

# Probabilistic Indistinguishability and the Quality of Validity in Byzantine Agreement

Guy Goren

Technion

Yoram Moses

Technion

Alexander Spiegelman

Novi Research

---

## Abstract

Lower bounds and impossibility results in distributed computing are both intellectually challenging and practically important. Hundreds if not thousands of proofs appear in the literature, but surprisingly, the vast majority of them apply to deterministic algorithms only. Probabilistic protocols have been around for at least four decades and are receiving a lot of attention with the emergence of blockchain systems. Nonetheless, we are aware of only a handful of randomized lower bounds.

In this paper we provide a formal framework for reasoning about randomized distributed algorithms. We generalize the notion of indistinguishability, the most useful tool in deterministic lower bounds, to apply to a probabilistic setting. We apply this framework to prove a result of independent interest. Namely, we completely characterize the quality of decisions that protocols for a randomized multi-valued Consensus problem can guarantee in an asynchronous environment with Byzantine faults. We use the new notion to prove a lower bound on the probability at which it can be guaranteed that honest parties will not decide on a possibly bogus value. Finally, we show that the bound is tight by providing a protocol that matches it.

**Keywords and phrases** Indistinguishability, probabilistic lower bounds, Byzantine agreement.

**Digital Object Identifier** 10.4230/LIPIcs...

## 1 Introduction

Randomized algorithms have a long tradition in distributed computing [42], where they have been applied to many different problems in a variety of models [34]. In the context of fault-tolerant agreement they have served to overcome the impossibility of agreement in asynchronous settings [25, 43, 9], and have significantly improved efficiency compared to deterministic solutions [22, 30]. With the recent prevalence of blockchain systems, Byzantine agreement algorithms that can overcome malicious parties have found renewed interest in both industry and academia. For obvious reasons, blockchain systems should strive to minimize the share of decisions that originate from malicious parties, and to increase the share originating from honest ones. A natural question, then, is what are the inherent limits on the quality of Byzantine agreement algorithms in this regard? Namely, what can we say about the probability with which an algorithm can guarantee that a good decision is made?

Given their practical importance, characterizing the power and limitations of randomized distributed algorithms for agreement has become ever more desirable. However, obtaining tight, or nontrivial, probabilistic bounds on properties in the asynchronous Byzantine setting can be a challenging task. As is well known, there are “*Hundreds of impossibility results for distributed computing*” [24]. But very few of them apply to randomized protocols. Unfortunately, there is currently a dearth of general tools for characterizing the properties of randomized algorithms.

The notion of indistinguishability has for years been one of the most useful tools for proving deterministic lower bounds and impossibility results in distributed computing [6]. Such deterministic lower bounds typically rely on the fact that if a correct party cannot distinguish between two executions of a deterministic protocol (i.e., its local state is the same in both), then it performs the same actions in both. In a randomized algorithm, the fact that two executions are indistinguishable to a given party up to a certain time does not ensure that its next action will be the same in both. Moreover, a single execution does not, by itself, provide information on the probability with which actions are performed. As a result, the classic notion of indistinguishability does not directly capture many of the probabilistic aspects of a randomized algorithm.

Of course, probabilistic properties of distributed algorithms such as “*the probability that the parties decide on a value proposed by an honest party is at least  $x$* ” or “*all honest parties terminate with probability 1*” cannot be evaluated based on an individual execution. Clearly, to make formal sense of such statements, we need to define an appropriate probability space. However, due to the nondeterminism inherent in our model, a probability space over the set of all executions cannot be defined (cf. [41]). This is because we can’t assume a distribution over the initial configurations, and similarly there is no well-defined distribution on the actions of the adversary, who is in charge of all the nondeterministic decisions. Once we fix the adversary’s strategy, we are left with a purely probabilistic structure, which we call an *ensemble*. An ensemble naturally induces a probability space. This allows us to formally state probabilistic properties of an algorithm  $\mathcal{A}$  of interest with respect to all of its ensembles (= adversary strategies). E.g., “*for every ensemble of algorithm  $\mathcal{A}$ , all honest parties terminate with probability 1.*”

In deterministic algorithms, indistinguishability among executions is determined based on a party  $p_i$ ’s local history, i.e., the sequence of local states that  $p_i$  passes through in the executions. We generalize the notion of an  $i$ -local history to a notion called an  *$i$ -local ensembles*. A local ensemble is a tree of local states, that captures subtle, albeit essential, aspects of probabilistic protocols. This facilitates the definition of a notion of **probabilistic indistinguishability** among ensembles, whereby two ensembles are considered indistinguishable to a process  $p_i$  if they induce identical  $i$ -local ensembles. Indistinguishability among ensembles provides a formal and convenient framework that can be used to simplify existing lower bound proofs in a probabilistic setting, and to prove new ones. A significant feature of this framework is its simplicity and ease of use, allowing similar arguments as in the deterministic case. The notions contain just enough structure beyond that of their deterministic analogues to capture the desired probabilistic properties.

Our original motivation for developing the above framework was to formally prove tight probabilistic bounds on the share of good decisions made by a randomized Byzantine agreement algorithm in an asynchronous setting. In Section 5 we use probabilistic indistinguishability to prove that, roughly speaking, no algorithm can guarantee that the probability to decide on a genuine input value is greater than  $1 - \frac{f}{n-t}$ . (As usual,  $n$  is the total number of parties here, while  $t$  and  $f$  are the maximal and actual number of failures, respectively.) Moreover, this bound is shown to be tight, by presenting an algorithm that achieves it.

This paper makes two distinct and complementary main contributions:

- We define a notion of indistinguishability that generalizes its deterministic counterpart, and is suitable for proving lower bounds in the context of probabilistic protocols. A new element in our definition is a purely probabilistic tree whose paths represent local histories of a given process. The resulting framework provides an intuitive and rigorous way to reason about probabilistic properties of such protocols.

- We introduce *Qualitative Validity*, a new probabilistic validity condition for the Byzantine agreement problem. It provides a probabilistic bound on the ability of corrupt parties to bias the decision values, which is of interest in the blockchain arena. We prove that, in a precise sense, it is the strongest achievable validity property in the asynchronous setting. Both the statement of the property and the proof are facilitated by our new framework.

## 2 Related Work

**Deterministic indistinguishability in distributed computing.** The main inherent limitation of distributed computing leveraged by most proofs is the lack of global knowledge [33, 21]. That is, each party needs to evaluate the global state of the system based on its local state and act accordingly. Deterministic indistinguishability captures exactly that: if two executions of a deterministic protocol looks the same from some party's local point of view, then this party performs exactly the same sequence of actions in both. Good surveys of techniques used in these proofs are presented in [6, 24, 33]. They all utilize indistinguishable executions, but differ in the way they construct them.

**Lower bound approaches for randomized distributed protocols.** There are several lower bound and impossibility results for randomized algorithms in the literature. One approach is to reduce the distributed problem into a cleaner mathematical one that abstracts away the issues of local knowledge, and then apply methodologies from other fields to the latter problem [3, 32, 8, 4]. An example of such a mathematical tool, used in [3, 8, 4], is a form of collective coin-flipping game [10], which is an algorithm for combining local coins into a single global one, whose bias should be small. Each of the participants flips local coins and submits them to an adversary. The adversary then gets to decide which coins to reveal and which to hide. The adversary's purpose is to bias the global coin's output, while hiding as few local flips as possible. Lower bounds on the number of coin flips required to tolerate an adversary as a function of his budget to hide coins are proven in [3, 8]. They relate these bounds to randomized distributed algorithms in the following way: Intuitively, each step of the coin-flipping game corresponds to an execution of the distributed algorithm up to some random event, which can be interpreted as the flipping of a local coin. The adversary's choice to hide or reveal this local coin corresponds to its power to kill the process that executes the random event or to let it run.

Aspnes [3] extended the valency arguments introduced in FLP [25] to the probabilistic setting and used coin flipping games to prove an  $\Omega(\frac{n^2}{\log^2(n)})$  lower bound on the expected step complexity of solving asynchronous randomized Consensus. Bar-Joseph and Ben-Or [8] applied this technique to the synchronous setting and proved an  $\Omega(\sqrt{n/\log(n)})$  lower bound on the number of rounds required for Consensus, in expectation, under a worst case adversary. Attiya and Censor-Hillel [4] closed a gap left in Aspnes's work [3] by showing that  $\Theta(n^2)$  is a tight bound on the step complexity for Consensus in the shared memory model. For the lower bound, they elegantly combined valency arguments and coin flipping games with a layering technique introduced in [38] for the deterministic case. They restricted their adversarial strategies to proceed in layers of at least  $n - t$  parties. That is, for any given configuration, parties flip local coins and then the adversary picks at least  $n - t$  parties and lets each of them perform a step, which collectively leads to the next configuration. To capture the algorithm's randomness, they compute the probability to decide  $v$  from some configuration under a given adversary by summarizing products of probability spaces induced by reachable configurations in executions in which  $v$  is decided. To define the valence of a configuration

under a set  $S$  of possible adversarial strategies, they check whether, starting from the given configuration, there are adversaries in  $S$  that lead to decisions with a probability higher than some threshold. While the elegant way in which they treat the randomness of Consensus algorithms allows them to distinguish among *bivalent*, *v-valent*, and *null-valent* configurations as required by their proof, it is not clear how their formalization can be conveniently applied to other problems.

Another approach to proving randomized distributed lower bounds is via reducing the argument to deterministic indistinguishability by considering a fixed random tape of coin flips to abstract away the randomness from an execution [5, 17]. An execution is deterministically defined by the initial configuration, scheduler, and a fixed tape of coin flips. Given two executions, the standard indistinguishability argument can be applied. Attiya and Censor-Hillel [5] proved a trade-off between termination probability and step complexity of randomized Consensus algorithms in asynchronous systems. To prove their result, they showed that for any random tape there is an indistinguishability chain that starts and ends with executions that do not allow the same decision value due to validity. Assuming a minimal probability for termination per scheduler and initial configuration, according to the tapes distribution, they show that for at least one tape all the executions in the chain terminate. By the indistinguishability chain (and the agreement condition) the decision value on the two ends of the chain is the same, which contradicts validity. Their argument can be restructured within our framework in a manner that perhaps brings their proof closer to the intuition. An ensemble gathers all possible tapes for a given adversarial strategy (i.e., initial configuration and scheduler). As a result, instead of constructing an indistinguishability chain for every random tape, we can construct a single indistinguishability chain among ensembles. Moreover, instead of explicitly referring to the termination probability induced by the tapes' distributions, we can directly use the termination probability defined on ensembles.

