifier  $J_{\text{IN}}$  and an output justifier  $J_{\text{OUT}}$ . All honest  $S_i$  have an input  $B_i$  for which  $J_{\text{IN}}(B_i) = \top$  at the time the input was given. The output of the protocol is a block  $C_i$  justified by  $J_{\text{OUT}}$ .*

**Liveness:** *If all honest servers start running the protocol with a  $J_{\text{IN}}$ -justified input then eventually all honest servers have a  $J_{\text{OUT}}$ -justified output.*

**Justified Output:**  *$J_{\text{IN}}(C_i) = \top$  holds for all possible  $J_{\text{OUT}}$ -justified outputs  $C_i$ .*

**Agreement:** *For all possible justified outputs  $C_i$  and  $C_j$  it holds that  $C_i = C_j$ .*

1. The input of  $S_i$  is  $B_i$  with  $J_{\text{IN}}(B_i) = \top$ . It initialises  $\text{GaveOutput}_i = \perp$ .
2. Let  $B_i^0 = B_i$  and  $g_i^0 = 0$  and  $\text{CauCast}(B_i^0, g_i^0)$ , which is justified if  $J_{\text{IN}}(B_i^0) = \top$  and  $g_i^0 = 0$ .
3. For rounds  $r = 1, \dots$  the servers run  $\Pi_{\text{GRADEDSELECTBLOCK}}$  where:
  - (a)  $S_i$  has input  $B_i^{r-1}$ .
  - (b) The input of  $S_i$  is justified by a justified  $(B_i^{r-1}, g_i^{r-1})$ .
  - (c)  $S_i$  eventually gets justified output  $(B_i^r, g_i^r)$ .
4. In addition to the above loop each  $S_i$  runs the following *echo rules*:
  - In the first round  $r$  where  $\text{GaveOutput}_i = \perp$  and  $g_i^r = 2$ , set  $\text{GaveOutput}_i = \top$  and output  $C_i = B_i^r$ . The output justifier is the justifier for  $(B_i^r, g_i^r)$ .
  - In the first round  $r$  where  $\text{GaveOutput}_i = \perp$  and where some justified  $(B_j^p, g_j^p)$  propagated from  $S_j \neq S_i$  with  $g_j^p = 2$ , set  $\text{GaveOutput}_i = \top$ , and output  $C_i = B_j^p$ . The output justifier is the justifier for  $(B_j^p, g_j^p)$ .

**Fig. 4.**  $\Pi_{\text{SELECTBLOCK}}$

**Theorem 4.** *If  $t < n/2$  then  $\Pi_{\text{SELECTBLOCK}}$  is a Justified Select Block Protocol. When  $\beta = \sum_{i=1}^n |B_i|$  and using the protocol  $\Pi_{\text{GRADEDSELECTBLOCK}}$  from Fig. 3 as sub-protocol the complexity is*

$$\mathcal{O}(\beta \text{IN} + n^2 \log(n) \text{RS} + n \log(n) \text{RS}_{\#} + \text{ELECT}) .$$

*Proof.* We start with the complexity. In Step 2 the servers need only CauCast the inputs  $B_i$  as  $B_i^0$  and  $g_i^0$  can be computed from  $B_i$ . This gives  $\beta_{\text{IN}}$ . Each run of  $\Pi_{\text{GRADEDSELECTBLOCK}}$  from Fig. 3 has complexity  $\mathcal{O}(n^2 \log(n) \text{RS} + n \log(n) \text{RS}_{\#} + \text{ELECT})$ , where we ignore the IN component as we are CauCasting computed values which have known message identifiers. Besides this the protocol only CauCasts computed values for which the receiver knows the message identifier, so there is no more information to CauCast. The protocol terminates in expected  $\mathcal{O}(1)$  rounds as argued below. This gives the desired complexity. Liveness follows from Positive Agreement: at some point some honest server will have  $g_i^r = 2$  and then the protocol will eventually terminate by construction of the echo rules. Justified Outputs is clear by the Justified Output rule of  $\Pi_{\text{GRADEDSELECTBLOCK}}$  which maintains that  $J_{\text{IN}}(B_i^r) = \top$  for all  $r$ . We then argue Agreement. Assume that some  $S_i$  outputs  $C_i$ . Then it saw a justified  $(B_j^r = C_i, 2)$ . Let  $r$  be the smallest  $r$  for which a justified  $(B_j^r, 2)$  was seen by an honest server. Then by graded agreement all justified  $(B_j^r, g)$  for round  $r$  will have  $B_j^r = C_i$ . Therefore, by Stability, it holds for all justified  $(B_j^\rho, g)$  for rounds  $\rho \geq r$  that  $B_j^\rho = C_i$ . Now consider any other honest server  $S_k$  which outputs  $C_k$ . Then it saw some justified  $(B_j^{r'} = C_k, 2)$ . Since we picked  $r$  to be minimal we have that  $r' \geq r$ . From this it follows that  $B_j^{r'} = C_i$ . Ergo  $C_j = C_i$ .  $\square$

### 3.5 Justified Agreement on a Core Set

We then present a protocol for Justified Agreement on a Core Set (JACS). It just lets each server propose a set and then picks  $n - t$  of them.

**Definition 16 (Justified Agreement on a Core Set).** *A protocol for  $n$  servers  $S_1, \dots, S_n$  with input and output justifiers  $J_{\text{IN}}$  and  $J_{\text{OUT}}$ . All honest  $S_i$  have an input  $B_i$  for which  $J_{\text{IN}}(B_i) = \top$  at the time of input.*

**Liveness:** *If all honest servers start running the protocol with a  $J_{\text{IN}}$ -justified input then eventually all honest servers have a  $J_{\text{OUT}}$ -justified output.*

**Validity:** *For all possible  $J_{\text{OUT}}$ -justified outputs  $U$  and all honest  $S_i$  and all  $(S_i, B_i) \in U$  it holds that  $S_i$  had input  $B_i$ .*

**Justified Blocks:** *For all possible justified outputs  $U$  and all (potentially corrupt)  $S_i$  and all  $(S_i, B_i) \in U$  it holds that  $J_{\text{IN}}(B_i) = \top$ .*

**Agreement:** *For all possible justified outputs  $U_i$  and  $U_j$  it holds that  $U_i = U_j$ .*

**Large Core:** *For all possible justified outputs  $U$  it holds that  $|S| \geq n - t$ .*

1. The input of  $S_i$  is  $B_i$  with  $J_{\text{IN}}(B_i) = \top$ .
2. Server  $S_i$  CauCasts  $B_i$ . This message is justified by  $J_{\text{IN}}(B_i) = \top$  and  $B_i$  having been reliably broadcast by  $S_i$ .
3. Server  $S_i$  collects at least  $n - t$  justified  $B_j$  from servers  $S_j \in P_i$  and lets  $U_i = \{(S_j, B_j)\}_{S_j \in P_i}$ . This value is justified by each  $B_j$  being justified and  $|U_i| \geq n - t$ . CauCast  $U_i$ .
4. Run  $\Pi_{\text{SELECTBLOCK}}$  where  $S_i$  inputs  $U_i$ . The input justifier of  $\Pi_{\text{SELECTBLOCK}}$  is to check that  $U_i$  is justifiable as defined in the above step.
5. Server  $S_i$  gets output  $C_i$  from  $\Pi_{\text{SELECTBLOCK}}$  and outputs  $C_i$ . The output justifier is that  $C_i$  is a justified output from the above  $\Pi_{\text{SELECTBLOCK}}$ .

**Fig. 5.** Protocol to Agree on a Core Set  $\Pi_{\text{ACS}}$

**Theorem 5.** *If  $t < n/2$  then  $\Pi_{ACS}$  is a JACS protocol. When  $\beta = \sum_{i=1}^n |B_i|$  and when using  $\Pi_{SELECTBLOCK}$  from Fig. 4 as sub-protocol the complexity is*

$$\mathcal{O}(\beta \text{IN} + n^2 \log(n) \text{RS} + n \log(n) \text{RS}_{\#} + \text{ELECT}) .$$

*Proof.* Safety and liveness properties follows directly from those of  $\Pi_{SELECTBLOCK}$ . We address the complexity. The CauCast of  $B_i$  costs  $\beta \text{IN}$ . CauCasting  $U_i$  costs  $n^2 \text{RS}$  to specify the sets  $P_i$ . Running  $\Pi_{SELECTBLOCK}$  from Fig. 4 costs expected  $\mathcal{O}(n^2 \log(n) \text{RS} + n \log(n) \text{RS}_{\#} + \text{ELECT})$ , where we ignore the  $\text{IN}$  component as it is run on computed messages. There are no further costs.  $\square$

### 3.6 Total Order Subcast

We now give a protocol for total order subcast assuming reliable subcast. The protocol uses an erasure code with reconstruction from  $n - t$  code words out of  $n$ . It also uses a PKI for  $\mathbb{G}$ . We first discuss some hairy details of how transactions are collected to not clutter the protocol description with these.

