

Perfect Information Leader Election in $\log^* n + O(1)$ Rounds

ALEXANDER RUSSELL*
Computer Science Division
University of California
Berkeley, CA 94720
acr@cs.berkeley.edu

DAVID ZUCKERMAN†
Department of Computer Science
University of Texas
Austin, TX 78712
diz@cs.utexas.edu

*Most of this work was done while the author was a postdoctoral fellow at the University of Texas at Austin. Supported by NSF NYI Grant No. CCR-9457799 and a David and Lucile Packard Fellowship for Science and Engineering.

†Part of this work was done while the author was on leave at the University of California at Berkeley. Supported in part by NSF NYI Grant No. CCR-9457799, a David and Lucile Packard Fellowship for Science and Engineering, and an Alfred P. Sloan Research Fellowship.

Abstract

In the leader election problem, n players wish to elect a random leader. The difficulty is that some coalition of players may conspire to elect one of its own members. We adopt the perfect information model: all communication is by broadcast, and the bad players have unlimited computational power. Protocols proceed in rounds: though players are synchronized between rounds, within each round the bad players may wait to see the inputs of the good players. A protocol is called resilient if a good leader is elected with probability bounded away from 0.

We give a simple, constructive leader election protocol that is resilient against coalitions of size βn , for any $\beta < 1/2$. Our protocol takes $\log^* n + O(1)$ rounds, each player sending at most $\log n$ bits per round. For any constant k , our protocol can be modified to take k rounds and offer resilience against coalitions of size $\epsilon n / (\log^{(k)} n)^3$, where ϵ is a small enough constant and $\log^{(k)}$ denotes the logarithm iterated k times. This is constructive for $k \geq 3$.

The primary component of the above protocols is a new collective sampling protocol: for a set S of large enough (polynomial) size, this protocol generates an element $s \in S$ in a single round so that for any subset $T \subset S$, $\Pr[s \in T] \leq |T|/|S|^{-\alpha(1-\beta)}$ for a constant $\alpha > 0$.

1 Introduction

This paper is about three related problems which arise naturally in the study of distributed computing: leader election, collective sampling, and collective coin-flipping. We begin with a discussion of coin-flipping, since this is perhaps the most basic of the three.

In a distributed computing environment common random bits may be required. Collective coin-flipping is the problem of obtaining such bits if some processors are faulty. If people are behind the processors, the faults may be malicious; this is the case, for example, when coin flips are needed to gamble over the Internet [HS97]. Following Ben-Or and Linial [BL90], we assume that faults may be malicious, that all communication is by broadcast, and that the sender of every message is known with certainty. Processors may broadcast messages simultaneously.

The simplest method for n processors, called *players*, to generate a collective random bit is as follows. A suitable function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is chosen in advance. Then each player broadcasts a random $r_i \in \{0, 1\}$, and the collective random bit is taken to be $r = f(r_1, \dots, r_n)$.

We allow a subset $B \subset [n]$ of bad players to collude to bias the resulting bit. In particular, they may not choose their r_i 's randomly. One obtains different models depending on whether the distributed environment is synchronous and whether the bad players' computational power is limited. This paper focuses on the most difficult of these possibilities.

In a synchronous environment, the players cannot see other players' choices for r_i . Thus, PARITY will output a perfectly unbiased bit if even one player is honest. On the other hand, our model assumes an asynchronous environment: although messages are supposed to be sent in parallel, they may be sent in any order. Therefore, the bad players may wait to see the honest players' choices before they act. In this case, PARITY is foiled by just one bad player.

If the bad players' computational power is restricted to polynomial-time, then the players can use cryptography to communicate with each other privately (assuming sufficiently strong cryptography). The resulting problem is related to Byzantine agreement. To avoid relying on unproven assumptions and to obtain the strongest possible results, our model allows unlimited computational power for the bad players. This is called the *perfect information model*, and was first introduced in the context of collective coin-flipping by Ben-Or and Linial [BL90].