A recent work [1] extends a lower bound by Dolev and Reischuk [19] on the communication complexity of Byzantine agreement to the randomized case. Their proof makes use of arguments about the indistinguishability between two adversaries, without providing a formal definition. In Appendix A we show that using our definitions can fill this gap in their presentation. The work of Fich, Herlihy and Shavit in [23] proves space lower bounds for randomized shared objects. Their technique exploits covering arguments, which, in turn, rely on deterministic indistinguishability. To this end, they remove the stochastic nature of the problem by considering the “*nondeterministic solo termination*” property that requires only a non-zero probability for solo termination. They also show that their bounds immediately apply to nondeterministic wait-free objects, but conjecture that the bounds are not tight for that case. A possible reason could be that their technique ignores the quantitative probabilistic nature of the problem. A recent paper by Ellen, Gelashvili, and Zhu [20] shows that nondeterministic solo termination and obstruction freedom are equivalent in the specific context of space lower bounds.

**Indistinguishability and equivalence in probabilistic systems.** Probabilistic protocols have a long tradition in the computer science. There is a broad literature concerned with rigorous probabilistic analysis of such protocols, both in cryptography [26, 36, 28, 48, 27] and in distributed systems [45, 3, 41, 29]. Moreover, notions of indistinguishability play an important role in the analysis of probabilistic systems. In cryptography, for example, notions of computational indistinguishability and statistical indistinguishability are routinely used in order to capture the fact that protocols do not unintentionally leak information (e.g., in zero-knowledge protocols and multi-party computation) and to formalize notions such as

pseudo-randomness. Indistinguishability in this context is defined in terms of the difference between families of distributions, and in terms of an agent's ability to tell them apart. In the context of probabilistic automata there are well established notions of simulation and bisimulation that provide definitions of equivalence between two systems [46, 15]. These facilitated notions of refinement and the verification of probabilistic systems.

It may be possible to formulate probabilistic arguments used in lower bound proofs for standard distributed systems protocols either in terms of the cryptographic notions of indistinguishability or in terms of bisimulation among I/O automata. However, this would require nontrivial technical adjustments, and it is not clear that it would provide new insights or inroads into the essence of the proofs at hand. Indeed, as reviewed above, none of the lower bound proofs on probabilistic protocols in distributed computing that we are aware of make use of these frameworks.

### 3 Model

We consider a standard message passing model with a set  $\Pi$  of  $n$  parties and an adversary [7]. Parties communicate via an asynchronous network of peer-to-peer communication links.<sup>1</sup> Each party maintains a well-defined *local state* at all times. We assume for simplicity that the local state of each party  $p_i$  at a particular time  $T$  consists of an initial state  $l_i^0$  and the finite sequence of local events at  $p_i$  up to time  $T$ . This sequence is composed of the actions that  $p_i$  has performed (including the messages it has sent) before time  $T$ , as well as the messages that  $p_i$  received until time  $T$ . In particular, its *initial* local state is  $\langle l_i^0, [] \rangle$ . (For example, in a Consensus algorithm the initial local state of each party is  $\langle v_i, [] \rangle$ , where  $v_i$  is its input value.) A *configuration* is a mapping from parties to their local states and from communication links to the set of pending messages therein. (A message that has been sent on a link but not yet delivered is *pending*.) An *initial configuration* associates with each party its initial local state and each link with an empty set of pending messages. An *algorithm* defines the actions that each party performs (local computations, decisions and message sends) as a function of its local state.

We assume an interleaving model where at each point in time a single local event occurs [34]. A local event consists either of a local step performed by a party according to the algorithm, or of the delivery of a pending message. The identity of the party that moves, or the pending message that is delivered, are determined by the adversary. Both the scheduling of local steps and the delivery of messages are asynchronous. I.e., while the adversary must schedule every correct party to move infinitely often, the relative rates by which parties move can be arbitrary. Moreover, while every message sent must be eventually delivered (exactly once), there is no bound on how long messages spend in transit. An *execution* of a deterministic algorithm  $\mathcal{A}$  is a (finite or infinite) sequence of the form  $e = \langle C_0, \phi_1, C_1, \phi_2, C_2, \phi_3, \dots \rangle$ , where  $C_0$  is an initial configuration,  $C_k$  is a configuration, and  $\phi_k$  is either a local step or a message delivery for every  $k > 0$ . In case  $\phi_k$  is a local step by a correct party, the configuration  $C_k$  is obtained from  $C_{k-1}$  by modifying this party's local state (and possibly an outgoing link) according to the algorithm  $\mathcal{A}$ . If  $\phi_k$  is the delivery of a message  $m$  to party  $p_i$  from party  $p_j$ , then  $C_k$  is similar to  $C_{k-1}$  except that  $m$  is removed from the link between  $p_i$  and  $p_j$  and appended to  $p_i$ 's local state.

We also wish to model settings in which failures can occur. In this case, the identity of faulty parties and their behavior are determined by the adversary. In any given execution, we

<sup>1</sup> Our definitions can be easily translated to synchronous communication and shared memory models.

associate with the adversary a *strategy*, which determines all of its decisions in the execution. Various failure assumptions exist in the literature (e.g., crash, Byzantine, authenticated Byzantine, etc.). Each failure model induces its own set of constraints on how the adversary's strategy affects failures.

In order to facilitate the study and analysis of randomized algorithms, we slightly extend the model by adding probabilistic objects (**p\_ob** for short) to the system. We add a new type of local action (local step) that consists of a party accessing a probabilistic object. The result of this action is that the party immediately receives a return value from the object. A probabilistic object has a local state that can change following an access. The return value obtained from accessing such an object is sampled from a given distribution, which may depend on the object's local state. The range of return values may be infinite but it must be countable. A randomized algorithm is associated with the set of **p\_obs** that it employs, and the specifications of these objects are part of the algorithm's definition. Moreover, a configuration now contains the local states of the probabilistic objects, in addition to the state of the parties and the communication links.

We can, for example, model randomization via a simple local coin by having a (stateless) probabilistic object that returns 1 or 0 with probability  $1/2$  for each access. Probabilistic objects can be used to model more complex situations in which there may be correlations among values received by different parties. The reason we use probabilistic objects is to facilitate the analysis of systems with e.g., shared coins, VRFs, etc. all within the same framework [31, 3, 16, 37, 14].

## 4 Probabilistic Indistinguishability

In this section, we generalize the notion of indistinguishability to account for probabilistic aspects. First, let us review a standard definition of indistinguishability. Recall that an algorithm determines a party's behavior as a function of its local state. Therefore, two executions  $e_1$  and  $e_2$  are indistinguishable to a party  $p_i$  if the sequence of local states it goes through in both is the same. More formally, we define the *i-local history* for a party  $p_i$  in a given execution  $e$  to be the (stuttering-free) sequence of  $p_i$ 's local states in the configurations of  $e = \langle C_0, \phi_1, C_1, \phi_2, C_2, \phi_3, \dots \rangle$ . I.e., the sequence of local states except that consecutive repetitions are removed. Executions  $e_1$  and  $e_2$  are considered *indistinguishable* to  $p_i$  if both executions induce the same *i-local history*.

In the context of randomized algorithms, one is often interested in probabilistic properties of the algorithm, such as the probability that a given action is taken. Since a single execution does not, in itself, contain such probabilistic information, indistinguishability between pairs of executions is not the appropriate notion for reasoning about such properties. Such reasoning requires assigning probabilities to events consisting of appropriately chosen sets of executions. Indeed, a probabilistic notion of indistinguishability should be based on relating such sets of executions.

There is typically no way to define a probability space over the set of all executions of a randomized algorithm in an asynchronous system with failures. The initial configuration, the scheduling of parties to move and of message deliveries, and the identities and behavior of faulty parties are considered genuinely nondeterministic decisions. No probabilistic distribution is assumed on these decisions. We typically consider the nondeterministic decisions to be governed by an *adversary*. Adversaries come in different types, e.g., static or adaptive, and their abilities may vary, e.g., in terms of the type of failures that they can cause—crash vs. Byzantine, etc. The notion of indistinguishability that will be presented



shortly applies to all of them; it is independent of the type of adversary under consideration. An adversary of a given type can employ many different concrete strategies. In all cases, fixing the adversary's strategy eliminates all nondeterminism. Consequently, every transition is then either deterministic or purely probabilistic, and this induces a well-defined probability space over subsets of the executions. In order to define the probability space of interest w.r.t. a given strategy of the adversary, we proceed as follows.

## 4.1 Ensembles

We define an *ensemble* to be a directed weighted tree in which each node is a configuration, and edges represent local events. When clear from the context, we may slightly abuse notation by writing  $v \in \text{ENS}$  to denote that node  $v$  appears in the ensemble  $\text{ENS}$ . For a node  $v \in \text{ENS}$  we define  $E_{\text{out}}(v)$  to be the set of edges connecting  $v$  to its children in  $\text{ENS}$ . We define *maximal path* to be a path that cannot be extended, i.e., starts at the root and ends at a leaf. For readability, when clear from context, we refer to them simply as paths. An ensemble  $\text{ENS}$  with respect to an algorithm  $\mathcal{A}$  must satisfy the following properties:

- The root of the ensemble (tree) is an initial configuration  $C_0$ .
- Each path in the ensemble consists of the sequence of configurations of a legal execution of  $\mathcal{A}$ .
- For every node  $v \in \text{ENS}$ , the weights on the edges of  $E_{\text{out}}(v)$  are positive and their sum is 1. Moreover,
  - If  $|E_{\text{out}}(v)| = 1$ , then the edge  $(v, u) \in E_{\text{out}}(v)$  represents a deterministic local event.
  - Otherwise, (when  $|E_{\text{out}}(v)| > 1$ ), the edges in  $E_{\text{out}}(v)$  represent a single local event consisting of accessing a probabilistic object. Each edge  $(v, u) \in E_{\text{out}}(v)$  represents a possible return value, where the weight on an edge is the probability that the object access returns this value.
- The return values and the weights of edges that represent accesses to a probabilistic object  $\text{p\_ob}$  in  $\text{ENS}$  satisfy the object's specifications. For example, consider an object  $\text{p\_ob}$  consisting of a biased coin returning 1 with probability  $x \in [0, 1]$  and returning 0 with probability  $1 - x$ . If a party accesses  $\text{p\_ob}$  at a configuration  $v \in \text{ENS}$ , then the node  $v$  has two children. One child, at the end of an edge labeled  $x$  corresponds to the coin returning 1. The other child corresponds to the coin returning 0, at the end of an edge labeled  $1 - x$ .

We note that similarly to how, in deterministic algorithms, an adversary's strategy determines an execution of  $\mathcal{A}$ , in randomized algorithms the adversary's strategy determines an ensemble of  $\mathcal{A}$ . Ensembles generalize executions in the sense that a deterministic algorithm yields ensembles that consist of a single path.