**Definition 17 (Transaction identifiers, Blocks, Message Sets).** *The first item establishes a FIFO delivery on reliable broadcast of transactions. The second item establishes a causal order on transaction sets using vector clocks. The last item collects all messages in the causal past of a transaction set.*

- In the  $\Pi_{TOS}$  protocol the servers will send around so-called transaction identifier  $(j, c_j)$  identifying transaction number  $c_j$  by  $\mathbf{G}_j$ . We say that  $(j, c_j)$  is justified at  $\mathbf{S}_i$  if it saw that some  $(j, c_j, m_j)$  was reliable subcast by  $\mathbf{G}_j$  and  $c_j = 1$  or  $(j, c_j - 1)$  is similarly justified at  $\mathbf{S}_i$ . This means that the transaction identifiers become justified in order  $c_j = 1, 2, \dots$  and that when  $(j, c_j)$  is justified, the message  $m_j$  is known. Let  $\text{Msg}(j, c_j) = m_j$ .
- Servers  $\mathbf{S}_i$  will send out (and relay) blocks  $(B_i^e, w_i^e)$ , where  $e$  is an epoch number,  $B_i^e$  is a set of transaction identifiers  $(j, c_j)$ , and  $w_i^e = (v_i^1, \dots, v_i^n)$  is a vector clock. The set  $B_i^e$  is justified at  $\mathbf{S}_k$  if all  $(j, c_j) \in B_i^e$  are justified and  $\mathbf{S}_k$  already received previous blocks  $((B_1^{v_i^1}, w_1^{v_i^1}), \dots, (B_n^{v_i^n}, w_n^{v_i^n}))$  from  $(\mathbf{S}_1, \dots, \mathbf{S}_n)$  which are similarly justified.<sup>3</sup>
- Justified blocks define message sets as follows. For justified  $B_i^e$  let  $\text{Msg}(B_i^e) = \{\text{Msg}(j, c_j) \mid (j, c_j) \in B_i^e\}$ . Let  $\text{Msg}(B_i^0, w_i^0) = \emptyset$  and for  $e > 0$  and justified  $(B_i^e, w_i^e)$  let  $\text{Msg}(B_i^e, w_i^e) = \text{Msg}(B_i^e) \cup \bigcup_{j=1}^n \text{Msg}(B_j^{v_i^j}, w_j^{v_i^j})$ .

**Theorem 6.** *If  $t < n/2$  then protocol  $\Pi_{TOS}$  is a total order subcast. When  $\beta$  is the total bit-length of all inputs,  $\iota$  is the number of inputs, and  $\rho$  is the number of epochs in the protocol, and when using  $\Pi_{ACS}$  from Fig. 5 as sub-protocol, and assuming that on average inputs have length at least  $\kappa \log(\kappa)$ , the communication complexity is*

$$\mathcal{O} \left( \beta \text{RS} + \iota \cdot \text{RS}_{\#} + \rho \cdot \left( n^2 \log(n) \text{RS} + n \log(n) \text{RS}_{\#} + \text{ELECT} \right) \right) .$$

*Proof.* The ledger is clearly monotone. Consider Agreement. By Agreement of  $\Pi_{ACS}$  there is agreement on  $U^e$ . As  $\text{Msg}$  is a function this implies agreement on  $M^e$ , and thus agreement on Ledger

<sup>3</sup> The blocks with  $e = 0$  are empty and always justified, so the recursion ends at  $e = 0$ .

**Init:** All  $S_i$  lets  $\text{Ledger}_i = ()$ ,  $\text{Scheduled}_i = \emptyset$ ,  $\text{OnGoing}_i = \perp$ , and  $e_i = 0$ . Define  $B_i^0 = \emptyset$ ,  $w_i^0 = (0, \dots, 0)$  and that  $(B_i^0, w_i^0)$  has been CauCast received from  $S_i$  by all  $S_j$  already. For all  $G_j \in \mathbb{G}$  let  $c_j = 0$ .

**Input:**  $G_i$ : On input  $m$  at  $G_i$  where  $J_{\text{IN}}(m) = \top$  let  $c_i \leftarrow c_i + 1$  and reliably subcast  $(c_i, m)$ .

**Schedule:**  $S_i$ : On arrival of RB of  $(c_j, m)$  from  $G_j$  where  $c_j = 1$  or  $(j, c_j - 1) \in \text{Scheduled}_i$ , add  $(j, c_j)$  to  $\text{Scheduled}_i$ .

**Keep vector clock:**  $S_i$ : Throughout the protocol, for each  $S_j$  let  $v_i^j$  be the largest integer such that  $B_j^{v_i^j}$  was received via CauCast, in particular  $v_i^i = e_i - 1$ . Let  $w_i := (v_i^1, \dots, v_i^n)$ . Let  $\text{InBlocks}_i := \bigcup_j \text{Msg} \left( B_j^{v_i^j} \right)$ .

**Propose Next Block:** Server  $S_i$ : If  $\text{OnGoing}_i = \perp$  and  $\text{Scheduled}_i \setminus \text{InBlocks}_i \neq \emptyset$ , let  $\text{OnGoing}_i = \top$  and atomically do the following:<sup>a</sup>

1. Let  $e_i \leftarrow e_i + 1$ .
2. Let  $B_i^{e_i} = \text{Scheduled}_i \setminus \text{InBlocks}_i$ ,  $w_i^{e_i} = w_i$ , and CauCast  $(B_i^{e_i}, w_i^{e_i})$  justified as explained in Definition 17.<sup>b</sup>
3. Start  $\Pi_{\text{ACS}}^{e_i}$  with input  $(B_i^{e_i}, w_i^{e_i})$  with the input justifier being that the value was CauCast.

**Extend Ledger:** If  $\text{OnGoing}_i = \top$  and  $\Pi_{\text{ACS}}^{e_i}$  produces output  $U^{e_i} = \{(S_j, (B_j^{e_i}, w_j^{e_i}))\}_{S_j \in P}$ , then let  $M^{e_i} = \bigcup_{S_j \in P} \text{Msg}(B_j^{e_i}, w_j^{e_i})$ , sort  $M^{e_i} \setminus \text{Ledger}_i$  using some deterministic rule to get a list  $M$ , and let  $\text{Ledger}_i \leftarrow \text{Ledger}_i \parallel M$ . Then let  $\text{OnGoing}_i = \perp$ .

<sup>a</sup> Here any non-deadlocking condition  $\text{Wait}_i$  can be added to let  $\text{Scheduled}_i \setminus \text{InBlocks}_i$  grow to some bigger size.

<sup>b</sup> Note that this involves only RB'ing identifiers of the values justifying  $(B_i^{e_i}, w_i^{e_i})$ . We discuss as part of analysing communication complexity exactly which values need to be sent.

**Fig. 6.** Totally Ordered Subcast Protocol  $\Pi_{\text{TOS}}$ .

by a simple induction. We then look at liveness. Liveness of subcast means that  $(j, c_j)$  eventually ends up in  $B_i^{e_i}$  at all honest  $S_i$ . This  $B_i^{e_i}$  will eventually reach all honest  $S_j$  which will set  $v_j^i \geq e_i$ . After that  $m$  is in  $\text{Msg}(B_j^{e_j}, w_j^{e_j})$  for all honest  $S_j$ . In the next run of  $\Pi_{\text{ACS}}$  message  $m_j$  will then be in  $\text{Msg}(B_j^{e_j}, w_j^{e_j})$  for all honest  $S_j \in P$ . Since  $|P| = n - t \geq t + 1$  there is an honest  $S_j$  in  $P$ , so  $m$  will be in  $M^{e_i}$  and then  $\text{Ledger}_i$ . We count complexity. Each input is reliably subcast. This is  $\mathcal{O}(\beta \text{RS} + \iota \text{RS}_{\#})$ . To CauCast the blocks we only need to CauCast the transaction identifiers and  $w_i$ , as it is clear from  $w_i$  which other values are needed to justify the block. We count the total length of the sets in all blocks, i.e.,  $\sum_{i,e} |\text{enc}(B_i^e)|$  for an asymptotically optimal encoding of the set of transaction identifiers. For each input we add  $(j, c_j)$  to at most  $n$  blocks. Since  $j$  and  $c_j$  are counters they will be  $\mathcal{O}(\text{poly}(\kappa))$  in any poly-time run, so we represent them with  $\mathcal{O}(\log(\text{poly}(\kappa))) = \mathcal{O}(\log(\kappa))$  bits. Therefore  $\sum_{i,e} |\text{enc}(B_i^e)| = \mathcal{O}(\iota n \log(\kappa))$ . This overall adds  $\mathcal{O}(\iota n \log(\kappa) \text{RS})$ . We assumed that inputs have length at least  $\kappa \log(\kappa)$  and that  $n = \Theta(\kappa)$ . Therefore  $\iota n \log(\kappa) = \mathcal{O}(\beta)$ , so the contribution is  $\mathcal{O}(\beta \text{RS})$ . In each epoch  $S_i$  CauCasts  $B_i$  along with the vector clock which also has size  $n \log(\kappa)$ . Ignoring the size of the blocks, which we already accounted for, this adds up to  $\mathcal{O}(\rho n^2 \log(\kappa) \text{RS} + \rho n \text{RS}_{\#})$ . In each of  $\rho$  epochs  $\Pi_{\text{ACS}}$  contributes an extra  $\mathcal{O}(n^2 \log(n) \text{RS} + n \log(n) \text{RS}_{\#} + \text{ELECT})$ . Sum and use that  $n = \Theta(\kappa)$ .  $\square$

### 3.7 Total Order Broadcast

We finally note that from TOS we can get TOB. We will use the protocol  $\Pi_{\text{TOS}}$  in a white-box manner. After each epoch we will outcast the new part of the ledger to the ground population to let them learn the update. The protocol is given in Fig. 7.

