Brief Announcement: Probabilistic Indistinguishability and The Quality of Validity in Byzantine Agreement

Guy Goren
The Viterbi Faculty of Electrical & Computer Engineering, Technion, Haifa, Israel

Yoram Moses
The Viterbi Faculty of Electrical & Computer Engineering, Technion, Haifa, Israel

Alexander Spiegelman
Novi Research, USA

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 work 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 guaranteed probability that honest parties will not decide on a possibly bogus value proposed by a malicious party. Finally, we show that the bound is tight by providing a protocol that matches it.

This brief announcement consists of an introduction to the full paper [6] by the same title. The interested reader is advised to consult the full paper for a detailed exposition.

2012 ACM Subject Classification Theory of computation → Distributed algorithms; Security and privacy → Distributed systems security

Keywords and phrases Indistinguishability, probabilistic lower bounds, Byzantine agreement

Digital Object Identifier 10.4230/LIPIcs.DISC.2021.57

Related Version *Full Version*: https://arxiv.org/abs/2011.04719 [6]

Funding *Guy Goren*: Supported by a grant from the Technion Hiroshi Fujiwara cyber security research center and the Israel cyber bureau.
Yoram Moses: This work was funded in part by ISF grant 2061/19. Yoram Moses is the Israel Pollak academic chair at the Technion.

1 Introduction

Randomized algorithms have a long tradition in distributed computing [10], where they have been applied to many different problems in a variety of models [8]. In the context of fault-tolerant agreement they have served to overcome the impossibility of agreement in asynchronous settings [5, 11, 2], and have significantly improved efficiency compared to deterministic solutions [3, 7]. 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*" [4]. 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 [1]. 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 the nondeterminism inherent in our model, a probability space over the set of all executions cannot be defined (cf. [9]). 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 ensemble*. 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 of [6] 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.

The 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.

We now provide an informal description of the Qualitative Validity condition, and present the lower bound and upper bound results regarding this condition that appear in the full version [6]. Roughly speaking, we use $\boldsymbol{f}$ to denote the maximal number of failures that a given adversary's strategy allows to occur in any of its possible executions. In addition, $max\_mult(\mathcal{V}_{in})$ denotes the multiplicity of the most frequent value in the input vector (consisting of the initial values of the n parties).

▶ **Definition 1** (Qualitative Validity). *If* $max\_mult(\mathcal{V}_{in}) - \boldsymbol{f} \geq 2t + 1$*, 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* $1 - \frac{\boldsymbol{f}}{n-t}$*.*

Our main lower bound, which we prove using the probabilistic indistinguishability formalism introduced in this work, is stated as follows:

▶ **Theorem 2.** *No asynchronous Byzantine Agreement algorithm satisfies a validity property* Φ *that is strictly stronger than Qualitative Validity even against a weak and static adversary.*

In order to show that the lower bound result of Theorem 2 is tight, we present a Byzantine agreement protocol that satisfies Qualitative Validity. As a result we obtain:

▶ **Theorem 3.** *There exists an asynchronous Byzantine agreement algorithm that satisfies Qualitative Validity against a strong and adaptive adversary.*

References

1. Hagit Attiya and Faith Ellen. Impossibility results for distributed computing. *Synthesis Lectures on Distributed Computing Theory*, 5(1):1–162, 2014.
2. Michael Ben-Or. Another advantage of free choice (extended abstract) completely asynchronous agreement protocols. In *Proceedings of the second annual ACM symposium on Principles of distributed computing*, pages 27–30, 1983.
3. Pesech Feldman and Silvio Micali. An optimal probabilistic protocol for synchronous byzantine agreement. *SIAM Journal on Computing*, 26(4):873–933, 1997.
4. Faith Fich and Eric Ruppert. Hundreds of impossibility results for distributed computing. *Distributed computing*, 16(2-3):121–163, 2003.
5. Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *JACM*, 1985.

6 Guy Goren, Yoram Moses, and Alexander Spiegelman. Probabilistic indistinguishability and the quality of validity in byzantine agreement, 2020. arXiv:2011.04719.

7 Jonathan Katz and Chiu-Yuen Koo. On expected constant-round protocols for byzantine agreement. In *Annual International Cryptology Conference*, pages 445–462. Springer, 2006.

8 Nancy A Lynch. *Distributed algorithms*. Elsevier, 1996.

9 Amir Pnueli. On the extremely fair treatment of probabilistic algorithms. In *Proceedings of the fifteenth annual ACM symposium on Theory of computing*, pages 278–290, 1983.

10 Michael O Rabin. Probabilistic algorithms in finite fields. *SIAM Journal on computing*, 9(2):273–280, 1980.

11 Michael O Rabin. Randomized byzantine generals. In *24th Annual Symposium on Foundations of Computer Science (sfcs 1983)*, pages 403–409. IEEE, 1983.