### Associating probabilities with configurations

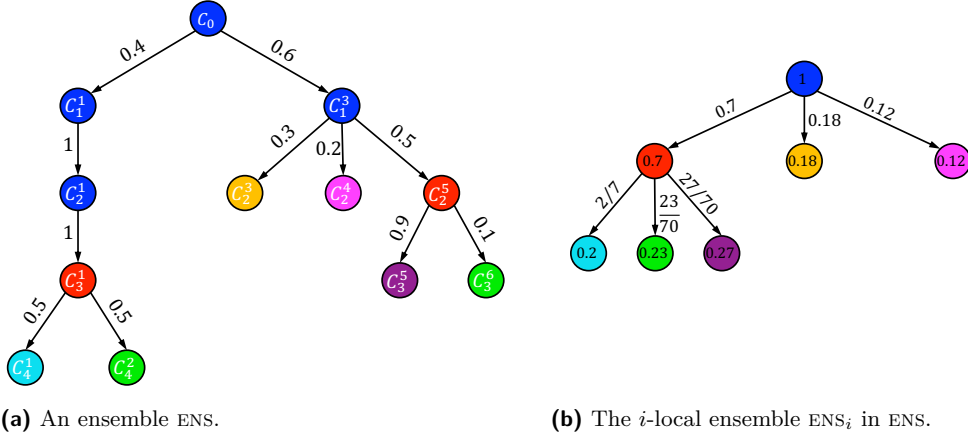
An ensemble  $\text{ENS}$  induces a probability space defined by the triplet  $(\Omega_{\text{ENS}}, \mathcal{F}, P_{\text{ENS}})$  which is specified as follows.  $\Omega_{\text{ENS}}$  is the set of paths (executions) in  $\text{ENS}$ . For each node  $v$  in the ensemble, define by  $S_v$  the set of executions (paths) in  $\text{ENS}$  that pass through  $v$ .  $\mathcal{F}$  is the sigma algebra generated by  $\{S_v : v \in \text{ENS}\}$ . (Closing under complement and countable unions.) Finally,  $P_{\text{ENS}}$  is the probability function defined by  $P_{\text{ENS}}(S_v) \triangleq$  the product of the edge weights along the path from the root to  $v$ . As required, our definitions satisfy both  $P_{\text{ENS}}(\bar{S}) = 1 - P_{\text{ENS}}(S)$  and that  $P_{\text{ENS}}$  is countably additive.

## 4.2 Local Ensembles

Our goal will be to define a notion of indistinguishability among ensembles, with respect to a particular party  $p_i$ . Roughly speaking, this is determined by  $p_i$ 's local histories in these ensembles. To this end, we consider a probabilistic tree consisting of the  $i$ -local histories in a given ensemble ENS. This is called an  $i$ -local ensemble, and is denoted by  $\text{ENS}_i$ . Since  $p_i$  can have the same local history in different paths of ENS, a node in  $\text{ENS}_i$  typically corresponds to several nodes of ENS. Consequently, the construction of  $\text{ENS}_i$  must be done carefully, to correctly account for the probabilistic transitions in  $\text{ENS}_i$ .

For each unique local state  $l_i$  of  $p_i$  that appears in ENS, there is a node  $v_{l_i} \in \text{ENS}_i$  labeled with the tuple  $\langle l_i, p_{l_i} \rangle$ , where  $p_{l_i} \in [0, 1]$  is the probability of  $p_i$  to reach  $l_i$  in ensemble ENS. More formally, for a local state  $l_i$ , define  $S_{l_i} = \{S_v \mid v \in \text{ENS} \text{ and } v \text{ is a configuration containing } l_i\}$ . The resulting probability assigned to  $v_{l_i} \in \text{ENS}_i$  labeled with  $l_i$  is  $p_{l_i} = P_{\text{ENS}}(\bigcup_{S_v \in S_{l_i}} S_v)$ . The

root of  $\text{ENS}_i$  is labeled with  $l_i^0$ , i.e.,  $p_i$ 's local state at the root of ENS, and  $p_{l_i^0} = 1$ . For each node  $v_{l_i} \in \text{ENS}_i$  labeled with an  $i$ -local state  $l_i$ , we define the set  $\text{Child}(v_{l_i})$  to consist of the nodes  $u_{\hat{l}_i} \in \text{ENS}_i$  labeled by local states  $\hat{l}_i \neq l_i$  that directly follow  $l_i$  in an execution contained in ENS. I.e., there are  $v, u \in \text{ENS}$  such that  $u$  is a child of  $v$ , and  $v, u$  represent configurations that contain  $l_i$  and  $\hat{l}_i$ , respectively. (Notice that since local states contain the full history of local events, once  $p_i$  has transitioned from a local state  $l_i$ , this state never repeats.) An illustrative example of an ensemble and its induced local ensemble for some party  $p_i$  is given in Figure 1. We remark that constructing an  $i$ -local ensemble  $\text{ENS}_i$  is linear in the size of ENS. An algorithm is provided in Appendix B for completeness.



■ **Figure 1** An ensemble ENS and its induced  $i$ -local ensemble for some party  $p_i$ . The ensemble corresponds to a given adversary strategy. It contains 6 possible executions, all start in in configuration  $C_0$ . Each node of ENS represents a unique configuration and the local state of party  $p_i$  in this configuration is color encoded in our figure. For example,  $p_i$  has the same local state, which encode in red, in  $C_3^1$  and  $C_2^5$ . The edges of ENS are labeled with their probabilities. The color of a node  $v_{l_i}$  in the  $i$ -local ensemble  $\text{ENS}_i$  represents its local state  $l_i$  and the node is labeled with the probability  $p_{l_i}$  of reaching this local state under the current strategy. For convenience, edges of  $\text{ENS}_i$  are weighted in a similar fashion to the edges of ENS.

► **Definition 1 (Probabilistic Indistinguishability).** Two ensembles are indistinguishable to a party  $p_i$  if they induce the same  $i$ -local ensemble.

For a given ensemble ENS and a local action  $\alpha$  of party  $p_i$ , we can assign a probability



for  $p_i$  to perform  $\alpha$  in ENS according to the probability space induced by ENS. A central feature of probabilistic indistinguishability is captured as follows:

► **Lemma 2.** *Let ENS-A and ENS-B be two ensembles of algorithm  $\mathcal{A}$  that are probabilistically indistinguishable to party  $p_i$ . For each action  $\alpha$  of process  $p_i$ , the probability that  $p_i$  performs  $\alpha$  is equal in ENS-A and ENS-B.*

**Proof.** Let ENS-A and ENS-B satisfy the assumptions, and let  $\alpha$  be an action of process  $p_i$ . Recall that the local state of process  $p_i$  contains the sequence of all actions that  $p_i$  has performed up to its current state. Denote by  $V_\alpha^A$  the set of nodes in ENS-A that represent configurations in which  $\alpha$  appears only once in the sequence of local events contained in  $p_i$ 's local state, and it is the last element in that sequence. As before, for a node  $v \in \text{ENS-A}$  define by  $S_v^A$  the set of executions in ENS-A that pass through  $v$ . Recall that the set  $S_v^A$  is an event in the probability space induced by ENS-A. Thus,  $S_\alpha^A \triangleq \bigcup_{v \in V_\alpha^A} S_v^A$  constitutes a measurable event

in the probability space, and  $P_A(S_\alpha^A)$  corresponds to the probability that process  $p_i$  performs action  $\alpha$  in the ensemble ENS-A. Similarly, define  $V_\alpha^B$ ,  $S_\alpha^B$ , and  $P_B(S_\alpha^B)$  with respect to ENS-B. It remains to show that  $P_A(S_\alpha^A) = P_B(S_\alpha^B)$ . Recall that for a local state  $l_i$ , we have defined the set of path sets  $S_{l_i}^{\text{ENS}} = \{S_v^{\text{ENS}} \mid v \in \text{ENS s.t. } v \text{ represents a configuration that contains } l_i\}$ .

Let  $LV_\alpha^{A_i}$  be the set of nodes in the  $i$ -local ensemble ENS-A $_i$  that are labeled with a local state  $l_i$  in which  $\alpha$  appears only once and is the last element in the sequence of local events. Similarly, define  $LV_\alpha^{B_i}$  with respect to ENS-B $_i$ .

$$P_A(S_\alpha^A) = P_A\left(\bigcup_{v \in V_\alpha^A} S_v^A\right) = P_A\left(\bigcup_{\langle l_i, p \rangle \in LV_\alpha^{A_i}} \left(\bigcup_{S_v^A \in S_{l_i}^A} S_v^A\right)\right) = \sum_{\langle l_i, p_{l_i} \rangle \in LV_\alpha^{A_i}} P_A\left(\bigcup_{S_v^A \in S_{l_i}^A} S_v^A\right).$$

Where the last equality follows from the fact that for any  $\langle l_i^1, p_{l_i^1} \rangle \neq \langle l_i^2, p_{l_i^2} \rangle \in LV_\alpha^{A_i}$  we have that  $S_{l_i^1}^A \cap S_{l_i^2}^A = \emptyset$ . This is because  $\alpha$  appears only once in the respective  $l_i$ , so a path cannot contain two different local states in which  $\alpha$  is both last and appears only once. By definition of ENS $_i$ , we have that  $p_{l_i} = P_{\text{ENS}}\left(\bigcup_{S_v^{\text{ENS}} \in S_{l_i}^{\text{ENS}}} S_v^{\text{ENS}}\right)$ . Therefore,

$$\sum_{\langle l_i, p_{l_i} \rangle \in LV_\alpha^{A_i}} P_A\left(\bigcup_{S_v^A \in S_{l_i}^A} S_v^A\right) = \sum_{\langle l_i, p_{l_i} \rangle \in LV_\alpha^{A_i}} p_{l_i}.$$

Similarly,

$$P_B(S_\alpha^B) = \sum_{\langle l_i, p_{l_i} \rangle \in LV_\alpha^{B_i}} p_{l_i}.$$

By definition of ensemble indistinguishability, ENS-A $_i = \text{ENS-B}_i$ , and so  $LV_\alpha^{A_i} = LV_\alpha^{B_i}$ . Hence,  $P_A(S_\alpha^A) = P_B(S_\alpha^B)$ , as claimed. ◀

In the next section we use ensembles and probabilistic indistinguishability to prove a lower bound on quality of decisions in Byzantine Agreement. That is, on the probability of deciding on a bogus value suggested by dishonest parties.

## 5 Byzantine Agreement with Qualitative Validity

Byzantine Agreement is one of the most fundamental problems in distributed computing. A set of  $n$  parties, some of which might be Byzantine, need to agree on the same value. Ideally,

we would like the decision to be on a value proposed by an honest party. And indeed, in the classic binary case [9], where the set of possible inputs is  $\{0, 1\}$ , this is exactly what the Validity property of Byzantine Agreement requires. However, in the multi-valued case, in which inputs come from some arbitrary domain  $\mathbb{V}$ , this is generally impossible to guarantee, because one or more Byzantine parties can propose a value that is not proposed by honest parties and otherwise act honestly [39]. Since multi-valued Byzantine Agreement protocols are the core of many Blockchain systems [12, 47], the issue of preventing malicious attacks on the “quality” of decisions is becoming more and more important. The question is, therefore, what is the best validity property a multi-valued Byzantine Agreement protocol can provide. That is, what are the conditions under which an algorithm can be guaranteed to decide on a value proposed by an honest party and what is the probability with which such a decision can be ensured if these conditions fail to hold.