**Init:** All  $G_i \in \mathbb{S}$  keeps an ordered list  $\text{Ledger}_i$  as part of  $\Pi_{\text{TOS}}$ . All  $G_i \in \mathbb{G} \setminus \mathbb{S}$  will keep their own  $\text{Ledger}_i$ , initially empty.

**Input:** On input  $m$  at  $G_i$  input  $m$  to  $\Pi_{\text{TOS}}$ .

**Outcast:**  $S_i \in \mathbb{S}$ : When computing  $\text{Ledger}_i \leftarrow \text{Ledger}_i \| M$  in  $\Pi_{\text{TOS}}$  in epoch  $e_i$  input  $M$  to  $\Pi_{\text{RELIABLEOUTCAST}}$  with session identifier  $e_i$ .

**Extend Ledger:**  $G_j \in \mathbb{G} \setminus \mathbb{S}$ : On output  $M$  from  $\Pi_{\text{RELIABLEOUTCAST}}^{e_i}$  let  $\text{Ledger}_i \leftarrow \text{Ledger}_i \| M$ .

**Fig. 7.** Total-Order Broadcast  $\Pi_{\text{TOB}}$

**Theorem 7.** *If  $t < n/2$  then protocol  $\Pi_{\text{TOB}}$  is a total order broadcast. When  $\beta$  is the total bit-length of all inputs,  $\iota$  is the number of inputs, and  $\rho$  is the number of epoch in  $\Pi_{\text{TOS}}$ , and when using  $\Pi_{\text{TOS}}$  from Fig. 6 and  $\Pi_{\text{OUTCAST}}$  from Section 2.2 as sub-protocol, and assuming that on average inputs have length at least  $\kappa \log(\kappa)$ , the communication complexity is*

$$\mathcal{O}\left(\beta(RS+N) + \iota \cdot RS_{\#} + \rho \cdot \left(n^2 \log(n)(RS+N) + n \log(n) RS_{\#} + \text{ELECT}\right)\right).$$

*Proof.* The sub-protocol  $\Pi_{\text{TOS}}$  has communication complexity

$$\mathcal{O}\left(\beta RS + \iota \cdot RS_{\#} + \rho \cdot \left(n^2 \log(n) RS + n \log(n) RS_{\#} + \text{ELECT}\right)\right)$$

and  $\Pi_{\text{OUTCAST}}$  has communication complexity  $\mathcal{O}(nN \log(n)\kappa + N|M|)$  for each  $M$ . We run  $\Pi_{\text{OUTCAST}}$  once each epoch and on  $M$ 's of total length  $\beta$ . This gives a total contribution of  $\mathcal{O}(\rho nN \log(n)\kappa + N\beta)$  from the outcasting. Then use that  $\rho nN \log(n) = \Theta(\rho N n^2 \log(n))$  and sum.  $\square$

## 4 RB with Subsampling

We now present a reliable subcast protocol for  $N$  ground members and  $n$  servers. We assume that at most  $T < N/3$  ground members are corrupt and at most  $t < n/2$  servers are corrupt.

As demonstrated in [ACKN23] this setting can be used to get sub-quadratic communication with good constants in a setting with static security from the assumption that less than  $N/3$  ground members are corrupt. The basic idea is that from the  $N$  ground members one samples a uniformly random subset  $\mathbb{S}$  of size  $n = \Theta(\kappa)$ . Since we are sampling from a set with a supermajority of honest servers and only need that  $\mathbb{S}$  has a majority of honest servers the size of  $\mathbb{S}$  can be practical. In our setting the set  $\mathbb{S}$  could be constructed simply by running  $n = |\mathbb{S}|$  justified leader elections in parallel and wait for all of them and let  $\mathbb{S}$  be the  $n$  winners. A party elected multiple times can be handled by proceeding with a smaller  $\mathbb{S}$  or letting the party run multiple servers. There are, however, many different ways in the literature for doing sub-sampling of committees. We therefore leave the concrete subsampling out of the description and just assume that it has been done and that there are at most  $t < n/2$  corruptions in  $\mathbb{S}$ . For now we just remark that if committees are sampled with replacement from the ground population with equal probability, the amount of honest parties in the committee follows the binomial distribution. In [ACKN23] this is used to show that 653 is the minimal committee size needed to get honest majority with 60 bits of statistical security, when assuming less than a third corrupted parties in the ground population. Using the same method—checking that the cumulative distribution function of the honest parties getting elected for any number  $\leq n/2$  of seats on the committee is negligible—we compute (cf. Appendix A) the minimal secure committee sizes for various choices of statistical security parameter ( $\sigma$ ) and maximal fraction

**Setup:** We assume a setup for a threshold signature scheme with verification key  $\text{vk}$  being public and each  $G_i$  holding key share  $\text{sk}_i$ . The reconstruction threshold is  $N - T$ . The protocol also use a Merkle-tree scheme  $(\text{HashTree}, \text{Path}, \text{VerPath})$  and uses an erasure code  $\text{EC} = (\text{Enc}, \text{Dec})$  with  $n$  code words and reconstruction threshold  $n - t$ .

**Input:**  $G_j$ : On input  $(\text{mid}, m)$  with  $G^{\text{mid}} = G_j$  (i.e.,  $G_j$  is the designated sender of  $\text{mid}$ ) and  $J^{\text{mid}}(m) = \top$  let  $(m_1, \dots, m_n) = \text{EC.Enc}(m)$ , let  $h = \text{HashTree}(m_1, \dots, m_n)$ , and send  $(\text{mid}, h)$  to all  $G_k$ .

**SubSign**  $G_k$ : On receiving  $(\text{mid}, h)$  from  $G_j = G^{\text{mid}}$ , where no  $(\text{mid}, \cdot)$  was received from  $G_j$  before and  $J^{\text{mid}}(m) = \top$ , send  $\sigma_k = \text{Sig}_{\text{sk}_k}((\text{mid}, h))$  to  $G_j$ .

**Send Shards:**  $G_j$ : On having received  $N - T$  valid signature shares compute  $\sigma = \text{Sig}_{\text{sk}}((\text{mid}, h))$  using Combine and send  $(\text{mid}, h, \sigma, m_i, \text{Path}_i = \text{Path}(m_1, \dots, m_n, i))$  to  $S_i$ .

**Echo and Record own Shard:**  $S_i$ : On having received  $(\text{mid}, h, \sigma, m_i, \text{Path}_i)$  with  $\text{Ver}_{\text{vk}}((\text{mid}, h), \sigma) = \top$  and  $\text{VerPath}(h, i, m_i, \text{Path}_i) = \top$  from  $G_j = G^{\text{mid}}$  or some server  $S_j \in \mathbb{S}$  send  $(\text{mid}, h, \sigma, m_i, \text{Path}_i)$  to all other servers and record  $(\text{mid}, h, i, m_i)$ .

**Record other Shards:**  $S_j$ : On having received  $(\text{mid}, h, \sigma, m_i, \text{Path}_i)$  with  $\text{Ver}_{\text{vk}}((\text{mid}, h), \sigma) = \top$  and  $\text{VerPath}(h, i, m_i, \text{Path}_i) = \top$  from  $S_i$  record  $(\text{mid}, h, i, m_i)$ . Note that we deliberately do not echo here.

**Combine Shards:**  $S_j$ : On having recorded  $(\text{mid}, h, i, m_i)$  for  $n - t$  servers  $S_i$  for the same  $\text{mid}$ , call the set of these servers  $S$ , compute  $m = \text{EC.Dec}(\{(i, m_i)\}_{S_i \in S})$  and  $(m_1, \dots, m_n) = \text{EC.Enc}(m)$  and  $h' = \text{HashTree}(m_1, \dots, m_n)$ . If  $h' = h$  then output  $(\text{mid}, m)$  and for each  $S_i \in \mathbb{S}$  send  $(\text{mid}, h, \sigma, m_i, \text{Path}_i)$  to  $S_i$ .

**Fig. 8.** RSS: A protocol for reliable subcast with  $N$  ground members,  $n$  servers, at most  $T < N/3$  corrupt ground members, and at most  $t < n/2$  corrupt servers.

of corruption tolerated in the ground population ( $c$ ). The results are shown in Table 1. In particular it shows that a committee of 173 parties would be resilient against a  $1/5$  of the ground population being corrupt with 60 bits of security. This compares very favourably against the Algorand setting in which less than a  $1/5$  of the ground population is assumed to be corrupted in order to get 56 bits of security with committees of 6000 parties (cf. [BBK<sup>+</sup>23]).

