Short Paper: On Game-Theoretically-Fair Leader Election

Rati Gelashvili, Guy Goren⋆, and Alexander Spiegelman

Novi Research

Abstract. This work studies the problem of game-theoretically-fair leader election. That is, provide fairness in the strong sense that the probability of any player being elected cannot be reduced even when facing an adversarial coalition of all other players. We extend a recent lower bound by [8] that shows that the tournament-tree protocol (based on Blum [5]) is optimal in the number of rounds, among the protocols that are restricted to immediately open the cryptographic commitments.

Our argument works even if commitments can be opened at arbitrary times, which is an open question left by [8]. To this end, we make two technical assumptions, one of which is weaker than the prior restriction and both of which are satisfied by the tournament-tree protocol, even if all players commit to their randomness for the entire execution in the beginning. The resulting proof is simple and streamlined, which we hope facilitates further research into an unconditional lower bound (or a new upper bound).

1 Introduction

Leader election is a fundamental task in distributed computing [16]. A natural predicate to compute in a multiparty setting, it provides the symmetry breaking power of a designated leader that often plays a key role in efficient distributed protocols for complex problems.

Leader election is closely related to *shared-coin* [4], another important distributed task of generating a shared value with guarantees about its distribution, based on the local randomness of protocol participants. Since an honest leader can toss a coin locally and share the outcome, the problem can be reduced to electing an honest leader. Leader election can also be viewed as an n-way shared coin-toss, and in fact, shared-coin is a commonly used building block for randomized distributed protocols [19].

There is a vast research spanning over four decades into the resilience of coin-flipping protocols against adversarial corruptions, e.g. [10,11,20,13]. The classic lower bound by Cleve [10] shows that a strong version of fairness called *unbiasability*, isn't achievable in the presence of a corrupt majority. A simple,

⋆ The work of Guy Goren was partly supported by a grant from the Technion Hiroshi Fujiwara cyber security research center and by the Israel Science Foundation under grant 2061/19.

well-known protocol of Blum [5], however, allows two participants to achieve a weaker, *game-theoretic* notion of fairness via leader election. Participants first commit to binary values, then they open the commitments, and the XOR of the values determines the leader (who determines the shared-coin value). The result can only be biased by deviating from the protocol and not revealing the commitment. In this case, the other participant becomes the leader, hence an adversarial participant can't improve the likelihood of its desired outcome.

Recent papers [9,21] have characterized the computational landscape of game-theoretic fairness and incentive compatibility of n-party shared-coin protocols. The impossibility results established for large coalitions of participants contrast with a tournament-tree generalization (standard construction akin [1,17,3]) of Blum's 2-party protocol, which maintains game-theoretic fairness even against coalitions of size $n-1$. In particular, adversarial behavior can never decrease the chance of any honest participant becoming the leader as compared to the system where all participants behave honestly.

The leader-driven nature of consensus in state machine replication [15,6,14,7], with the emerging economics and incentives in blockchain systems [18], further motivates exploring fair leader election protocols. The tournament tree protocol requires $O(\log n)$ rounds of communication. Chung, Chan, Wen and Shi [8] showed a clever lower bound argument that $\Omega(\log n)$ rounds are required for game-theoretically fair leader election in the standard broadcast model with a perfect commitment scheme, albeit when committed values must be opened immediately after they are broadcast (a restriction that still includes Blum's protocol and a variant of the tournament tree generalization). This is a foundational result, and in general, proving tight logarithmic lower bounds for similar tournament-based leader election protocols is notoriously difficult [1,2].

Not requiring commitments to be immediately opened significantly weakens the adversary controlling the coalition and complicates lower bound arguments, as the corrupted participants can no longer determine their messages based on the actual execution (*adaptive* adversary), but only based on the distribution of all possible executions (*oblivious* adversary). To contrast the adversarial power, note that in standard asynchronous shared memory with crash failures, a protocol with $O(\log^\star n)$ step complexity exists against the oblivious adversary [12], while the best known protocol against the adaptive adversary is the $O(\log n)$ tournament-tree of [1]. The authors in [8] emphasize the important dependency of their lower bound on immediately opening the commitments, and the need for more sophisticated techniques to overcome these in the context of their proof.

We take a step in better understanding this dependency on opening commitments and adversarial power for the problem of game-theoretically fair leader election. We remove the restriction on the commitment scheme in [8] but our lower bound also makes two assumptions about the protocol. Our lower bound applies to a different set of protocols, including the tournament-based protocol even if all commitments are made in the beginning of the protocol.