Two incomparable validity definitions, called weak Validity and external Validity, have been proposed for the multi-valued case. As in the binary case, *weak Validity* [40, 18] requires that if all honest parties propose the same value  $v$ , then  $v$  is the only value that can be decided. However, if honest parties propose different values, then they can decide on some pre-defined default value (which we denote by  $\perp$ ). The initial motivation for weak validity was a spaceship cockpit with four sensors, one of which might be broken [40]. However, from a contemporary practical point of view, such a definition is useless for building Byzantine state machine replication (SMR) (e.g., as in blockchains) [2, 35] since a decision of  $\perp$  in such a setting does not allow the system to make progress. Hence, in order to guarantee progress, all honest parties must input the same value (agree *a priori*) even in failure-free runs.

To deal with this issue, Cachin, Kursawe, Petzold and Shoup [13] introduced the *external Validity* property, which allows the decision to be any value as long as it is valid according to some external predicate (e.g., a valid transaction in a blockchain system). In particular, external validity does not preclude a situation in which the decision value does not originate from an honest party. To overcome this deficiency, Abraham, Malkhi, and Spiegelman [2] extended the definition of external validity with a *decision quality* requirement, which bounds the probability of the decision being a value proposed by the Byzantine parties. Specifically, they provide an algorithm that guarantees probability of at least  $1/2$  for the decision value to be an input of an honest party. Moreover, they claim in the paper that no algorithm provide a stronger quality guarantee in the worst case scenario, but provide no proof. While their claim is very intuitive, it is not obvious how to prove it without a notion such as probabilistic indistinguishability.

Note that the two variants of multi-valued validity are incomparable. On the one hand, with external Validity parties never agree on  $\perp$  and thus SMR progress is guaranteed by reaching a meaningful decision in every slot. On the other hand, honest parties may agree on a bogus value proposed by malicious participants even if they agree *a priori*. In addition, note that neither definition takes into account the actual number of failures in the execution  $f \leq t$ . Below we define the Qualitative Validity property, which promises progress and is stronger than each of these validity conditions.

## 5.1 Problem definition

In this section we assume a computationally bounded adversary that can corrupt up to  $t$  of the  $n$  parties, where  $n = 3t + 1$ . Parties corrupted by the adversary are called *Byzantine* and may arbitrarily deviate from the protocol. Other parties are *honest*.

Given an ensemble, we denote by  $f \leq t$  the maximal number of parties the adversary corrupts in any of the paths in the ensemble. In addition, every party  $p_i$  starts with an

initial input value  $v_i$  from some domain  $\mathbb{V}$ , i.e.,  $p_i$ 's local state in the initial configuration (the root of the ensemble) is  $s_i = v_i$ . We denote by  $\mathcal{V}_{in} = \{\{v_i \mid p_i \in \Pi\}\}$  the multiset of all input values. For every multiset  $\mathcal{M}$  and value  $v \in \mathcal{M}$ , we denote by  $\text{mult}(v, \mathcal{M})$  the multiplicity of  $v$  in  $\mathcal{M}$ . The maximum multiplicity in a multiset  $\mathcal{M}$  is denoted by  $\text{max\_mult}(\mathcal{M}) \triangleq \max(\{\text{mult}(v, \mathcal{M}) \mid v \in \mathcal{M}\})$ .

We distinguish between static and adaptive adversaries and between weak and strong ones. A weak adversary does not observe the local states of honest parties, whereas a strong one does. A static adversary knows the input values but must determine the corrupted parties at the start, i.e., immediately after the root. An adaptive one is allowed to corrupt parties on the fly. To strengthen our result, we consider a weak and static adversary for the lower bound, and a strong and adaptive one for the upper bound (the algorithm).

The Agreement problem exposes an API by which a party can *propose* the input value and output a *decision* from the domain  $\mathbb{V}$ . An Agreement algorithm is one that satisfies the Agreement, Termination and Validity properties. As deterministic solutions in failure-prone asynchronous systems are impossible by FLP [25], we are interested in algorithms that never compromise safety, and ensure liveness almost surely. That is, we require that every ensemble ENS of the algorithm must satisfy the following properties:

- Agreement: In every path (i.e., execution) of ENS, all honest parties that decide, output the same decision value.
- Probabilistic Termination: Every honest party decides with probability 1.

As for validity, we extend previous definitions [40, 18, 13] to capture the optimal conditions under which parties decide on a value proposed by an honest party. Recall that input values are determined by the initial configuration, which is at the root of an ensemble. It follows that all executions in an ensemble share the same input vector. The ensemble notation allows us to require the following non-deterministic property:

- Qualitative Validity: If  $\text{max\_mult}(\mathcal{V}_{in}) - f \geq 2t + 1$ , then all honest parties that decide, output decision values in  $\mathcal{V}_{in}$ . Otherwise, the probability that they decide on a value in  $\mathcal{V}_{in}$  is at least  $1 - \frac{f}{n-t}$ .

One important feature of Qualitative Validity is that in ensembles without failures honest parties can only decide on a value in  $\mathcal{V}_{in}$ . Moreover, parties never decide on a pre-defined  $\perp$  and the probability to decide on a value in  $\mathcal{V}_{in}$  is proportional to the number of Byzantine parties that actually occur in the ensemble. Many Validity definitions are phrased with relation to the input of correct processes. We, in contrast, consider all inputs and give the adversary a choice to corrupt the parties based on their input. Our phrasing does not weaken the Validity property (a direct proof appears in ??).

## 5.2 Tight bounds on Qualitative Validity

In this section we prove that no algorithm for multi-valued Byzantine Agreement that satisfies Agreement and Probabilistic Termination can provide a better guarantee than Qualitative Validity. Moreover, we then show that this validity condition is the best we can achieve, by presenting an algorithm that satisfies this validity property.

The following lemma states that if the condition in the Qualitative Validity definition ( $\text{max\_mult}(\mathcal{V}_{in}) - f \geq 2t + 1$ ) does not hold, then we can always find a multiset  $\mathcal{M} \subset \mathcal{V}_{in}$  of size  $n - t - f$  such that no value in  $\mathcal{M}$  has multiplicity higher than  $t$ .

## XX:12 Probabilistic Indistinguishability & Qualitative Validity

► **Lemma 3.** Consider a multiset  $\mathcal{V}_{in}$  of  $n = 3t + 1$  values and  $0 < f \leq t$ .

If  $\max\_mult(\mathcal{V}_{in}) - f < 2t + 1$ , then there is a multiset  $\mathcal{M} \subset \mathcal{V}_{in}$  such that  $|\mathcal{M}| = n - t - f$  and  $\max\_mult(\mathcal{M}) \leq t$ .

**Proof.** Consider two cases:

- If  $\max\_mult(\mathcal{V}_{in}) \leq t$  then the claim's conclusion is true and the lemma holds.
- Otherwise, let  $v \in \mathcal{V}_{in}$  such that  $\text{mult}(v, \mathcal{V}_{in}) = \max\_mult(\mathcal{V}_{in}) > t$ . Since, by assumption,  $\max\_mult(\mathcal{V}_{in}) - f < 2t + 1$ , we obtain  $\text{mult}(v, \mathcal{V}_{in}) \leq 2t + f$ . This means that there are at least  $n - (2t + f) = t - f + 1$  values distinct from  $v$  in  $\mathcal{V}_{in}$ . Define  $\mathcal{M}$  to contain  $t$  copies of  $v$  along with  $t - f + 1$  values distinct from  $v$ . Then  $|\mathcal{M}| = 2t - f + 1 = n - t - f$ , and  $\max\_mult(\mathcal{M}) \leq t$  as desired, since  $f > 0$ . ◀

Clearly, if  $\max\_mult(\mathcal{V}_{in}) - f \geq 2t + 1$ , then Qualitative Validity guarantees that the decision value is in  $\mathcal{V}_{in}$ , which is the most that a validity property can ensure. Therefore, for a validity property  $\Phi$  to be stronger than Qualitative Validity, there must be  $\hat{f}$  and  $\mathcal{V}_{in}$  for which the probability to decide on a value in  $\mathcal{V}_{in}$  according to  $\Phi$  is strictly higher than  $1 - \frac{\hat{f}}{n-t}$ . However, for a validity property to be strictly stronger than Qualitative Validity it must, in addition, satisfy Qualitative Validity for all other values of  $f \leq t$  and all other  $\mathcal{V}_{in}$ . Formally, we say that a Validity property  $\Phi$  is *strictly stronger* than Qualitative Validity if an algorithm  $\mathcal{A}$  satisfying  $\Phi$  guarantees that:

1. For all  $f \leq t$  and  $\mathcal{V}_{in}$  such that  $\max\_mult(\mathcal{V}_{in}) - f \geq 2t + 1$ , in every ensemble of  $\mathcal{A}$  with  $f$  corrupted parties and the input multiset  $\mathcal{V}_{in}$  an honest party that decides, decides on a value in  $\mathcal{V}_{in}$ .
2. For all  $f \leq t$  and  $\mathcal{V}_{in}$  such that  $\max\_mult(\mathcal{V}_{in}) - f < 2t + 1$ , in every ensemble of  $\mathcal{A}$  with  $f$  corrupted parties and the input multiset  $\mathcal{V}_{in}$  the probability to decide on a value in  $\mathcal{V}_{in}$  is at least  $1 - \frac{f}{n-t}$ .
3. There exist some  $\hat{f} \leq t$  and  $\mathcal{V}_{in}$  that satisfies  $\max\_mult(\mathcal{V}_{in}) - \hat{f} < 2t + 1$ , such that in every ensemble of  $\mathcal{A}$  with  $\hat{f}$  corrupted parties and the input multiset  $\mathcal{V}_{in}$  the probability to decide on a value in  $\mathcal{V}_{in}$  is strictly higher than  $1 - \frac{\hat{f}}{n-t}$ .