$\sigma$	30			40			60			80		
$c$	$1/5$	$1/4$	$1/3$	$1/5$	$1/4$	$1/3$	$1/5$	$1/4$	$1/3$	$1/5$	$1/4$	$1/3$
$n$	81	127	307	111	173	423	173	269	653	235	363	887

**Table 1.** Minimal committee sizes ( $n$ ) that guarantee an honest majority (except with probability  $2^{-\sigma}$ ) when parties are sampled according to the binomial distribution from a ground population in which less than a fraction  $c$  of the parties are corrupted.

**Theorem 8.** *If  $t < n/2$  and  $T < N/3$ , then RSS is a reliable subcast. If messages have length  $\beta \geq \kappa \log(\kappa)$ , then the complexity is  $\mathcal{O}(\beta n + N\kappa + n^2 \log(n)\kappa)$ , which is of the form  $\mathcal{O}(\beta RS + RS_{\#})$  for  $RS = n$  and  $RS_{\#} = N\kappa + n^2 \log(n)\kappa$ . For large populations  $N \geq n^2 \log(n)$  or large messages  $\beta \geq n \log(n)\kappa$  the complexity is  $\mathcal{O}(\beta n + N\kappa)$  such that  $RS_{\#} = N\kappa$ .*

*Proof.* We argue agreement. If a server accepts  $(\text{mid}, h, \sigma, m_i, \text{Path}_i)$  then it was signed by  $N - T$  ground members and therefore at least  $T + 1$  honest ground members. Therefore  $h$  is unique for  $\text{mid}$ . And if  $S_j$  outputs  $(\text{mid}, m)$  then  $h = \text{HashTree}(\text{EC.Enc}(m))$  and therefore they all output the same  $m$  if they output something. Eventual Output 1 is trivial. We argue eventual Output 2. Assume some honest  $S_j$  outputs  $(\text{mid}, m)$ . Then it sends  $(\text{mid}, h, \sigma, m_i, \text{Path}_i)$  to each  $S_i$  and therefore each honest  $S_i$  sends  $(\text{mid}, h, \sigma, m_i, \text{Path}_i)$  to all servers in **Echo and Record own Shard**. There are  $n - t$  honest



servers doing this, so all honest will end up recording  $n-t$  shards and output  $(\text{mid}, m)$  in **Combine Shards**. We count complexity. Getting the signature shares from  $N$  ground members  $G_j$  costs  $N\kappa$ . Each  $m_i$  and  $\text{Path}_i$  it sent to and from  $S_i$  at most  $\mathcal{O}(n)$  times which contributes  $\mathcal{O}(n|m_i| + n \log(n)\kappa)$ . Summing over  $S_i$  this gives  $\mathcal{O}(n|m| + n^2 \log(n)\kappa)$ . For large ground populations  $N \geq n^2 \log(n)$  we have that  $\mathcal{O}(\beta n + N\kappa + n^2 \log(n)\kappa) = \mathcal{O}(\beta n + N\kappa + N\kappa) = \mathcal{O}(\beta n + N\kappa)$ . For large messages  $\beta \geq n \log(n)\kappa$  we have that  $\mathcal{O}(\beta n + N\kappa + n^2 \log(n)\kappa) = \mathcal{O}(\beta n + N\kappa + n\beta) = \mathcal{O}(\beta n + N\kappa + n\beta)$ .  $\square$

**Corollary 1.** *Protocol  $\Pi_{TOS}$  when using  $\Pi_{ACS}$  for ACS,  $\Pi_{CONSTANTINE}$  for leader election, and RSS for RB is a total order subcast. Let  $\beta$  be the total length of inputs and let  $\iota$  be the number of inputs. Then the communication complexity is*

$$\beta n + \iota \cdot (N\kappa + n^2 \log(n)\kappa) + \rho \cdot (Nn \log(n)\kappa + n^3 \log^2(n)\kappa) . \quad (1)$$

*If there is a large ground population  $N \geq n^2 \log(n)$  or messages which are large on average (i.e.  $\beta/\iota \geq n \log(n)\kappa$ ), and if the protocol on average consumes at least  $n \log(n)$  messages per epoch, then the complexity is*

$$\beta n + \iota N\kappa . \quad (2)$$

*Proof.* The complexity is  $\beta \text{RS} + \iota \cdot \text{RS}_{\#} + \rho \cdot (n^2 \log(n) \text{RS} + \text{ELECT} + n \log(n) \text{RS}_{\#})$ . Plug in  $\text{RS} = n$ ,  $\text{RS}_{\#} = N\kappa + n^2 \log(n)\kappa$ , and  $\text{ELECT} = Nn\kappa$  to get

$$\beta n + \iota \cdot (N\kappa + n^2 \log(n)\kappa) + \rho \cdot (Nn \log(n)\kappa + n^3 \log^2(n)\kappa) .$$

If on average we consume  $n \log(n)$  messages per epoch then  $\rho n \log(n) \leq \iota$ , so we get

$$\rho \cdot (Nn \log(n)\kappa + n^3 \log^2(n)\kappa) \leq \iota \cdot (N\kappa + n^2 \log(n)\kappa) .$$

So we can drop the LHS asymptotically. If  $N \geq n^2 \log(n)$  then  $N\kappa \geq n^2 \log(n)\kappa$ . Equivalently, if  $\beta/\iota \geq n \log(n)\kappa$  then  $\beta n \geq \iota \cdot n^2 \log(n)\kappa$ . In both cases we can drop  $\iota \cdot n^2 \log(n)\kappa$  asymptotically.  $\square$

**Corollary 2.** *Protocol  $\Pi_{TOB}$  when using  $\Pi_{TOS}$  from Corollary 1 for TOS and  $\Pi_{OUTCAST}$  for outcasting is a TOB. Let  $\beta$  be the total length of inputs and let  $\iota$  be the number of inputs. Then the communication complexity is*

$$\beta N + \iota \cdot (N\kappa + n^2 \log(n)\kappa) + \rho \cdot (Nn \log(n)\kappa + n^3 \log^2(n)\kappa) . \quad (3)$$

*For a large ground population  $N \geq n^2 \log(n)$  or messages which are large on average (i.e.,  $\beta/\iota \geq n \log(n)\kappa$ ), and if the protocol on average consumes at least  $n \log(n)$  messages per epoch, then the complexity is*

$$\beta N . \quad (4)$$

*Proof.* The contribution of outcasting is  $\mathcal{O}(\beta N + \rho Nn \log(n)\kappa)$ . The term  $\rho Nn \log(n)\kappa$  is already dominated by Eq. (1). Adding  $\beta N$  to Eq. (1) gives Eq. (3). Adding  $\beta N$  to Eq. (2) gives  $\beta N + \iota N\kappa$ . Then use that  $\iota\kappa \leq \beta$  to get Eq. (4).  $\square$

## 5 RB with Dual Threshold and Asymmetric Synchrony Assumptions

We now present a RB for the servers only. We will not focus on communication complexity but resilience. The protocol can be turned into a communication efficient reliable subcast using threshold signatures and the sharding technique from RSS, but we leave this out of the description to focus on the main contribution. The protocol will have two corruption thresholds  $t_A$  and  $t_S$  where  $t_A \leq t_S \leq n/2$  and where there are at most  $t_S$  corruptions.

The protocol uses one timeout per server—server  $S_i$  waits  $\Delta_{\text{WAIT}}^i$  seconds. The protocol has the property that if there are at most  $t_A$  corruptions then the running time of protocol does not depend on the  $\Delta_{\text{WAIT}}^i$ . However, if the actual number of corruptions  $t$  is in the interval  $t_A < t \leq t_S$  then the running time of the protocol does depend on the  $\Delta_{\text{WAIT}}^i$ .

**Input** We assume a PKI for a signature scheme. The input of the designated sender  $S_s$  is  $m$  with  $J_{\text{IN}}(m) = \top$ . In response to this input  $S_s$  sends  $(m, \sigma_s = \text{Sig}_{\text{sk}_s}(m))$  to all servers. Create initially empty sets SignedAsync and SignedSync.

**Asynchronous Echo**  $S_i$ : On receiving  $(m, \sigma_s)$  from  $S_s$ , where  $J_{\text{IN}}(m) = \top$  and  $\text{Ver}_{\text{vk}_s}(m, \sigma_s) = \top$  and where no such message was received before and there is no  $(S_j, m_j, \sigma_j) \in \text{SignedAsync}$  with  $m_j \neq m$ , proceed as follows:

1. Let  $\sigma_i = \text{Sig}_{\text{sk}_i}(\text{ASYNC}, m)$  and send  $(m, \sigma_s, \sigma_i)$  to all servers.
2. Add  $(S_i, m, \sigma_i)$  to SignedAsync.
3. Set a timeout  $\text{Timeout}(\text{COLLECTFROMHONEST}, \Delta_{\text{WAIT}}^i)$ .