Our first assumption is that in any sub-protocol (a protocol that can be reached by rounds of honest execution), any player that has an overall positive probability to win when all players are honest, also has a chance to win regardless of the set of messages of the other players (in the first round of the sub-protocol). On the other hand, immediately opening commitments restriction in [8] implies that all sub-protocols are also game-theoretically fair, which in turn implies a stronger version of our assumption (that, in addition, the positive probabilities against fixed message vectors of other players are all equal to each other).

Our second assumption is more technical and captures the idea of an essential inductive ingredient in the previous proof of [8]. While we cannot claim that this assumption is weaker or follows from the previous commitment requirement, it is trivially satisfied by the tournament-tree protocol and leads to a simple, streamlined argument. Proving an unconditional lower bound (or a better upper bound) is an important open problem, and having a proof based on a different set of assumptions could help gain intuition about the general problem.

2 Model

We strive to remain close to the structure and notation of [8]. Therefore, we consider a standard synchronous round-based broadcast model with n participants that will be called *players* that communicate via a broadcast channel. As in [8], we restrict our attention to the case when the set of messages that a player i may send in each round is finite — denoted by M_i — to avoid non-measurable and other technical issues. Without loss of generality, we restrict to the case when in each round, an honest player i uniformly samples a message to send from M_i. Also like in [8], we assume that the $|M_i|$ is the same in every round. This is justified since we can construct equivalent protocols by sampling over multiple copies of every message.

In [8], it was assumed that a protocol could use a perfect commitment scheme to make the adversary commit to its randomness. However, the adversary could determine the message in a round based on the transcript of all previous rounds. This corresponds to the restriction on the protocol to immediately open every commitment. Without this restriction, the adversary must determine the messages of corrupted players ahead of time, and can only rely on the distribution of possible executions as opposed to the actual execution unfolding. This is a known, major distinction between the *adaptive* and *oblivious* randomized adversarial models.

In our setting, a round consists of each (non-crashed) player attempting to broadcast a message, while the adversary can *rush* to *crash*: it observes the messages and decides which players to crash[1]. After this, the messages of all non-crashed players appear on the broadcast channel. A crashed player remains crashed in all subsequent rounds, with its messages treated as $\perp$. Finally, notice that we assumed that a corrupted player j always sends a message from M_j. This is without loss of generality since otherwise, the corruption would be detectable on the broadcast channel, so the adversary could instead just crash the player j.

[1] Intuitively, this corresponds to not opening a commitment

2.1 Coalition Resistant Protocols

A 0-round leader election protocol must elect a unique, single winner among n players (without any communication). An r-round protocol is defined recursively, where processes engage in a round of communication, and proceed to an $(r-1)$-round leader election protocol.

Given a leader election protocol φ, let $p_i(\varphi)$ be the probability of player i winning over all failure-free executions. We call $p(\varphi) \in [0,1]^n$ the *winning probability distribution* of φ. We simply write p (or p_i) whenever the parameter φ is clear from the context. Since the adversary may choose not to corrupt any players, p_i upper bounds the minimum probability of player i winning against the adversary that can corrupt all players except i. We call a protocol φ *coalition resistant* if player i's probability of winning is p_i regardless of any adversarial strategy, which may control coalitions of size up to $n-1$.

For any multi-round protocol φ, we say φ' is a *sub-protocol* of φ if it can be reached by a finite number of rounds in which all players act honestly according to φ. Since a sub-protocol φ' is reachable by an all-honest execution, we can define the winning probability distribution p'_i for φ' analogous to the definition of p_i for φ (i.e., considering failure-free executions only). For a sub-protocol φ, let $S(\varphi)$ denote the support of the winning probability distribution of φ. Formally, for a probability distribution $p(\varphi) \in [0,1]^n$, we have $S(\varphi) \coloneqq \{i \in [n] : p_i > 0\}$. For any vector μ of possible messages for all n players, let $\varphi(\mu)$ denote the sub-protocol reached by one round of φ in which each player sends the corresponding message.

We will use the notation $\varphi(X_1 \leftarrow x_1, \dots, X_k \leftarrow x_k)$ when $X_1, \dots, X_k$ is a partition of all players and x_i is a vector of messages sent by players in X_i. I.e., this is the same as $\varphi(\mu)$ for a μ determined by $x_1, \dots, x_k$.