Intuitively, if  $\max\_mult(\mathcal{V}_{in}) - f < 2t + 1$ , then the adversary can delay  $t$  honest parties until the decision is made such that the remaining honest parties have input values in a multiset  $\mathcal{M}$  that satisfies the property in Lemma 3. Therefore, we say that a “fair share” probability for a value  $v \in \mathcal{M}$  to be decided is proportional to  $\text{mult}(v, \mathcal{M})$  and equal to  $\frac{\text{mult}(v, \mathcal{M})}{n-t}$ . Roughly speaking, by (3), for a validity condition to be strictly stronger than Qualitative Validity, there must be a value in  $v \in \mathcal{M}$  that is decided with a probability that is higher than its “fair share”. We show that in this case there is a probabilistically indistinguishable ensemble in which  $v$  is proposed only by corrupted parties and the probability to decide on  $v$  is the same. As a result, in that ensemble the corrupted parties get more than their “fair share”, and thus violate (2). Formally, to show that Qualitative Validity is optimal we prove the following:

► **Theorem 4.** No asynchronous Byzantine Agreement algorithm satisfies a validity property  $\Phi$  that is strictly stronger than Qualitative Validity even against a weak and static adversary.

**Proof.** Assume, by way of contradiction, that there is such an algorithm  $\mathcal{A}$ . In particular, there exist  $\hat{f} \leq t$  and  $\mathcal{V}_{in}$  that satisfy  $\max\_mult(\mathcal{V}_{in}) - \hat{f} < 2t + 1$ , and in every ensemble of  $\mathcal{A}$  with  $\hat{f}$  (maximal corrupted parties in any path in the ensemble) and the input multiset  $\mathcal{V}_{in}$  the probability to decide on a value in  $\mathcal{V}_{in}$  is strictly higher than  $1 - \frac{\hat{f}}{n-t}$ . To show the

contradiction we use  $\mathcal{V}_{in}$  to construct an ensemble ENS of  $\mathcal{A}$  with a multiset  $\mathcal{V}'_{in} \neq \mathcal{V}_{in}$  of input values and  $f' \leq t$  corrupted parties such that the probability to agree on a value in  $\mathcal{V}'_{in}$  in ENS is strictly lower than  $1 - \frac{f'}{n-t}$ . This contradicts either condition (2) or (3) from the properties that  $\Phi$  must satisfy in order to be strictly stronger than Qualitative Validity. We next describe a few ensembles under different adversarial strategies and use Probabilistic Indistinguishability between them to prove the theorem.

► **Ensemble 1.**  $\mathcal{V}_{in}^1 = \mathcal{V}_{in}$  and  $f^1 = \hat{f}$ . By Lemma 3, there is a multiset  $\mathcal{M} \subset \mathcal{V}_{in}^1$  such that  $|\mathcal{M}| = n - t - \hat{f}$  and  $\max\_mult(\mathcal{M}) \leq t$ . Let  $M \subset \Pi$  be the set of parties with inputs in  $\mathcal{M}$ , i.e., the multiset  $\mathcal{M} = \{v_i \mid p_i \in M\}$ . Let  $F, T \subset \Pi$  be two sets of parties such that  $|F| = \hat{f}$ ,  $|T| = t$ , and  $\mathcal{V}_{in} \setminus \mathcal{M} = \{v_i \mid p_i \in F \cup T\}$ . The adversary (statically) corrupts the parties in  $F$  and delays all messages from parties in  $T$  until all honest parties in  $M$  decide. Messages sent among parties in  $M \cup F$  are immediately delivered. Let  $v_f \notin \mathcal{V}_{in}^1$ . The corrupted parties act like honest parties that get  $v_f$  as an input.

► **Ensemble 2.** Consider a multiset of input values  $\mathcal{V}_{in}^2$  that is identical to  $\mathcal{V}_{in}^1$  except that parties in  $F$  get  $v_f$  as an input. The adversary corrupts parties in  $T$ . Messages sent among parties in  $F \cup M$  are immediately delivered. Finally the corrupted parties send no messages.



(a) Ensemble 1. Parties in  $F$  are Byzantine that act as if they are honest with input  $v_f \notin \mathcal{V}_{in}^1$ . Messages from parties in  $T$  are delayed.

(b) Ensemble 2. Parties in  $F$  are honest and  $v_f \in \mathcal{V}_{in}^2$ . Parties in  $T$  are Byzantine that send no messages.

■ **Figure 2** Sets  $F, T$  and  $M$  divide the parties into 3 disjoint sets. In both ensembles the set  $M$  contains parties with inputs from  $\mathcal{M}$ . Ensemble 1 and Ensemble 2 are probabilistically indistinguishable to all parties in  $M$ .

Clearly, by Definition 1, Ensemble 1 and Ensemble 2 are probabilistically indistinguishable to all processes in  $M$  (see illustration in Figure 2). By the Probabilistic Termination property, all parties in  $M$  decide in Ensemble 2 with probability 1. Therefore, Lemma 2 implies that, in Ensemble 1, all parties in  $M$  decide with probability 1 as well.

For every value  $v \in \mathcal{V}_{in}$  denote by  $P^1(v)$  the probability that the honest parties in  $M$  decide on  $v$  in Ensemble 1. By assumption, the probability to decide on a value in  $\mathcal{V}_{in}$  is strictly higher than  $1 - \frac{\hat{f}}{n-t}$  in Ensemble 1, i.e.,  $\sum_{v \in \mathcal{V}_{in}} P^1(v) > 1 - \frac{\hat{f}}{n-t}$ . Therefore, there are two cases for Ensemble 1:


- First, there is a value  $u \in \mathcal{M}$  s.t.  $P^1(u) > \frac{\max\_mult(u, \mathcal{M})}{n-t}$ .
- Otherwise, there is a value  $w \notin \mathcal{M} \cup \{v_f\}$  such that  $P^1(w) > 0$ .

**First case:** For the first case, let  $U = \{p_i \in M \mid v_i = u\}$ . Since  $\max\_mult(\mathcal{M}) \leq t$ , we get  $|U| \leq t$ . Consider the following ensemble:


► **Ensemble 3.** The multiset of input values  $\mathcal{V}_{in}^3$  is identical to  $\mathcal{V}_{in}^1$  except (1) parties in  $F$  get  $v_f$  as an input; and all other parties that get  $u$  as an input in Ensemble 1 get some value  $v' \neq u$ . (Note that  $u \notin \mathcal{V}_{in}^3$ ). The adversary corrupts all parties in  $U$  s.t. corrupted parties

## XX:14 Probabilistic Indistinguishability & Qualitative Validity

act as honest parties that got  $u$  as an input.  $f^3 = |U|$ . Messages from parties in  $T$  are delayed until all parties in  $M \setminus U$  decide, and messages sent among parties in  $M \cup F$  are immediately delivered.



(a) Ensemble 1. Parties in  $F$  are Byzantine that act as if they honest with input  $v_f \notin \mathcal{V}_{in}^1$ . Messages from parties in  $T$  are delayed.




(b) Ensemble 3. Parties in  $F$  are honest and  $v_f \in \mathcal{V}_{in}^3$ . Messages from parties in  $T$  are delayed. The set  $U \subset M$  contains parties that input  $u$  in Ensemble 1. Here they are Byzantine that act as if they honest with input  $u$ . The value  $u$  is not in  $\mathcal{V}_{in}^3$ .

■ **Figure 3** In both ensembles parties in  $M \setminus U$  input the same values. Ensemble 1 and Ensemble 3 are probabilistically indistinguishable to all processes in  $M \setminus U$ .


Note that Ensemble 1 and Ensemble 3 are probabilistically indistinguishable for parties in  $M \setminus U$  (see illustration in Figure 3). Therefore, by Lemma 2, the probability of parties in  $M \setminus U$  to agree on  $u$  in Ensemble 3 is higher than  $\frac{|U|}{n-t}$ . Consequently, the probability to agree on a value in  $\mathcal{V}_{in}^3$  in Ensemble 3 is strictly lower than  $1 - \frac{|U|}{n-t} = 1 - \frac{f^3}{n-t}$ . This contradicts the assumption that  $\mathcal{A}$  satisfies a strictly stronger validity property than Qualitative Validity.

**Second case:** In this case there is a value  $w \notin \mathcal{M} \cup \{v_f\}$  such that  $P^1(w) > 0$ . Consider the following ensemble:

► **Ensemble 4.** Consider  $\mathcal{V}_{in}^4$  to be the multiset of input values that is identical to  $\mathcal{V}_{in}^1$  except parties in  $F$  get  $v_f$  as an input and all other parties that are assigned  $w$  from  $\mathcal{V}_{in}^1$  as an input, are assigned instead some  $v \neq w$  ( $w \notin \mathcal{V}_{in}^4$ ). The adversary corrupts no parties ( $f^4 = 0$ ). Messages from parties in  $T$  are delayed until all parties in  $M$  decide, and messages sent among parties in  $M \cup F$  are immediately delivered.



(a) Ensemble 1. Parties in  $F$  are Byzantine that act as if they honest with input  $v_f \notin \mathcal{V}_{in}^1$ . Messages from parties in  $T$  are delayed.



(b) Ensemble 4. All parties are honest. Messages from parties in  $T$  are delayed. The value  $v_f \in \mathcal{V}_{in}^4$ , whereas the value  $w \notin \mathcal{V}_{in}^4$ .

■ **Figure 4** Ensemble 1 and Ensemble 4 are probabilistically indistinguishable to all processes in  $M$ .

Note that Ensemble 1 and Ensemble 4 are probabilistically indistinguishable for parties in  $M$  (see illustration in Figure 4). Therefore, by Lemma 2, the probability of a party in  $M$  to agree on  $w$  in Ensemble 4 is higher than 0. Thus, since  $w \notin \mathcal{V}_{in}^4$ , the probability to agree



on a value in  $\mathcal{V}_{in}^4$  in Ensemble 4 is strictly lower than 1. Since  $f^4 = 0$  in Ensemble 4, we get that the probability to agree on a value in  $\mathcal{V}_{in}^4$  in Ensemble 4 is strictly lower than  $1 - \frac{f^4}{n-t}$ . This contradicts the assumption that  $\mathcal{A}$  satisfies a validity property  $\Phi$  that is strictly stronger than Qualitative Validity, completing the proof.  $\blacktriangleleft$

The lower bound result of Theorem 4 is tight, as the following theorem shows:

► **Theorem 5.** *There exists an asynchronous Byzantine Agreement algorithm that satisfies Qualitative Validity against a strong and adaptive adversary.*

In fact, this theorem shows a bit more than that the bound is tight. While Theorem 4 showed that no better than Qualitative Validity can be achieved even against a weak and static adversary, Theorem 5 shows that it *is* achievable, and this can be done against a much stronger adversary.

We defer the proof of Theorem 5 to Appendix C, where we show that a sequential composition of two known algorithms, from [44] and [2], yields the first Agreement protocol that satisfies the Agreement, Probabilistic Termination, and Qualitative Validity properties. In a nutshell, we first run the algorithm in [44]. By its weak validity property, if all honest parties start with the same value, then they all decide on it. Otherwise, they decide on  $\perp$ , in which case we invoke the algorithm in [2] and output its decision value. The resulting asynchronous algorithm achieves Qualitative Validity. Moreover, it does so in expected constant number of rounds, using  $O(n^2)$  communication complexity, and is resilient against  $t < n/3$  Byzantine parties. Each of these parameters is known to be optimal in this setting [25, 2, 11].