**Collect Asynchronous Echos** All servers: On receiving  $(m_j, \sigma_s, \sigma_j)$  from  $S_j$ , where  $\text{Ver}_{\text{vk}_s}(m_j, \sigma_s) = \top$  and  $\text{Ver}_{\text{vk}_j}(\text{ASYNC}, m_j, \sigma_j) = \top$  and no such value was received from  $S_j$  before, add  $(S_j, m_j, \sigma_j)$  to SignedAsync.

**Synchronous Echo**  $S_i$ : If COLLECTFROMHONEST occurred and there exists  $m$  such that there are at least  $n - t_S$  values  $(S_j, m, \cdot) \in \text{SignedAsync}$  and there does not exist  $(S_k, m', \cdot) \in \text{SignedAsync}$  where  $m' \neq m$ , then let  $\sigma_i = \text{Sig}_{\text{sk}_i}(\text{SYNC}, m)$  and send  $(m, \sigma_i)$  to all servers.

**Collect Synchronous Echos** All servers: On receiving  $(m_j, \sigma_j)$  from  $S_j$ , where  $\text{Ver}_{\text{vk}_j}(\text{SYNC}, m_j, \sigma_j) = \top$  and no such value was received from  $S_j$  before, add  $(S_j, m_j, \sigma_j)$  to SignedSync.

**Asynchronous Output** All servers: If there exists  $m$  such that there are  $n - t_A$  values  $(S_j, m, \sigma_j) \in \text{SignedAsync}$  then let  $\Sigma = \{(S_j, \sigma_j)\}_{(S_j, m, \sigma_j) \in \text{SignedAsync}}$ , output  $m$ , send  $(m, \Sigma)$  to all servers, and terminate.

**Synchronous Output** All servers: If there exists  $m$  such that there are  $n - t_S$  values  $(S_j, m, \sigma_j) \in \text{SignedSync}$  then let  $\Sigma = \{(S_j, \sigma_j)\}_{(S_j, m, \sigma_j) \in \text{SignedSync}}$ , output  $m$ , send  $(m, \Sigma)$  to all servers, and terminate.

**Output by Relay** On receiving  $(m, \Sigma)$  from any server where either

- $\Sigma$  contains  $n - t_S$  values  $(S_j, m, \sigma_j)$  for distinct  $S_j$  such that  $\text{Ver}_{\text{vk}_j}(\text{SYNC}, m, \sigma_j) = \top$  (call such a value synchronous-valid), or
- $\Sigma$  contains  $n - t_A$  values  $(S_j, m, \sigma_j)$  for distinct  $S_j$  such that  $\text{Ver}_{\text{vk}_j}(\text{ASYNC}, m, \sigma_j) = \top$  (call such a value asynchronous-valid),

output  $m$ , send  $(m, \Sigma)$  to all servers, and terminate.

**Fig. 9.** RBD: A protocol for RB with dual thresholds  $t_S$  and  $t_A$ . For conciseness we do not explicitly mentioning the messages identifier mid and we let  $S_s = S^{\text{mid}}$  and let  $J_{\text{IN}} = J^{\text{mid}}$ .

The protocol can be proven secure in two settings. One is a setting where the network is always synchronous and where  $t_S + 2t_A < n$ . In this setting we need that  $\Delta_{\text{WAIT}}^i \geq 2\Delta_{\text{NET}}$ . We call this the *optimistic model*. In the other setting we can tolerate that the network is sometimes asynchronous. Here we only need  $\Delta_{\text{WAIT}}^i \geq \Delta_{\text{NET}}$ . However, we need that  $2t_S + t_A < n$ . As a strengthening we can here tolerate that either the network is synchronous and there are at most  $t \leq t_S$  corruptions or the

network is asynchronous and there are at most  $t \leq t_A$  corruptions. We call this the *fallback model* below.

## 5.1 Asymmetric Synchrony Assumptions

We construct our TOBs for a model with asymmetric synchrony assumptions, which is meant as a model making it easier to implement the needed notion of synchrony in practice. Consider a setting where a group of servers  $S_1, \dots, S_n$  have just been thrown together, maybe sampled at random from a larger ground population. They want to run a synchronous protocol to be able to tolerate  $t < n/2$ . Assume that each  $S_i$  can set a sound timeout length  $\Delta_{\text{GUESS}}^i$ , i.e., all messages sent to  $S_i$  are received within time  $\Delta_{\text{GUESS}}^i$ . This still opens the question of what timeout length to use in the protocol if a common timeout is needed. Note that we cannot broadcast the values  $\Delta_{\text{GUESS}}^i$  to help us pick a common value, as broadcast is the problem we are trying to solve. This motivates implementing TOB in the following model where the parties do not agree on a common timeout value.

**Definition 18 (Asymmetric Synchrony Assumption (ASA) Model).** *The ASA model considers  $n$  servers  $S_1, \dots, S_n$ . Each  $S_i$  gets its own  $\Delta_{\text{GUESS}}^i$  as private input, i.e., the other servers are not given  $\Delta_{\text{GUESS}}^i$ . The adversary schedules messages, but it is guaranteed that all messages sent at time  $t$  from an honest server to an honest server  $S_i$  are delivered no later than at time  $t + \Delta_{\text{GUESS}}^i$ . Note the only messages to  $S_i$  are delivered in time  $\Delta_{\text{GUESS}}^i$ . Messages to other honest servers may be slower. Round complexity is measured in units of  $\Delta_{\text{GUESS}}^{\text{MAX}} := \max_{\text{honest } S_i} \Delta_{\text{GUESS}}^i$  and  $\Delta_{\text{NET}}$ , where  $\Delta_{\text{NET}}$  is the longest it took to send a message from an honest server to an honest server. In the weakly asymmetric synchrony assumption (WASA) model we make the stronger assumption that all messages between honest servers are delivered faster than any honest  $\Delta_{\text{GUESS}}^i$ , i.e.,  $\Delta_{\text{NET}} \leq \Delta_{\text{GUESS}}^{\text{MIN}} := \min_{\text{honest } S_i} \Delta_{\text{GUESS}}^i$ .*

**Definition 19.** *Let  $\text{RBD}^{\text{OPT}}$  be the protocol RBD in Fig. 9 with  $\Delta_{\text{WAIT}}^i = 2\Delta_{\text{GUESS}}^i$  and  $t_A \leq t_S < n/2$  and  $t_S + 2t_A < n$ . Let  $\text{RBD}^{\text{FALLBACK}}$  be the protocol RBD with  $\Delta_{\text{WAIT}}^i = \Delta_{\text{GUESS}}^i$  and  $t_A \leq t_S < n/2$  and  $2t_S + t_A < n$ .*

## 5.2 Synchronous Security with Asynchronous Fallback

We now analyse  $\text{RBD}^{\text{FALLBACK}}$  in the ASA model.

**Theorem 9 (Fallback).**  *$\text{RBD}^{\text{FALLBACK}}$  is a justified RB protocol for the model where either the network is ASA synchronous and there are at most  $t_S$  corruptions or the network is asynchronous and there are at most  $t_A$  corruptions. All servers terminate within time  $2\Delta_{\text{NET}} + \Delta_{\text{GUESS}}^{\text{MAX}}$ . Furthermore, if  $t \leq t_A$  and the sender is honest then all honest servers terminate within time  $2\Delta_{\text{NET}}$ .*

It is not hard to see that the protocol has eventual output and validity. The running time is also straight forward. The main observation is that when  $t \leq t_A$  then within time  $\Delta_{\text{NET}}$  the  $n - t_A$  honest servers trigger **Asynchronous Echo** and then within  $\Delta_{\text{NET}}$  all servers have a asynchronous-valid output. We sketch why the protocol has agreement. The pivotal property which the protocol has by design is the following.