For any sub-protocol φ', we call a subset $A \subseteq S(\varphi')$ of players a *winning subset* of φ' if for any possible vector a' of the non-A players, there exists a vector of messages a for the players in A, such that $S(\varphi'(A \leftarrow a, [n] \setminus A \leftarrow a')) \subseteq A$, and in addition, for some a' there exists an a such that $S(\varphi'(A \leftarrow a, [n] \setminus A \leftarrow a')) = A$. In other words, players in A always have messages that eliminate all other players from contention (regardless of the messages of non-A players), and there is at least one possibility that all players in A maintain a chance to win.

Our lower bound applies to any coalition resistant protocol φ that satisfies the following two conditions.

Assumption 1. *For any sub-protocol φ' of φ, any player $i \in S(\varphi')$, and any possible vector x of messages for the non-$\{i\}$ players $X = [n] \setminus \{i\}$, there exists a message m_i of player i such that $i \in S(\varphi'(X \leftarrow x, \{i\} \leftarrow m_i))$.*

In other words, if player i had a positive probability of winning, there is no possible combination of messages that the other players may send such that i can no longer win after one round. As noted in the introduction, this requirement is weaker than having the sub-protocol φ' being also coalition resistant (as is implied by the immediately opening every commitment constraint in [8]).

Assumption 2. *For any sub-protocol φ' of φ with a support size $|S(\varphi')| > 1$, we can find two disjoint winning subsets of players $A, B \subset S(\varphi')$.*

Note that the tournament-tree based protocol satisfies these assumptions regardless of whether the players commit to messages at each round or at the beginning of the protocol. A sub-protocol is just a level in the tournament tree consisting of pairs of players engaging in 2-player Blum mechanism. For each of these pairs, we can place one player in subset A and the other in B, satisfying Assumption 2.

3 Lower Bound

Let φ be a coalition resistant protocol that satisfies Assumption 1 and Assumption 2. We prove by induction that for any r-round sub-protocol φ' of φ (consisting of the last r rounds of φ), it holds that $S(\varphi') \leq 2^r$. Consequently, a protocol that elects one out of n possible leaders requires at least $\log n$ rounds. Clearly, 0-round sub-protocols do not send any messages so there is just one possible execution and since the protocol must elect a unique winner, the induction base is satisfied for $r = 0$.

For the induction step, let us consider any $(r+1)$-round sub-protocol φ' and define subsets $A, B \subset S(\varphi')$ satisfying Assumption 2. By the induction hypothesis and the definition of a winning subset, $|A| \leq 2^r$ and $|B| \leq 2^r$. To complete the argument, we show that $A \cup B = S(\varphi')$, which will give $S(\varphi') = |A| + |B| \leq 2^{r+1}$ as desired.

Suppose for contradiction that disjoint subsets A and B do not include all players in $S(\varphi')$, and let i be some player among the rest of the players (in $S(\varphi') \setminus (A \cup B)$). Let D be the (possibly empty) set of all remaining players (in $S(\varphi') \setminus (A \cup B \cup \{i\})$) — for these players we will set a vector of messages d throughout the following argument.

Let $M(A, B)$ be a set of pairs of vectors (a, b) of messages for players in A and B, that allow only players in A or only players in B to retain a chance to win. Formally, $(a, b) \in M(A, B)$ iff there exists m with $S(\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m, D \leftarrow d)) \subseteq A$ or $S(\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m, D \leftarrow d)) \subseteq B$. A (and B) are winning subsets by Assumption 2, thus, $M(A, B)$ is non-empty.

We assign a *valency* to each element $(a, b) \in M(A, B)$, defined as the number of different messages m_i for player i, such that i retains a chance to win, i.e. $i \in S(\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m_i, D \leftarrow d))$. For the rest of the argument, let (a, b) be the element in $M(A, B)$ with the minimum valency. Suppose, without loss of generality that there exists i's message m_A such that only players in A retain a chance to win, i.e. $S(\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m_A, D \leftarrow d)) \subseteq A$ (the case for B is symmetrical, and one of these cases hold by the definition of $M(A, B)$).

First, we prove that for any message m' of i for which i does not retain a chance to win in $\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m', D \leftarrow d)$, only players in A retain a chance to win.