## 6 Discussion

Validity is one of the essential properties that agreement algorithms are required to satisfy. For multi-valued agreement, several distinct versions of validity have been studied in the literature over the years. Indeed, the desire to provide a variety of quality and fairness features in the blockchain world has given rise to new validity properties that are especially suited to randomized agreement algorithms. In this work we have introduced a new, probabilistic, validity property for multi-valued agreement. Called *Qualitative Validity*, this notion is strictly stronger than two popular validity conditions, which are incomparable to one another. Intuitively, it bounds the probability that corrupted parties will cause the algorithm to decide on a bogus value. Our main theorem is that, in a precise sense, Qualitative Validity is the strongest validity property that can be satisfied by an asynchronous Byzantine Agreement algorithm.

In order to prove our lower bound, we represented adversary strategies in terms of a mathematical object called an ensemble, and introduced the notion of probabilistic indistinguishability between ensembles. This framework facilitates the statement of probabilistic properties of algorithms, and the proof of lower bounds on such properties. Our framework is applicable beyond the proof of our theorem. For example, as discussed in Section 2, the lower bound proof by Attiya and Censor-Hillel in [5] constructs multiple (deterministic) indistinguishability chains to account for different random tapes. Their construction can be replaced by a single probabilistic indistinguishability chain among ensembles. Another place where our framework fits seamlessly is the probabilistic lower bound theorem of [1]. Their technical argument shows that two adversaries are indistinguishable to a particular party. While they do not provide a formal definition of indistinguishability between adversaries, interpreting their proof using our definitions fills this gap perfectly. Moreover, as we show in

Appendix A, the added transparency into their proof that is obtained by couching it using our framework allows us to strengthen their claim: Their probabilistic lower bound holds for a strictly weaker adversary than is claimed in their theorem.

Probabilistic indistinguishability captures the fact that a given party will perform the same actions, with the same probability, in both ensembles. Indeed, our theorem provides a tight bound on the probability that an honest party will decide on a good value. However, in the design and analysis of randomized distributed algorithms one may be interested in the correlation between actions of several parties. E.g., whether parties decide on the same value (i.e., satisfy Agreement) with a sufficiently high probability. The notion of probabilistic indistinguishability does not account for such correlation among actions. An interesting topic for future investigation is to formulate notions that will account for the correlation among actions of different parties. We expect that such notions will be generalizations of our definition of probabilistic indistinguishability. Their precise form is left as an open problem.

---

## References

---

- 1 Ittai Abraham, TH Hubert Chan, Danny Dolev, Kartik Nayak, Rafael Pass, Ling Ren, and Elaine Shi. Communication complexity of byzantine agreement, revisited. In *PODC*, 2019.
- 2 Ittai Abraham, Dahlia Malkhi, and Alexander Spiegelman. Asymptotically optimal validated asynchronous byzantine agreement. In *PODC*, PODC '19, New York, NY, USA, 2019. ACM.
- 3 James Aspnes. Lower bounds for distributed coin-flipping and randomized consensus. *JACM*.
- 4 Hagit Attiya and Keren Censor. Tight bounds for asynchronous randomized consensus. *JACM*.
- 5 Hagit Attiya and Keren Censor-Hillel. Lower bounds for randomized consensus under a weak adversary. *SIAM Journal on Computing*, 39(8):3885–3904, 2010.
- 6 Hagit Attiya and Faith Ellen. Impossibility results for distributed computing. *Synthesis Lectures on Distributed Computing Theory*, 5(1):1–162, 2014.
- 7 Hagit Attiya and Jennifer Welch. *Distributed computing: fundamentals, simulations, and advanced topics*, volume 19. John Wiley & Sons, 2004.
- 8 Ziv Bar-Joseph and Michael Ben-Or. A tight lower bound for randomized synchronous consensus. In *PODC*, volume 98, pages 193–199, 1998.
- 9 Michael Ben-Or. Another advantage of free choice (extended abstract) completely asynchronous agreement protocols. In *PODC*, pages 27–30, 1983.
- 10 Michael Ben-Or and Nathan Linial. Collective coin flipping, robust voting schemes and minima of banzhaf values. In *SFCS*. IEEE, 1985.
- 11 Gabriel Bracha. Asynchronous byzantine agreement protocols. *Inf. Comput.*, 75(2), 1987.
- 12 Christian Cachin et al. Architecture of the hyperledger blockchain fabric. In *DCCL*, 2016.
- 13 Christian Cachin, Klaus Kursawe, Frank Petzold, and Victor Shoup. Secure and efficient asynchronous broadcast protocols. In *Advances in Cryptology*, 2001.
- 14 Christian Cachin, Klaus Kursawe, and Victor Shoup. Random oracles in Constantinople: Practical asynchronous byzantine agreement using cryptography. *Cryptology*, 2005.
- 15 Ran Canetti, Ling Cheung, Dilsun Kirli Kaynar, Moses D. Liskov, Nancy A. Lynch, Olivier Pereira, and Roberto Segala. Time-bounded task-pioas: A framework for analyzing security protocols. In *DISC*. Springer, 2006.
- 16 Ran Canetti and Tal Rabin. Fast asynchronous byzantine agreement with optimal resilience. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, 1993.
- 17 Benny Chor, Michael Merritt, and David B Shmoys. Simple constant-time consensus protocols in realistic failure models. *Journal of the ACM (JACM)*, 36(3):591–614, 1989.
- 18 Miguel Correia, Nuno Ferreira Neves, and Paulo Veríssimo. From consensus to atomic broadcast: Time-free byzantine-resistant protocols without signatures. *Computer*, 2006.
- 19 Danny Dolev and Rudiger Reischuk. Bounds on information exchange for byzantine agreement. *JACM*, 1985.

- 20 Faith Ellen, Rati Gelashvili, and Leqi Zhu. Revisionist simulations: a new approach to proving space lower bounds. In *PODC*, 2018.
- 21 Ronald Fagin, Yoram Moses, Joseph Y Halpern, and Moshe Y Vardi. *Reasoning about knowledge*. MIT press, 2003.
- 22 Peasech Feldman and Silvio Micali. An optimal probabilistic protocol for synchronous byzantine agreement. *SIAM Journal on Computing*, 26(4):873–933, 1997.
- 23 Faith Fich, Maurice Herlihy, and Nir Shavit. On the space complexity of randomized synchronization. *Journal of the ACM (JACM)*, 45(5):843–862, 1998.
- 24 Faith Fich and Eric Ruppert. Hundreds of impossibility results for distributed computing. *Distributed computing*, 16(2-3):121–163, 2003.
- 25 Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *JACM*, 1985.
- 26 Oded Goldreich. *Foundations of cryptography: volume 2, basic applications*. 2009.
- 27 Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In *Theory of Computing*, 1987.
- 28 Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on computing*, 18(1):186–208, 1989.
- 29 Sergiu Hart, Micha Sharir, and Amir Pnueli. Termination of probabilistic concurrent program. *ACM Trans. Program. Lang. Syst.*, 5(3):356–380, 1983. doi:10.1145/2166.357214.
- 30 Jonathan Katz and Chiu-Yuen Koo. On expected constant-round protocols for byzantine agreement. In *Annual International Cryptology Conference*, pages 445–462. Springer, 2006.
- 31 Valerie King and Jared Saia. Byzantine agreement in expected polynomial time. *JACM*, 2016.
- 32 Eyal Kushilevitz, Yishay Mansour, Michael O Rabin, and David Zuckerman. Lower bounds for randomized mutual exclusion. In *Theory of computing*, 1993.
- 33 Nancy Lynch. A hundred impossibility proofs for distributed computing. In *PODC*, 1989.
- 34 Nancy A Lynch. *Distributed algorithms*. Elsevier, 1996.
- 35 Dahlia Malkhi. *Concurrency: The Works of Leslie Lamport*. Morgan & Claypool, 2019.
- 36 Ueli M. Maurer. Indistinguishability of random systems. In *EUROCRYPT*, 2002.
- 37 Silvio Micali, Michael Rabin, and Salil Vadhan. Verifiable random functions. In *STOC*, 1999.
- 38 Yoram Moses and Sergio Rajsbaum. The unified structure of consensus: a layered analysis approach. In *PODC*, 1998.
- 39 Gil Neiger. Distributed consensus revisited. *Information processing letters*, 1994.
- 40 Marshall Pease, Robert Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *Journal of the ACM (JACM)*, 27(2):228–234, 1980.
- 41 Amir Pnueli. On the extremely fair treatment of probabilistic algorithms. In *Theory of computing*, 1983.
- 42 Michael O Rabin. Probabilistic algorithms in finite fields. *Journal on computing*, 1980.
- 43 Michael O Rabin. Randomized byzantine generals. In *SFCS*, 1983.
- 44 Michel Raynal and Achour Mostéfaoui. Signature-free asynchronous byzantine systems: from multivalued to binary consensus with  $t < n/3$ ,  $o(n^2)$  messages, and constant time. *Acta Informatica*, 54(5), 2017.
- 45 Roberto Segala. *Modeling and verification of randomized distributed real-time systems*. PhD thesis, Cambridge, 1995.
- 46 Roberto Segala and Nancy A. Lynch. Probabilistic simulations for probabilistic processes. *Nord. J. Comput.*, 2(2):250–273, 1995.
- 47 VMware. Concord: An open source decentralized trust infrastructure, 2018.
- 48 Andrew Chi-Chih Yao. How to generate and exchange secrets. In *SFCS*, 1986.

## A Proving Theorem 4 from [1] (ACDNPLS)

In (binary) *Byzantine Broadcast*, an *a priori* fixed designated *sender* starts out with an input bit  $b \in \{0, 1\}$ . An algorithm  $\mathcal{A}$  solves binary Byzantine Broadcast with probability at least  $x$  if, in every ensemble of  $\mathcal{A}$ , with the probability of at least  $x$  all of the three following properties hold.

- Agreement. All honest parties that output a bit, output the same bit.
- Termination. Every honest party outputs a bit.
- Validity. If the sender is honest and the sender's input is  $b$ , then all honest parties output  $b$ .

The following theorem from [1] considers a model with a non-uniform p.p.t. strongly adaptive adversary. That is, the adversary can (1) leverage randomness in its favor, and (2) observe that a message is sent at time  $T$  by any party  $p_i$  to any other party, decide to adaptively corrupt  $p_i$ , and remove the messages sent by  $p_i$  at time  $T$ .

► **Theorem 6** (ACDNPLS [1]). *If a protocol solves Byzantine Broadcast with  $\frac{3}{4} + \epsilon$  probability against a non-uniform p.p.t. strongly adaptive adversary, then in expectation, honest nodes collectively need to send at least  $(\epsilon t)^2$  messages.*

Using our definitions from this paper, we can now better formalize their statement (and slightly strengthen it). Our proof follows the outline of the proof in [1].