**Lemma 3 (Synchronous Echo Agreement).** *If two honest  $S_i$  and  $S_j$  send  $(m_i, \sigma_i)$  and  $(m_j, \sigma_j)$  in **Synchronous Echo** then  $m_i = m_j$ .*

*Proof.* Consider  $S_i$  and  $S_j$  sending  $(m_i, \sigma_i)$  and  $(m_j, \sigma_j)$  in **Synchronous Echo**. Then obviously they sent some  $(m_i, \sigma_s, \sigma'_i)$  and  $(m_j, \sigma'_s, \sigma'_j)$  in **Asynchronous Echo**. Assume  $S_i$  sent  $(m_i, \sigma_s, \sigma'_i)$  first. Then  $S_j$  started  $\text{Timeout}(\text{COLLECTFROMHONEST}, \Delta_{\text{WAIT}}^j)$  for  $\Delta_{\text{WAIT}}^j = \Delta_{\text{GUESS}}^j$  after  $(m_i, \sigma_s, \sigma'_i)$  was sent. Assume that the network is synchronous. Since we are in the ASA model  $S_j$  thus saw  $(m_i, \sigma_s, \sigma'_i)$  before  $\text{COLLECTFROMHONEST}$  occurred and  $(m_j, \sigma_j)$  was sent. Therefore  $m_j = m_i$ , or  $(m_i, \sigma_s, \sigma'_i)$  would have blocked the sending of  $(m_j, \sigma_j)$ . Assume then that the network is asynchronous. Then by assumption  $t \leq t_A$ . Recall that  $2t_s + t_A < n$ . If  $S_i$  sent  $(m_i, \sigma_i)$  then it saw  $n - t_s$  values  $(S_k, m_i, \cdot) \in \text{SignedAsync}$ . If  $S_j$  sent  $(m_j, \sigma_j)$  then it saw  $n - t_s$  values  $(S_k, m_j, \cdot) \in \text{SignedAsync}$ . This means they saw values from  $n - 2t_s > t_A$  common parties. So they saw  $(S_k, m_i, \cdot)$  and  $(S_k, m_j, \cdot)$  from at least one joint honest  $S_k$ . Therefore  $m_i = m_j$ .  $\square$

**Lemma 4 (Fallback Agreement).** *If  $S_i$  and  $S_j$  are honest and output  $m_i$  and  $m_j$  then  $m_i = m_j$ .*

*Proof.* If  $S_i$  outputs  $m_i$  then it saw a valid  $(m_i, \Sigma_i)$  and if  $S_j$  outputs  $m_j$  then it saw a valid  $(m_j, \Sigma_j)$ . If any of the servers saw a synchronous-valid value, then rename the servers such that  $S_i$  saw one. This gives three cases on the validity flavour of  $(m_i, \Sigma_i)$ - $(m_j, \Sigma_j)$ : synchronous-synchronous, synchronous-asynchronous, and asynchronous-asynchronous. Assume first they both are synchronous-valid. Recall that  $2t_s - t_A < n$ . Among the  $n - t_s$  servers in  $\Sigma_i$  there is at least one honest server as  $n - t_s > t$ , where  $t$  is the actual number of corruptions. Similarly, among the  $n - t_s$  servers in  $\Sigma_j$  there is at least one honest server. Agreement then follows from Lemma 3. Assume then that both  $(m_i, \Sigma_i)$  and  $(m_j, \Sigma_j)$  are asynchronous-valid. Then among the  $n - t_A$  servers in  $\Sigma_i$  and the  $n - t_A$  servers in  $\Sigma_j$  there are at least  $n - t_A - t_A > 2t_s - t_A > t_s \geq t$  common servers. Therefore there is at least one common honest server. Honest servers sign at most one message  $m$ . Assume then that  $(m_i, \Sigma_i)$  is synchronous-valid and  $(m_j, \Sigma_j)$  is asynchronous-valid. Among the  $n - t_s$  servers in  $\Sigma_i$  and the  $n - t_A$  servers in  $\Sigma_j$  there are at least  $n - t_s - t_A > t_s$  common servers. Since  $t_s \geq t$ , where  $t$  is the actual number of corruptions, it follows that there is at least one honest server in common among  $\Sigma_i$  and  $\Sigma_j$ . Clearly, if an honest server signs both  $(\text{SYNC}, m)$  and  $(\text{ASYNC}, m')$  then  $m' = m$ . Therefore  $m_i = m_j$ .  $\square$

### 5.3 Synchronous Security with Optimistic Responsiveness

**Theorem 10 (Optimistic).** *RBD<sup>OPT</sup> is a justified RB protocol for the WASA model with at most  $t_s$  corruptions. All servers terminate within time  $2\Delta_{\text{NET}} + \Delta_{\text{GUESS}}^{\text{MAX}}$ . Furthermore, if  $t \leq t_A$  and the sender is honest then all honest servers terminate within time  $2\Delta_{\text{NET}}$ .*

It is not hard to see that the protocol has eventual output, validity, and the stated time complexity. We sketch why the protocol has agreement. Note that Lemma 3 still holds. Namely, we now wait  $2\Delta_{\text{GUESS}}^i$  instead of  $\Delta_{\text{GUESS}}^i$ , we are in the WASA model which gives stronger guarantees, and the network is always synchronous, so the preconditions used for proving Lemma 3 are all stronger.

**Lemma 5 (Optimistic Agreement).** *If  $S_i$  and  $S_j$  are honest and output  $m_i$  and  $m_j$  then  $m_i = m_j$ .*

*Proof.* As in the proof of Lemma 4 we break into cases. The proofs of the cases synchronous-synchronous and asynchronous-asynchronous can basically be repeated verbatim, so we skip them. This leaves us with the case where  $(m_i, \Sigma_i)$  is synchronous-valid and  $(m_j, \Sigma_j)$  is asynchronous-valid. We have that  $\Delta_{\text{WAIT}}^i \geq 2\Delta_{\text{GUESS}}^i \geq 2\Delta_{\text{NET}}$ , as we are in the WASA model. Since  $n - t_s > t$ , clearly,

at least one honest  $S_k$  signed  $(\text{SYNC}, m_i)$  which in turn implies that it earlier signed  $(\text{ASYNC}, m_i)$ , say at time  $t_i$ . Furthermore, at least one honest  $S_\ell$  signed  $(\text{ASYNC}, m_j)$ , say at time  $t_j$ . Assume for the sake of contradiction that  $m_i \neq m_j$ . If  $t_i \leq t_j - \Delta_{\text{NET}}$  then  $(\text{ASYNC}, m_i)$  would by definition of  $\Delta_{\text{NET}}$  have reached  $S_\ell$  before  $t_j$  and then  $S_\ell$  would by construction not have signed  $(\text{ASYNC}, m_j)$ . So we can assume that  $t_j < t_i + \Delta_{\text{NET}}$ . This means that the  $(\text{ASYNC}, m_j)$  signed by  $S_\ell$  reached  $S_k$  by  $t_j + \Delta_{\text{NET}} < t_i + 2\Delta_{\text{NET}}$ . When  $S_k$  signed  $(\text{ASYNC}, m_j)$  by time  $t_i$  then it did not sign  $(\text{SYNC}, m_i)$  until time  $t_i + 2\Delta_{\text{NET}}$ . But by  $t_i + 2\Delta_{\text{NET}}$  it received  $(\text{ASYNC}, m_j)$ . And then by construction of **Synchronous Echo** and  $m_i \neq m_j$  it will not sign  $(\text{SYNC}, m_i)$ , a contradiction.  $\square$

## 6 Corollaries

In this section we mention a few easy corollaries of our result.

### 6.1 Sub-Quadratic Asynchronous MPC with $T < N/3$

As shown in Lemma 6.1 in [Coh16], given threshold fully homomorphic encryption and total order broadcast and  $t < n/2$  corruptions one can implement asynchronous multiparty computation for  $t < n/2$ . Since we implement total order broadcast among the servers for  $t < n/2$  from RB we can directly run [Coh16] in our framework and get AMPC from RB and  $t < n/2$ . This gives MPC for the fallback model and the optimistic model. We can also run the protocol in the sub-sampling setting with  $N$  ground members and  $n$  servers, with corruption threshold  $t < n/2$  and  $T < N/3$ . To avoid having specialised setup among the servers  $\mathbb{S}$  we can use that [Coh16] uses threshold decryption as a blackbox: it only uses that if all parties agree on a ciphertext  $c$  and that it should be decrypted, then they can eventually learn  $y = \text{Dec}_{\text{sk}}(c)$ . We can therefore secret share  $\text{sk}$  among  $\mathbb{G}$  with reconstruction threshold  $T + 1$  and let them provide decryption as a service for the servers. The servers outcast the encryption  $c$  of the output and the servers send a decryption share to each party. If we have a large ground population  $N \geq n^2 \log(n)$  and they all give one input we can enforce that they give inputs in the same rounds, and then our total order subcast has communication complexity  $\mathcal{O}(\beta n)$ , where  $\beta$  is the total length of the encrypted input. And outcasting  $y$  has complexity  $\mathcal{O}(N|y|)$  and returning the decryption shares has complexity  $Nn|y|$ . This gives sub-quadratic AMPC for  $T < N/3$  corruptions.

**Theorem 11 (informal).** *Assume  $t < n/2$  corruption in  $\mathbb{S}$  and  $T < N/3$  corruptions in  $\mathbb{G}$ , assume  $N \geq n^2 \log(n)$ , let  $f$  be an  $N$ -party function, let  $\beta$  be the total length of the inputs  $x_i$  and let  $\gamma$  be the length of the output  $y = f(x_1, \dots, x_N)$ . Then there is an AMPC protocol for  $\mathbb{G}$  with communication complexity  $\mathcal{O}(n\beta + nN\gamma)$ .*

Note that we can use the expensive  $\Pi_{\text{CONSTANTINE}}$  with communication  $\mathcal{O}(Nn\kappa)$  to implement AMPC and then use the AMPC to do distributed key generation for  $\Pi_{\text{CONSTANTINE}}$  among the  $n$  servers and thereafter do get communication  $\mathcal{O}(n^2\kappa)$ . It is an interesting open problem to propose concretely efficient asynchronous distributed key distributions for the setting with  $t < n/2$ .