A function f is called *resilient* if it gives rise to a robust coin-flipping protocol:

Definition 1. A family of functions $f_n : \{0, 1\}^n \rightarrow \{0, 1\}$, $n = 1, 2, \dots$, is called $b(n)$ -resilient if there exists $\gamma > 0$ such that for all n and $B \subseteq [n]$ with $|B| \leq b(n)$, regardless of the strategy of the players in B ,

$$\gamma \leq \Pr[f_n(r_1, \dots, r_n) = 1] \leq 1 - \gamma.$$

Thus, for example, MAJORITY is $c\sqrt{n}$ -resilient, for any positive c . The most resilient functions known were shown to exist by Ajtai and Linial (there are non-constructive parts to their proof):

Theorem 1 ([AL93]). There exists a family of functions which is $\epsilon n / \log^2 n$ -resilient, for a small enough positive constant ϵ .

There is also a lower bound:

Theorem 2 ([KKL88]). If $b(n) = \omega(n / \log n)$, then no family of functions is $b(n)$ -resilient.

In order to achieve larger resilience, we enrich the class of protocols under consideration, allowing the protocols to last many *rounds* and allowing players to send many bits in each round. Each round is asynchronous: within a round, the bad players may wait to see the communication of the good players. Between rounds, the processors are synchronized. The notion of resilience is extended in a natural way to this multi-round scenario.

We now broaden the discussion to include leader election protocols. In this case, the protocol is supposed to pick a uniformly random *leader* among the n processors. Resilience is then defined as follows:

Definition 2. A leader election protocol is called $b(n)$ -resilient if there is a constant $\gamma < 1$ which upper bounds the probability that any coalition of size $b(n)$ can elect one of its own members.

Note that if there is a k -round leader election protocol, then there is a $k + 1$ -round coin-flipping protocol with the same resilience: in the last round the leader may flip the coin.

Source	Resilience	Rounds	Constructive?	Bits/Round (Each Player)
[Sak89]	$O(\frac{n}{\log n})$	n	Yes	$\log n$
[AN93]	$O(n)$	$n^{O(1)}$	Yes	1
	$(\frac{1}{3} - \epsilon)n$	n	No	1
[BN]	$(\frac{1}{2} - \epsilon)n$	n	No	1
[CL95]	$O(n)$	$(\log n)^{O(1)}$	Yes	1
[ORV94]	$O(n)$	$O(\log n)$	Yes	$n^{O(1)}$
	$(\frac{1}{2} - \epsilon)n$	$O(\log n)$	No	$n^{O(1)}$
[Zuc97]	$(\frac{1}{2} - \epsilon)n$	$O(\log n)$	Yes	$\log n$
THIS PAPER	$(\frac{1}{2} - \epsilon)n$	$\log^*(n) + O(1)$	Yes	$\log n$
	$O\left(\frac{n}{(\log^{(k)} n)^3}\right)$	k	For $k \geq 3$	$\log n$

Figure 1: Historical summary.

One example of a leader election protocol is the *baton passing* protocol. Initially, player 1 holds the baton. In each round, the player holding the baton passes it to a player who has not yet held the baton. The last player to hold the baton is called the leader. Saks [Sak89] showed that if the honest players toss the baton randomly (among those players who have not yet touched the baton), this protocol is $\epsilon n / \log n$ -resilient for a small enough positive constant ϵ . Saks also observed that no protocol can be $\lceil n/2 \rceil$ -resilient (see [BN] for a proof).

The last decade has witnessed remarkable improvement in our understanding of this problem, culminating in constructive, $O(\log n)$ -round protocols which are βn -resilient [ORV94, Zuc97] for any fixed $\beta < 1/2$. The historical summary in Figure 1 briefly charts this progress. We present a constructive leader election protocol requiring only $\log^* n + O(1)$ rounds to achieve βn -resilience, for any $\beta < 1/2$. This protocol can be modified to yield improved constant round protocols, offering $\epsilon n / (\log^{(k)} n)^3$ -resilience in k rounds for a small enough constant ϵ . This is constructive for $k \geq 3$.

These protocols rely on a new protocol for *collective sampling*. The collective sampling problem is a generalization of the problems discussed above: the objective of a *collective sampling protocol for S* is to produce an element $s \in S$ in a suitably robust fashion. Typically, the set S varies with the number of players (as in the leader election problem), and a collective sampling protocol for S guarantees that for every target subset $T \subset S$, $\Pr[s \in T]$ is suitably small.