**Theorem 2'.** *If an algorithm  $\mathcal{A}$  solves (in a model with a strongly adaptive polynomial time adversary) Byzantine Broadcast with probability at least  $\frac{3}{4} + \epsilon$ , then there exists an ensemble of  $\mathcal{A}$  in which honest parties collectively send at least  $(\epsilon t)^2$  messages in expectation.*

**Proof.** Let  $\mathcal{A}$  be an algorithm that solves Byzantine Broadcast with probability at least  $\frac{3}{4} + \epsilon$ . Assume by way of contradiction that in **every ensemble** of  $\mathcal{A}$ , the honest parties collectively send fewer than  $(\epsilon t)^2$  messages in expectation.

Without loss of generality, assume that there exist  $\lceil n/2 \rceil$  parties each of which outputs 0 with probability at most  $1/2$  if they receive no messages. (Otherwise, then there must exist  $\lceil n/2 \rceil$  nodes that output 1 with at most  $1/2$  probability if they receive no messages, and the entire proof follows from a symmetric argument.) Formally, the set  $N_0 \subset \Pi$  is of size  $\lceil n/2 \rceil$ , and for every  $p_i \in N_0$  it holds that in every ensemble  $\text{ENS}$  of  $\mathcal{A}$ ,  $P_{\text{ENS}}[p_i \text{ outputs } 0 \mid p_i \text{ receives no messages in the path}] \leq 1/2$ . Let  $F \subset N_0$  be a set of  $t/2$  these parties, not containing the designated sender. Note that these nodes may output 1 or they may simply not terminate if they receive no messages. (We can always find such an  $F$  because  $t/2 < \lceil n/2 \rceil$ .) Consider the following ensemble  $\text{ENS-A}$ .

**Ensemble A.** *The sender's input bit is 0, all messages that are sent, are synchronously delivered, and the parties in  $F$  are corrupted. Specifically, a party in  $F$  behaves honestly (according to its protocol) except for (a) not sending messages to any other party in  $F$ , and (b) ignoring the first  $t/2$  messages it receives from parties in  $\Pi \setminus F$ .*

Let  $z_A$  be a random variable, in  $\text{ENS-A}$ , denoting the number of messages sent by parties in  $\Pi \setminus F$  to  $F$ . By the assumption we have that  $\mathbb{E}[z_A] < (\epsilon t)^2$ . Let  $X_1$  be the event that  $z_A \leq \frac{(\epsilon t)^2}{2\epsilon}$ . By Markov's inequality,  $P_{\text{ENS-A}}[z_A > \frac{1}{2\epsilon} \mathbb{E}[z_A]] < 2\epsilon$ . Thus, the probability of the event  $X_1$  in  $\text{ENS-A}$  is  $P_{\text{ENS-A}}[z_A \leq \frac{(\epsilon t)^2}{2}] \geq P_{\text{ENS-A}}[z_A \leq \frac{1}{2\epsilon} \mathbb{E}[z_A]] > 1 - 2\epsilon$ .

Let  $p_f$  be the party in  $F$  with the highest probability of receiving at most  $t/2$  from the first  $\frac{\epsilon t^2}{2}$  messages sent by honest parties in ENS-A. Formally, denote by  $x_i$  the random variable, in ENS-A, corresponding to the number of messages received by  $p_i$  out of the first  $\frac{\epsilon t^2}{2}$  messages sent by honest parties. Then  $p_f \triangleq \arg \max_{p_i \in F} P_{\text{ENS-A}}[x_i \leq t/2]$ . Notice that in each path in ENS-A there are a total of at most  $\frac{\epsilon t^2}{2}$  messages to distribute among the  $t/2$  parties in  $F$ . Therefore, fewer than  $\epsilon t$  parties in  $F$  receive more than  $t/2$  messages, and at least  $|F| - \epsilon t = (1/2 - \epsilon)t$  parties receive at most  $t/2$  of the mentioned messages. This implies that the expected probability of a party in  $F$  to receive at most  $t/2$  of the messages is at least  $\frac{|F| - \epsilon t}{|F|} = 1 - 2\epsilon$ . Let  $X_{p_f}$  be the event in ENS-A that  $x_f \leq t/2$ . By choosing  $p_f$  in the way we did, we obtain that  $P_{\text{ENS-A}}[X_{p_f}] \geq 1 - 2\epsilon$ .

In ENS-A, the probability that at most  $\frac{(\epsilon t)^2}{2\epsilon}$  messages from honest parties are sent to  $F$ , whilst party  $p_f$  receives at most  $t/2$  of those messages is:

$$P_{\text{ENS-A}}[X_1 \cap X_{p_f}] = P_{\text{ENS-A}}[X_1] + P_{\text{ENS-A}}[X_{p_f}] - P_{\text{ENS-A}}[X_1 \cup X_{p_f}] > (1 - 2\epsilon) + (1 - 2\epsilon) - 1 = 1 - 4\epsilon.$$

Now consider the following ensemble ENS-B, which is very similar to ENS-A.

**Ensemble B.** *The sender's input bit is 0, and parties in  $F \setminus \{p_f\}$  are corrupted and behave exactly as in ENS-A. In addition, parties in  $\Pi \setminus F$  behave as in ENS-A (according to the algorithm  $\mathcal{A}$ ) except that the first  $t/2$  messages that are supposed to be sent by  $\Pi \setminus F$  to  $p_f$  are now omitted from ENS-B. In order to do so, at most  $t/2$  parties in  $\Pi \setminus F$  are also corrupted (in an adaptive manner) in ENS-B. Other than this, the corrupted parties in  $\Pi \setminus F$  behave exactly as in ENS-A (including sending later messages to  $p_f$ ).*

We observe that for a party  $p_i \in \Pi \setminus F$  that is honest in ENS-B, the ensembles ENS-A and ENS-B are probabilistically indistinguishable. Therefore, by Lemma 2, their protocols prescribe the same (possibly probabilistic) behavior in both of the ensembles. The fact that all honest parties act according to  $\mathcal{A}$ , together with the fact that the maximal number of corrupted parties in ENS-B at most  $|F| - 1 + t/2 = t - 1 \leq t$ , mean that ENS-B is an ensemble of  $\mathcal{A}$ .

By construction,  $p_f$  receives no messages in ENS-B with the same probability as  $P_{\text{ENS-A}}[X_1 \cap X_{p_f}] = 1 - 4\epsilon$ . Recall that (by the definition of  $F$ ), party  $p_f$  outputs 0 with probability at most  $1/2$  on the set of paths in which it receives no messages whatsoever. Let  $Y_f$  be the complementary event, in ENS-B, to  $p_f$  outputting 0. That is,  $Y_f$  is the event that  $p_f$  outputs 1 or does not terminate at all. Then  $P_{\text{ENS-B}}[Y_f] > \frac{1}{2}(1 - 4\epsilon)$ .

Moreover, since ENS-A and ENS-B are probabilistically indistinguishable to honest parties in  $\Pi \setminus F$ , by Lemma 2, the probability of these parties to output 0 is equal in ENS-A and ENS-B, and is at least  $\frac{3}{4} + \epsilon$  by assumption. Let  $Y_0$  be the event in ENS-B that all honest parties in  $\Pi \setminus F$  output 0. If both  $Y_0$  and  $Y_f$  occur, then this path (execution) in ENS-B violates either Agreement of Termination. We get

$$P_{\text{ENS-B}}[Y_0 \cap Y_f] = P_{\text{ENS-B}}[Y_0] + P_{\text{ENS-B}}[Y_f] - P_{\text{ENS-B}}[Y_0 \cup Y_f] > (\frac{3}{4} + \epsilon) + \frac{1}{2}(1 - 4\epsilon) - 1 = \frac{1}{4} - \epsilon.$$

This contradicts the assumption that  $\mathcal{A}$  solves Byzantine Broadcast with probability at least  $\frac{3}{4} + \epsilon$ .  $\blacktriangleleft$

This proof followed the outline of the original one and made use of the same underlying structure. But the use of ensembles and probabilistic indistinguishability provided the

arguments with a rigorous foundation and removed the ambiguity in their use of indistinguishability. The improved transparency also made it obvious that enabling the adversary to employ randomization is unnecessary. Hence, the result could have been slightly strengthened by weakening the adversary's abilities.

## **B Pseudo-code for Construction of a local ensemble**

We construct  $\text{ENS}_i$  in two reiterating steps: (1) We connect a node directly to its closest descendants in which  $i$ 's local state does not change (and remove all nodes in between). (2) For each node – starting from the root and going down the tree – we combine sons that represent the same local state by merging their subtrees and assigning its root (the combined node) the sum of the probabilities. From a complexity perspective this equals to two simultaneous BFS runs on the (possibly infinite)  $\text{ENS}$  tree.

## **C Consensus with Qualitative Validity is Solvable**

In this section we show that a sequential composition of two known algorithms, from [44] and [2], yields an Agreement protocol that satisfies the Agreement, Probabilistic Termination, and Qualitative Validity properties. We next overview the properties guaranteed by each of the algorithms, then show how to combine them to achieve Qualitative Validity, and finally prove correctness and analyse complexity.

### **Overview of the RM protocol's properties.**

The asynchronous Agreement protocol proposed by Raynal and Mostefaoui (RM) [44] satisfies Agreement, Probabilistic termination, Weak validity, and Non-intrusion against an adaptive adversary. The Weak validity property is a variant of the first part of the Qualitative Validity. That is, if all input values are the same ( $\max\_mult(\mathcal{V}_{in}) = 3t+1$ ), then honest parties can only decide on this value. However, if parties start with different values ( $\max\_mult(\mathcal{V}_{in}) < 3t+1$ ), then they are allowed to agree on a pre-defined  $\perp$  value. The Non-intrusion property requires that honest parties decides on values in  $\mathcal{V}_{in} \cup \{\perp\}$ . That is, honest parties never decide on a value promoted by the adversary.

The complexity of RM [44] is the following: the protocol (1) tolerates up to  $t < n/3$  Byzantine parties, (2) runs in expected constant number of rounds, and (3) sends  $O(n^2)$  words in  $O(n^2)$  messages where a word contains a constant number of signatures and values.

### **Overview of the AMS protocol's properties.**

The asynchronous agreement protocol proposed by Avraham, Malkhi, and Spiegelman (AMS) [2] satisfies the Agreement, Probabilistic termination, External validity, and Quality properties against an adaptive adversary. The External validity property requires that honest parties decide on values that are valid by some external predicate. The Quality property requires that the probability that the decision value is in  $\mathcal{V}_{in}$  is at least  $1/2$ . The complexity of AMS is similar to that of RM.