Lemma 1. *For m' with $i \notin S(\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m', D \leftarrow d))$, we have $S(\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m', D \leftarrow d)) \subseteq A$.*

Proof. Suppose for contradiction that m' allows a player $j \neq i$ that is not in A to retain a chance to win, i.e. $j \in S(\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m', D \leftarrow d))$.

We show an adversary that contradicts the coalition resistance of protocol φ. The adversary acts as follows in sub-protocol φ': it observes the messages of all players once revealed, and then it might choose to crash player i so that its message is not delivered.

Let p_A be the winning probability distribution of $\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m_A, D \leftarrow d)$ and let p' be the winning probability distribution of $\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m', D \leftarrow d)$. Recall that, by definition of m_A and the lemma assumption on m', player i has probability 0 both in p_A and p'. Moreover, the probabilities of players in A sum to 1 in p_A, and player j has a non-zero probability in p'. Hence, the probabilities of A-players in p_A summed with the probability of j in p_j is larger than 1 and cannot be a probability vector. To determine the precise adversarial strategy, we consider a winning probability distribution $p_\perp$ for sub-protocol $\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow \perp, D \leftarrow d)$, where player i crashes.

In $p_\perp$, either player j has lower probability than in p', or some player in A has a lower probability than in p_A. In the first case, the adversary crashes player i when it observes messages $A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m', D \leftarrow d$, reducing the probability of player j winning. Otherwise, the adversary crashes player i when it observes $A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m_A, D \leftarrow d$, reducing the probability of a player in A winning.

This contradicts the fact that φ' is a sub-protocol of a coalition resistant protocol φ. The adversary only crashes player i when an all-honest execution reaches sub-protocol φ', which by definition of a sub-protocol occurs by a positive probability. This still reduces the overall winning probability of some honest player in the original protocol, giving the desired contradiction. □

Let m_i be a message for which player i retains a chance to win in $\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m_i, D \leftarrow d)$. By Assumption 1, all valencies are positive, so such an m_i exists. Because by Assumption 2 B is a winning subset, there exists a vector b' of messages for B-players such that $S(\varphi'(A \leftarrow a, B \leftarrow b', \{i\} \leftarrow m_i, D \leftarrow d) \subseteq B$.

Next, we prove that

Lemma 2. *$i \in S(\varphi'(A \leftarrow a, B \leftarrow b', \{i\} \leftarrow m_i, D \leftarrow d)$.*

Proof. We start by showing that for any m' that satisfies $S(\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m', D \leftarrow d)) \subseteq A$, we have $i \notin S(\varphi'(A \leftarrow a, B \leftarrow b', \{i\} \leftarrow m', D \leftarrow d))$. For contradiction, assume $S(\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m', D \leftarrow d)) \subseteq A$ and $i \in S(\varphi'(A \leftarrow a, B \leftarrow b', \{i\} \leftarrow m', D \leftarrow d))$ for some message m'. However, the same adversarial strategy as in the proof of Lemma 1 but by replacing the role of player j in the previous lemma by player i in this lemma and crashing players in B (in the current lemma) instead of player i (in the previous lemma), contradicts the coalition-resistance of the protocol φ.

Applying Lemma 1, we get that for any m' such that $i \notin S(\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m', D \leftarrow d))$, we have $i \notin S(\varphi'(A \leftarrow a, B \leftarrow b', \{i\} \leftarrow m', D \leftarrow d))$.

Note that (a, b') is in $M(A, B)$ by definition of b'. By the choice of (a, b) with the minimum valency, the valency of (a, b') is at least as large as the valency of (a, b). Since player i always sends one out of the same number of possible messages, the valency of (a, b) and (a, b') are the same. Moreover the set of messages for which the valency is counted is also the same. In particular, since for m_i we have $i \in S(\varphi'(A \leftarrow a, B \leftarrow b, \{i\} \leftarrow m_i, D \leftarrow d))$ we also get that $i \in S(\varphi'(A \leftarrow a, B \leftarrow b', \{i\} \leftarrow m_i, D \leftarrow d))$. □

However, since $S(\varphi'(A \leftarrow a, B \leftarrow b', \{i\} \leftarrow m_i, D \leftarrow d)) \subseteq B$ and $i \notin B$ we get the desired contradiction and complete the induction.

4 Conclusion

The elegant proof by [8] that showed a lower bound of $\log n$ rounds for coalition-resistant leader election, left open a question of relaxing a restriction on the protocols to immediately open all cryptographic commitments.