### 6.2 Asynchronous Covert TOB with $t < n/2$

We can also get total order broadcast for the model with covert security [AL07]. We assume that every  $G_i$  and  $S_i$  can be Byzantine but does not do anything which will lead to an eventual common detection, where all honest  $S_j$  output a proof that  $S_i$  was corrupted. We only

need to assume that at most  $t < n$  servers are corrupted for the RB to work. Our protocol works as follows. When  $S_i$  sends  $(\text{mid}, m)$  it sends along a signature  $\sigma_i$  on  $(\text{mid}, \text{Hash}(m))$ . On receiving  $(\text{mid}, m, \sigma_i)$  with a valid signature  $S_j$  outputs  $(\text{mid}, m)$  and forwards  $(\text{mid}, h, \sigma_i)$ . On receiving two valid  $(\text{mid}, h', \sigma'_i)$  and  $(\text{mid}, h, \sigma_i)$  with  $h' \neq h$  a server forwards them and outputs  $(S_i \text{ IS CORRUPT: } (\text{mid}, h', \sigma'_i), (\text{mid}, h, \sigma_i))$ . The message  $m$  can be relayed with linear complexity using an erasure code EC as in RSS.

**Theorem 12 (informal).** *Assume that  $t < n$  and the servers are Byzantine but covert. Then CRSS is a secure subcast protocol. If messages have length  $\beta \geq \kappa \log(\kappa)$ , then the complexity is  $\mathcal{O}(\beta n + n^2 \log(n)\kappa)$ , which is of the form  $\mathcal{O}(\beta RS + RS_{\#})$  for  $RS = n$  and  $RS_{\#} = n^2 \log(n)\kappa$ .*

We can use this to get a total order subcast protocol secure against  $t < n/2$  Byzantine corruptions in the covert model with communication  $\beta n + \iota \cdot (n^2 \log(n)\kappa) + \rho \cdot (n^3 \log^2(n)\kappa)$ .

**Input** The input of the designated sender  $S_s$  is  $m$  with  $J_{\text{IN}}(m) = \top$ . In response to this input  $S_s$  sends  $(m, \sigma_s = \text{Sig}_{\text{sk}_s}(m))$  to all servers. Create initially empty set **SignedAsync**.

**Asynchronous Echo**  $S_i$ : On receiving  $(m, \sigma_s)$  from  $S_s$ , where  $J_{\text{IN}}(m) = \top$  and  $\text{Ver}_{\text{vk}_s}(m, \sigma_s) = \top$  and where there are no  $(S_j, m_j, \sigma_j) \in \text{SignedAsync}$  with  $m_j \neq m$  proceed as below, let  $\sigma_i = \text{Sig}_{\text{sk}_i}(\text{ASYNC}(m))$ , send  $(m, \sigma_s, \sigma_i)$  to all servers, and add  $(S_i, m, \sigma_i)$  to **SignedAsync**.

**Collect Asynchronous Echos** All servers: On receiving  $(m_j, \sigma_s, \sigma_j)$  from  $S_j$ , where  $\text{Ver}_{\text{vk}_j}(\text{ASYNC}(m_j), \sigma_j) = \top$  and  $\text{Ver}_{\text{vk}_s}(m_j, \sigma_s) = \top$  and no such value was received from  $S_j$  before, add  $(S_j, m_j, \sigma_j)$  to **SignedAsync**.

**Output** All servers: If there exists  $m$  such that there are  $n - t$  values  $(S_j, m, \sigma_j) \in \text{SignedAsync}$  then let  $\Sigma = \{(S_j, \sigma_j)\}_{(S_j, m, \sigma_j) \in \text{SignedAsync}}$ , output  $m$ , send  $(m, \Sigma)$  to all servers, and terminate.

**Output by Relay** On receiving  $(m, \Sigma)$  from any server where  $\Sigma$  contains  $n - t$  values  $(S_j, m, \sigma_j)$  for distinct  $S_j$  such that  $\text{Ver}_{\text{vk}_j}(\text{ASYNC}(m), \sigma_j) = \top$ , output  $m$ , send  $(m, \Sigma)$  to all servers, and terminate.

**Fig. 10.** RBMA: A protocol for RB against mixed adversaries. For conciseness we do not explicitly mentioning the messages identifier  $\text{mid}$  and we let  $S_s = S^{\text{mid}}$  and let  $J_{\text{IN}} = J^{\text{mid}}$ .

### 6.3 Mixed Adversary ATOB

We can also get ATOB for the model with mixed adversaries. We assume servers can either silently crash or be fully Byzantine corrupted. We assume there are at most  $t_{\text{BYZ}}$  fully Byzantine parties and  $t_{\text{CRASH}}$  additional crash-silent corruptions. We let  $t = t_{\text{BYZ}} + t_{\text{CRASH}}$  in the protocol. The protocol is given in Fig. 10.

**Theorem 13.** *Assume that  $2t_{\text{CRASH}} + 3t_{\text{BYZ}} < n$ . Then RBMA is a justified reliable broadcast protocol.*

*Proof.* Eventual output is trivial as there are  $n - t$  honest parties. Agreement follows from the fact that any two sets  $\Sigma$  and  $\Sigma'$  will overlap on at least  $n - t - t$  parties and from  $2t_{\text{CRASH}} + 3t_{\text{BYZ}} < n$  and  $t = t_{\text{BYZ}} + t_{\text{CRASH}}$  we have that  $n - t - t > t_{\text{BYZ}}$ . Validity follows using similar arguments.  $\square$

## References

- ACKN23. Orestis Alpos, Christian Cachin, Simon Holmgaard Kamp, and Jesper Buus Nielsen. Practical large-scale proof-of-stake asynchronous total-order broadcast. Cryptology ePrint Archive, Paper 2023/1103, 2023. <https://eprint.iacr.org/2023/1103>.

- AJM<sup>+</sup>21. Ittai Abraham, Philipp Jovanovic, Mary Maller, Sarah Meiklejohn, Gilad Stern, and Alin Tomescu. Reaching consensus for asynchronous distributed key generation. *CoRR*, abs/2102.09041, 2021.
- AL07. Yonatan Aumann and Yehuda Lindell. Security against covert adversaries: Efficient protocols for realistic adversaries. In Salil P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 137–156, Amsterdam, The Netherlands, February 21–24, 2007. Springer, Heidelberg, Germany.
- AO12. Gilad Asharov and Claudio Orlandi. Calling out cheaters: Covert security with public verifiability. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 681–698, Beijing, China, December 2–6, 2012. Springer, Heidelberg, Germany.
- BBK<sup>+</sup>23. Fabrice Benhamouda, Erica Blum, Jonathan Katz, Derek Leung, Julian Loss, and Tal Rabin. Analyzing the real-world security of the algorand blockchain. Cryptology ePrint Archive, Paper 2023/1344, 2023. <https://eprint.iacr.org/2023/1344>.
- BCG93. Michael Ben-Or, Ran Canetti, and Oded Goldreich. Asynchronous secure computation. In *25th ACM STOC*, pages 52–61, San Diego, CA, USA, May 16–18, 1993. ACM Press.
- BKL19. Erica Blum, Jonathan Katz, and Julian Loss. Synchronous consensus with optimal asynchronous fallback guarantees. In Dennis Hofheinz and Alon Rosen, editors, *Theory of Cryptography - 17th International Conference, TCC 2019, Nuremberg, Germany, December 1-5, 2019, Proceedings, Part I*, volume 11891 of *Lecture Notes in Computer Science*, pages 131–150. Springer, 2019.
- BKLL20. Erica Blum, Jonathan Katz, Chen-Da Liu-Zhang, and Julian Loss. Asynchronous byzantine agreement with subquadratic communication. In Rafael Pass and Krzysztof Pietrzak, editors, *Theory of Cryptography - 18th International Conference, TCC 2020, Durham, NC, USA, November 16-19, 2020, Proceedings, Part I*, volume 12550 of *Lecture Notes in Computer Science*, pages 353–380. Springer, 2020.
- Bra87. Gabriel Bracha. Asynchronous byzantine agreement protocols. *Inf. Comput.*, 75(2):130–143, nov 1987.
- Cac21. Christian Cachin. Asymmetric distributed trust. In *ICDCN '21: International Conference on Distributed Computing and Networking, Virtual Event, Nara, Japan, January 5-8, 2021*, page 3. ACM, 2021.
- Can95. Ran Canetti. Studies in secure multiparty computation and applications. 1995.
- CHL21. Annick Chopard, Martin Hirt, and Chen-Da Liu-Zhang. On communication-efficient asynchronous MPC with adaptive security. *IACR Cryptol. ePrint Arch.*, page 1174, 2021.
- CKS00. Christian Cachin, Klaus Kursawe, and Victor Shoup. Random oracles in constantipole: practical asynchronous byzantine agreement using cryptography (extended abstract). In Gil Neiger, editor, *Proceedings of the Nineteenth Annual ACM Symposium on Principles of Distributed Computing, July 16-19, 2000, Portland, Oregon, USA*, pages 123–132. ACM, 2000.
- CKS20. Shir Cohen, Idit Keidar, and Alexander Spiegelman. Not a coincidence: Sub-quadratic asynchronous byzantine agreement WHP. In Hagit Attiya, editor, *34th International Symposium on Distributed Computing, DISC 2020, October 12-16, 2020, Virtual Conference*, volume 179 of *LIPICs*, pages 25:1–25:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- Coh16. Ran Cohen. Asynchronous secure multiparty computation in constant time. In Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang, editors, *Public-Key Cryptography - PKC 2016 - 19th IACR International Conference on Practice and Theory in Public-Key Cryptography, Taipei, Taiwan, March 6-9, 2016, Proceedings, Part II*, volume 9615 of *Lecture Notes in Computer Science*, pages 183–207. Springer, 2016.
- CP15. Ashish Choudhury and Arpita Patra. Optimally resilient asynchronous MPC with linear communication complexity. In Sajal K. Das, Dilip Krishnaswamy, Santonu Karkar, Amos Korman, Mohan J. Kumar, Marius Portmann, and Srikanth Sastry, editors, *Proceedings of the 2015 International Conference on Distributed Computing and Networking, ICDCN 2015, Goa, India, January 4-7, 2015*, pages 5:1–5:10. ACM, 2015.
- CT05. Christian Cachin and Stefano Tessaro. Asynchronous verifiable information dispersal. In Pierre Fraignaud, editor, *Distributed Computing, 19th International Conference, DISC 2005, Cracow, Poland, September 26-29, 2005, Proceedings*, volume 3724 of *Lecture Notes in Computer Science*, pages 503–504. Springer, 2005.
- CT19. Christian Cachin and Björn Tackmann. Asymmetric distributed trust. In Pascal Felber, Roy Friedman, Seth Gilbert, and Avery Miller, editors, *23rd International Conference on Principles of Distributed Systems, OPODIS 2019, December 17-19, 2019, Neuchâtel, Switzerland*, volume 153 of *LIPICs*, pages 7:1–7:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- CZ21. Christian Cachin and Luca Zanolini. Asymmetric asynchronous byzantine consensus. In Joaquín García-Alfaro, Jose Luis Muñoz-Tapia, Guillermo Navarro-Arribas, and Miguel Soriano, editors, *Data Privacy Management, Cryptocurrencies and Blockchain Technology - ESORICS 2021 International Workshops*,