### **Sequential composition**

We show that a sequential composition  $\mathcal{A}^s$  of the RM and AMS protocols satisfies Agreement, Probabilistic termination, and Qualitative Validity with an optimal resilience and complexity in an asynchronous setting with an adaptive adversary.



■ **Algorithm 1** Building local ensemble from ensemble ENS for party  $p_i$

*NodesQueue* - An initially empty queue.  
*root* is the root of ENS

---

```

1: create node  $v_{root_i}$ 
2:  $l_i^{root} \leftarrow root.configuration[i]$ 
3:  $Start_{root_i} \leftarrow \{\langle root, 1 \rangle\}$ 
4:  $NodesQueue.enqueue(\langle v_{root_i}, l_i^{root}, Start_{root_i} \rangle)$ 
5: while NodesQueue not empty do
6:    $\langle v_{l_i}, l_i, Start_{l_i} \rangle \leftarrow NodesQueue.dequeue()$ 
7:    $LabelNode(v_{l_i}, l_i, Start_{l_i})$ 

8: procedure LABELNODE( $v_{l_i}, l_i, Start_{l_i}$ )
9:    $SetOfChildren \leftarrow \{\}$ 
10:   $p_{l_i} \leftarrow \sum_{v \in Start_{l_i}} p_v$ 
11:  for each  $v \in Start_{l_i}$  do
12:    for each  $u \in Child(v)$  do
13:       $propagationQueue.enqueue(\langle u, p_v \cdot weight(v, u) \rangle)$ 
14:  while propagationQueue not empty do
15:     $\langle u, p_u \rangle \leftarrow propagationQueue.dequeue()$ 
16:    if  $u.configuration[i] = l_i$  then
17:      for each  $w \in Child(u)$  do
18:         $propagationQueue.enqueue(\langle w, p_u \cdot weight(u, w) \rangle)$ 
19:    else  $\triangleright u.configuration[i] \neq l_i$ 
20:       $SetOfChildren \leftarrow SetOfChildren \cup \{u.configuration[i]\}$ 
21:      if the set  $Start_{u.configuration[i]}$  does not exist then
22:         $Start_{u.configuration[i]} \leftarrow \{\}$ 
23:         $Start_{u.configuration[i]} \leftarrow Start_{u.configuration[i]} \cup \{\langle u, p_u \rangle\}$ 
24:    label node  $v_{l_i}$  with  $\langle l_i, p_{l_i} \rangle$ 
25:    for each  $l_i^j \in SetOfChildren$  do
26:      create node  $v_{l_i^j}$  and add it to  $Child(v_{l_i})$ 
27:       $NodesQueue.enqueue(\langle v_{l_i^j}, l_i^j, Start_{l_i^j} \rangle)$ 
28: end procedure

```

---

The pseudocode of  $\mathcal{A}^s$  appears as Algorithm 2. First, parties try to reach agreement via the RM protocol and if its decision value is not  $\perp$ , then the parties decide on this value. Otherwise, they propose their input value in the AMS protocol and decide on its decision value. Although the Weak validity property that is proved for RM in [44] does not imply the first part of Qualitative Validity, a small modification of the proof proves that the protocol indeed satisfies it. Moreover, the non-intrusion property guarantees that even if  $\max\_mult(\mathcal{V}_{in}) - f < 2t + 1$  then, in ensembles of  $\mathcal{A}^s$  with  $f$  maximal corrupted parties and an input multiset  $\mathcal{V}_{in}$ , parties never decide in line 3 of Algorithm 2 on a value that is not in  $\mathcal{V}_{in}$ . To prove that the protocol in Algorithm 2 satisfies qualitative validity, we need to show that the AMS protocol satisfies the second part of Qualitative Validity. That is, in every ensemble of  $\mathcal{A}^s$  with  $f$  maximal corrupted parties and an input multiset  $\mathcal{V}_{in}$ , the probability to decide on a value in  $\mathcal{V}_{in}$  is at least  $1 - \frac{f}{n-t}$ . By the quality property of AMS we get a probability of at least  $1/2$  in the worst case when  $f = t$ . Below we overview the main part of the AMS algorithm and prove that the protocol indeed satisfies a stronger property, i.e., the second part of Qualitative Validity.

■ **Algorithm 2** An agreement algorithm with Qualitative validity: protocol for party  $p_i$

---

```

1:  $y_i \leftarrow RM\text{-propose}(v_i)$ 
2: if  $y_i \neq \perp$  then
3:   decide  $y_i$ 
4: else
5:   decide  $AMS\text{-propose}(v_i)$ 

```

---

## C.1 Analysis

The Agreement and Probabilistic-termination properties of the composition follows immediately from the ones in RM and AMS since  $y_i \neq \perp$  in line 2 is true or false for all correct processes categorically. As for Qualitative Validity, it follows from the weak validity proof in RM that if  $\max\_mult(\mathcal{V}_{in}) - f \geq 2t + 1$ , then all honest parties that decide, decide on values in  $\mathcal{V}_{in}$ . Moreover, by the non-intrusion property of RM, honest parties can only decide in line 3 on a value in  $\mathcal{V}_{in}$ . Thus, to prove Qualitative Validity we need to show that AMS satisfies that the probability to decide on a value in  $\mathcal{V}_{in}$  is at least  $1 - \frac{f}{n-t}$ . Below we overview the relevant parts of AMS and then prove that it indeed satisfies the required property.

In a nutshell, in every round of the AMS protocol, parties concurrently promote their values via some broadcast algorithm until at least  $n - t$  broadcasts complete. Then, using a shared global coin, parties elect one broadcast instance uniformly at random and ignore the rest. If a completed broadcast is elected, then its value is *fixed* and parties will eventually decide on this value. Otherwise, parties continue to the next round with either their value from the previous round or with the value of the elected broadcast in this round. Given the above description we prove the following lemma:

► **Lemma 7.** *The probability to decide on a value proposed by an honest party in the AMS protocol is at least  $1 - \frac{f}{n-t}$ .*

**Proof.** We prove the lemma by showing that the probability to decide on a value proposed only by Byzantine parties is at most  $\frac{f}{n-t}$ . To bound this probability from above, we assume that if a broadcast by a Byzantine sender is elected in some round  $r$ , then all honest parties decide on its value in round  $r$ . That is, we assume that all Byzantine parties complete their broadcasts in all rounds before the parties randomly elect one broadcast instance. By the

AMS protocol, for every round  $r$ , if a not completed broadcast is elected, then honest parties continue to the next round with either their values from the previous round or with the value of the elected broadcast. Therefore, it follows by induction that in all rounds until a broadcast with a Byzantine sender is elected honest parties broadcast values proposed by honest parties.

Denote by  $x_1, x_2, x_3, \dots$  the number of completed broadcasts in rounds  $r_1, r_2, r_3, \dots$ , respectively. Thus, for every round  $r_i$ ,  $\frac{x_i}{n}$  is the probability to elect a completed broadcast and fix its decision value, whereas  $\frac{n-x_i}{n}$  is the probability to elect an uncompleted broadcast and continue to the next round without fixing a decision. Let  $Q$  be the probability to decide on a value proposed only by Byzantine parties. We get that:

$$Q = \frac{x_1}{n} \frac{f}{x_1} + \frac{n-x_1}{n} \left( \frac{x_2}{n} \frac{f}{x_2} + \frac{n-x_2}{n} \left( \frac{x_3}{n} \frac{f}{x_3} + \frac{n-x_3}{n} \left( \frac{x_4}{n} \frac{f}{x_4} + \frac{n-x_4}{n} (\dots) \right) \right) \right)$$

Let  $x = \min(\{x_i \mid i \in \mathbb{N}\})$ , we get that

$$\begin{aligned} Q &\leq \frac{x}{n} \frac{f}{x} + \frac{n-x}{n} \left( \frac{x}{n} \frac{f}{x} + \frac{n-x}{n} \left( \frac{x}{n} \frac{f}{x} + \frac{n-x}{n} \left( \frac{x}{n} \frac{f}{x} + \frac{n-x}{n} (\dots) \right) \right) \right) = \\ &\frac{f}{n} \left( 1 + \frac{n-x}{n} + \left( \frac{n-x}{n} \right)^2 \dots \right) = \frac{f}{n} \left( \prod_{i=0}^{\infty} \left( \frac{n-x}{n} \right)^i \right) = \frac{f}{n} \frac{n}{n - (n-x)} = \frac{f}{x} \end{aligned}$$

By the protocol,  $x_i \geq n - t$  for every round  $r_i$ . Therefore,  $Q \leq \frac{f}{n-t}$ .

◀

### Complexity.

Since the sequential composition uses one instance of RM and one of AMS, its asymptotic complexity and resilience are equal to that of RM and AMS, which are proven to be optimal [25, 2, 11]. That is, it (1) tolerates up to  $t < n/3$  Byzantine parties, (2) runs in expected constant number of rounds, and (3) sends  $O(n^2)$  words in  $O(n^2)$  messages where a word contains a constant number of signatures and values.

## D Qualitative Validity vs Weak Validity

A hasty reader might believe “Weak Validity (plus agreement) is actually stronger than the first statement of the Qualitative Validity” which can lead him to the false conclusion that “Qualitative Validity is incomparable to Weak Validity.”

Although it might not be obvious in a first glance, Qualitative Validity is strictly stronger than Weak Validity. Namely, (1) every algorithm that satisfies Qualitative Validity (QV) necessarily satisfies Weak Validity, while (2) there are algorithms that satisfy weak validity and do not satisfy QV. Part (2) follows immediately from the second part of the QV definition. Part (1) follows from Theorem 4. In particular, Theorem 4 shows that if the first part of QV is not satisfied and  $f > 0$ , then there is always a positive probability to decide on a value not in  $V_{in}$ . For better transparency, we provide here a direct proof of the following:

▷ **Claim 8.** A protocol that satisfies Qualitative Validity, also satisfies Weak validity.

**Proof.** Let  $\mathcal{A}$  be a protocol that satisfies QV, and let ENS be any ensemble of  $\mathcal{A}$  in which  $2t+1$  honest parties start with  $v$ . We show that  $v$  is the only possible decision in ENS. Consider another ensemble ENS' which is the same as ENS except that all corrupted parties (the same as parties in ENS) have  $v$  as their input but act exactly as in ensemble ENS. In ensemble ENS'

## XX:24 Probabilistic Indistinguishability & Qualitative Validity

we have  $\max\_mult(\mathcal{V}_{in}) \geq 2t + 1 + \mathbf{f}$  and therefore  $\max\_mult(Vin) - \mathbf{f} \geq 2t + 1$ . By the first condition of QV, we get that  $v$  is the only possible decision in  $\text{ENS}'$ . Clearly,  $\text{ENS}$  and  $\text{ENS}'$  are probabilistically indistinguishable for all honest parties, hence, the only possible decision in  $\text{ENS}$  is  $v$ . ◀