We take a step in this direction by removing this restriction. In particular, our lower bound captures the standard tournament-tree protocol even if all message commitments are made in the beginning ("static" adversarial behavior). However, we require a new assumption for our proof that may help viewing the open problem of the unconditional round complexity of coalition-resistant leader election in a different light - i.e. when attempting to circumvent this assumption by a clever algorithm or a stronger lower bound.

References

1. Yehuda Afek, Eli Gafni, John Tromp, and Paul M. B. Vitányi. Wait-free test-and-set (extended abstract). In *Proceedings of the 6th International Workshop on Distributed Algorithms*, WDAG '92, pages 85–94, 1992.
2. Dan Alistarh, Rati Gelashvili, and Giorgi Nadiradze. Lower bounds for shared-memory leader election under bounded write contention. In *Proceedings of the 35th International Symposium on Distributed Computing*, DISC '21, 2021.
3. Massimo Bartoletti and Roberto Zunino. Constant-deposit multiparty lotteries on bitcoin. In *International Conference on Financial Cryptography and Data Security*, pages 231–247, 2017.
4. Michael Ben-Or and Nathan Linial. Collective coin flipping, robust voting schemes and minima of banzhaf values. In *Proceedings of the 26th Symposium on Foundations of Computer Science*, FOCS '85, pages 408–416, 1985.
5. Manuel Blum. Coin flipping by telephone a protocol for solving impossible problems. *ACM SIGACT News*, 15(1):23–27, 1983.
6. Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In *Proceedings of the 3rd Symposium on Operating Systems Design and Implementation*, OSDI '99, pages 173–186, 1999.

7. Benjamin Y Chan and Elaine Shi. Streamlet: Textbook streamlined blockchains. In *Proceedings of the 2nd Conference on Advances in Financial Technologies*, AFT '20, pages 1–11, 2020.
8. Kai-Min Chung, T-H Hubert Chan, Ting Wen, and Elaine Shi. Game-theoretic fairness meets multi-party protocols: The case of leader election. In *Annual International Cryptology Conference*, pages 3–32, 2021.
9. Kai-Min Chung, Yue Guo, Wei-Kai Lin, Rafael Pass, and Elaine Shi. Game theoretic notions of fairness in multi-party coin toss. In *Theory of Cryptography Conference*, pages 563–596, 2018.
10. Richard Cleve. Limits on the security of coin flips when half the processors are faulty. In *Proceedings of the 18th ACM symposium on Theory of Computing*, STOC '86, pages 364–369, 1986.
11. Uriel Feige. Noncryptographic selection protocols. In *Proceedings of the 40th Symposium on Foundations of Computer Science*, FOCS '99, pages 142–152, 1999.
12. George Giakkoupis and Philipp Woelfel. Efficient randomized test-and-set implementations. *Distributed Computing*, 32(6):565–586, 2019.
13. Iftach Haitner and Yonatan Karidi-Heller. A tight lower bound on adaptively secure full-information coin flip. In *Proceedings of the 61st Symposium on Foundations of Computer Science*, FOCS '20, pages 1268–1276, 2020.
14. Ramakrishna Kotla, Lorenzo Alvisi, Mike Dahlin, Allen Clement, and Edmund Wong. Zyzzyva: speculative byzantine fault tolerance. In *Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles*, pages 45–58, 2007.
15. Leslie Lamport. Paxos made simple. *ACM Sigact News*, 32(4):18–25, 2001.
16. Nancy A Lynch. *Distributed algorithms*. Elsevier, 1996.
17. Andrew Miller and Iddo Bentov. Zero-collateral lotteries in bitcoin and ethereum. In *EuroS&PW Workshop*, pages 4–13, 2017.
18. Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review*, page 21260, 2008.
19. Michael Rabin. Randomized byzantine generals. In *Proceedings of the 24th Symposium on Foundations of Computer Science*, FOCS '83, pages 403–409, 1983.
20. Alexander Russell, Michael Saks, and David Zuckerman. Lower bounds for leader election and collective coin-flipping in the perfect information model. *SIAM Journal on Computing*, 31(6):1645–1662, 2002.
21. Ke Wu, Gilad Asharov, and Elaine Shi. A complete characterization of game-theoretically fair, multi-party coin toss. *https://eprint.iacr.org/2021/748*, 2021.