- DPM 2021 and CBT 2021, Darmstadt, Germany, October 8, 2021, Revised Selected Papers*, volume 13140 of *Lecture Notes in Computer Science*, pages 192–207. Springer, 2021.
- DDFN07. Ivan Damgård, Yvo Desmedt, Matthias Fitzi, and Jesper Buus Nielsen. Secure protocols with asymmetric trust. In Kaoru Kurosawa, editor, *Advances in Cryptology - ASIACRYPT 2007, 13th International Conference on the Theory and Application of Cryptology and Information Security, Kuching, Malaysia, December 2-6, 2007, Proceedings*, volume 4833 of *Lecture Notes in Computer Science*, pages 357–375. Springer, 2007.
- DMM<sup>+</sup>20. Thomas Dinsdale-Young, Bernardo Magri, Christian Matt, Jesper Buus Nielsen, and Daniel Tschudi. Afgjort: A partially synchronous finality layer for blockchains. In Clemente Galdi and Vladimir Kolesnikov, editors, *Security and Cryptography for Networks - 12th International Conference, SCN 2020, Amalfi, Italy, September 14-16, 2020, Proceedings*, volume 12238 of *Lecture Notes in Computer Science*, pages 24–44. Springer, 2020.
- DMM<sup>+</sup>22. Bernardo David, Bernardo Magri, Christian Matt, Jesper Buus Nielsen, and Daniel Tschudi. Gearbox: Optimal-size shard committees by leveraging the safety-liveness dichotomy. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS 2022, Los Angeles, CA, USA, November 7-11, 2022*, pages 683–696. ACM, 2022.
- FHM98. Matthias Fitzi, Martin Hirt, and Ueli M. Maurer. Trading correctness for privacy in unconditional multi-party computation (extended abstract). In Hugo Krawczyk, editor, *Advances in Cryptology - CRYPTO '98, 18th Annual International Cryptology Conference, Santa Barbara, California, USA, August 23-27, 1998, Proceedings*, volume 1462 of *Lecture Notes in Computer Science*, pages 121–136. Springer, 1998.
- GHM<sup>+</sup>17. Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nikolai Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. Cryptology ePrint Archive, Paper 2017/454, 2017. <https://eprint.iacr.org/2017/454>.
- GMR88. Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, April 1988.
- GP92. Juan A. Garay and Kenneth J. Perry. A continuum of failure models for distributed computing. In Adrian Segall and Shmuel Zaks, editors, *Distributed Algorithms, 6th International Workshop, WDAG '92, Haifa, Israel, November 2-4, 1992, Proceedings*, volume 647 of *Lecture Notes in Computer Science*, pages 153–165. Springer, 1992.
- HKL20. Martin Hirt, Ard Kastrati, and Chen-Da Liu-Zhang. Brief announcement: Multi-threshold asynchronous reliable broadcast and consensus. In Hagit Attiya, editor, *34th International Symposium on Distributed Computing, DISC 2020, October 12-16, 2020, Virtual Conference*, volume 179 of *LIPICs*, pages 48:1–48:3. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- KKNS21. Idit Keidar, Eleftherios Kokoris-Kogias, Oded Naor, and Alexander Spiegelman. All you need is DAG. In Avery Miller, Keren Censor-Hillel, and Janne H. Korhonen, editors, *PODC '21: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, July 26-30, 2021*, pages 165–175. ACM, 2021.
- MMNT19. Bernardo Magri, Christian Matt, Jesper Buus Nielsen, and Daniel Tschudi. Afgjort - A semi-synchronous finality layer for blockchains. *IACR Cryptol. ePrint Arch.*, page 504, 2019.
- MP91. F.J. Meyer and D.K. Pradhan. Consensus with dual failure modes. *IEEE Transactions on Parallel and Distributed Systems*, 2(2):214–222, 1991.
- PS17. Rafael Pass and Elaine Shi. Hybrid consensus: Efficient consensus in the permissionless model. In Andréa W. Richa, editor, *31st International Symposium on Distributed Computing, DISC 2017, October 16-20, 2017, Vienna, Austria*, volume 91 of *LIPICs*, pages 39:1–39:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

## A Script for calculating minimal committee sizes

The following python script computes the minimal secure committee for various security parameters and corruption thresholds.

```
from scipy.stats import binom
import math
from fractions import Fraction
```



```

# Check if a committee of size n has an honest majority
# with probability at least  $1-2^{-\text{seccparam}}$ 
# when sampled from a ground population with at most t corruptions
def check_committee(n, seccparam, t = 1/3):
    negl_prob = Fraction(1, pow(2, seccparam))
    least_majority = math.ceil((n + 1) / 2)
    # probability that each member is honest
    p_honest = 1-t
    # cdf over honest parties having 0,1,...,least_majority-1 spots
    p_corrupt_majority = binom(n, p_honest).cdf(least_majority-1)
    return p_corrupt_majority < negl_prob

def smallest_safe_committee(seccparam, t = 1/3):
    lower_bound = 1
    while not check_committee(lower_bound * 2, seccparam, t):
        lower_bound *= 2
    upper_bound = 2 * lower_bound
    while lower_bound < upper_bound:
        test = (lower_bound + upper_bound) // 2
        if check_committee(test, seccparam, t):
            upper_bound = test
        else:
            lower_bound = test
    if upper_bound - lower_bound < 2:
        break
    # Heuristically subtract 10 from lower bound and try each possibility
    # because the predicate is not monotone.
    # The relative difference between odd and even size committees is
    # larger for smaller committees and we overshoot by < 4 for the
    # examples in main, so 10 should suffice for realistic security parameters
    for i in range(lower_bound - 10, upper_bound + 1):
        if check_committee(i, seccparam, t):
            print(
                i, "in-committee-and-up-to", t,
                "of-ground-population-corrupted-gives-honest-majority-with",
                seccparam, "bits-security")
            return i

if __name__ == '__main__':
    security_parameters = [30, 40, 60, 80]
    # Get numbers for t = 1/3 optimal resiliency
    print("Committee-size-with-optimal-resilience:")
    for seccparam in security_parameters:
        smallest_safe_committee(seccparam)
    # t = 0.25
    print("Committee-size-with-<1/4-of-GP-corrupted:")
    for seccparam in security_parameters:
        smallest_safe_committee(seccparam, 1/4)
    # algorand setting, t = 0.2
    print("Committee-size-<1/4-of-GP-corrupted-(Algorand-setting):")
    for seccparam in security_parameters:
        smallest_safe_committee(seccparam, 1/5)

```