Goldreich, Goldwasser, and Linial [GGL91] introduced the collective sampling problem, and demonstrated a collective sampling protocol for which

$$\Pr[s \in T] \leq (|T| / |S|)^{1-c\mu(B)}$$

where $\mu(B)$ is the fraction of corrupt players and $c > 0$ is some constant. This is optimal up to the constant c . Note that such a bound on $\Pr[s \in T]$ gives a “negligibility property”: if $|T|/|S| = o(1)$ then $\Pr[s \in T] = o(1)$.

Their protocol has a couple of disadvantages. First, $\mu(B)$ has to be a small enough constant (less than $1/c$). Second, their protocol takes many rounds, consisting of $\log |S|$ metarounds where each metaround consists of a polynomial number of sequential calls to a collective coin-flipping subroutine.

Here we remove these disadvantages, and give a one-round protocol achieving

$$\Pr[s \in T] \leq |T||S|^{-\alpha(1-\mu(B))}$$

for some $\alpha > 0$ and large enough polynomial $|S|$. The running time is polynomial in $|S|$, unless $|S| \geq 2^n$, in which case a simple algorithm running in time linear in $\log |S|$ will suffice. Although our bound on $\Pr[s \in T]$ is useful for any $\mu(B) < 1$, it doesn’t yield the negligibility property. Observe, however, that it is unrealistic to achieve their bound in one-round: if this were possible, then taking $|T| = 1$ and $|S| = 2$ would yield a one round collective coin-flipping protocol.

We note that subsequent to our work, Feige [Fei] gave a simpler leader election protocol requiring the same number of rounds as ours. Although he discusses sampling under the term selection, his work does not appear to offer a comparable collective sampling protocol.

Finally, we remark that if the bad players’ computational power were restricted to polynomial-time, and if sufficiently strong cryptography exists, then the Byzantine agreement protocol of Feldman and Micali [FM97] may be used to achieve an $\lfloor (n-1)/3 \rfloor$ -resilient leader election protocol that takes a constant number of expected rounds.

The paper is organized as follows. In Section 2 we present the necessary background; in Section 3 we present the one-round collective sampling protocol; in Section 4 we present the leader election protocol; and in Section 5 we present constant-round variants of these protocols.

2 Preliminaries

We denote the set $\{1, \dots, n\}$ by $[n]$. The logarithm base 2 is denoted $\log n$ and the natural logarithm $\ln n$. In general, we ignore rounding errors when their effect is insignificant.

Two combinatorial constructions shall be instrumental in the development of our protocol: a “balanced” poly-logarithmic set system and a hitting set for combinatorial rectangles. These are introduced below in §2.1 and §2.2. As a final preparatory step, §2.3 is devoted to bounding a class of recurrence relations related to the protocol.

2.1 Balanced Set Systems and Committee Sampling via Extractors

A paradigm appearing frequently in the leader election literature is the recursive application of “committee” selection. Briefly, the description of the n -player protocol includes a collection of (overlapping) committees of the n players, each of size $n' \ll n$. A collective sampling protocol is invoked to select a committee from this collection, which removes from consideration all players but those in the selected committee. The remaining players then carry out the n' -player protocol to elect the final leader. Assuming that some β fraction of the players are corrupt, a natural property to desire on the part of this family of committees is that regardless of which subset of the players are corrupt, very few of the committees have much more than a β fraction of corrupt members. If the sampling protocol we apply is suitably robust, we can then recurse on an appropriately balanced collection of players. Specifically, the committees we use shall have the properties outlined in Definition 4, below.

Definition 3. A subset $B \subseteq [n]$ has *density* $\mu(B) = |B|/n$. A subset $C \subseteq [n]$ is called *B-saturated* if $|C \cap B| \geq (\mu(B) + 1/\log n) |C|$.

Definition 4. $\mathfrak{C}_n \subset 2^{[n]}$ is a *balanced set system* if

1. $\forall C \in \mathfrak{C}_n, |C| = (\log n)^{O(1)}$,
2. for any $B \subset X$, the number of B -saturated committees is $O(n^{1.1})$.

As one would expect, a random collection of $n^{O(1)}$ such sets can easily be shown to satisfy the above properties with high probability, proving existence. We need an explicit construction, which is supplied by extractor constructions (see [Nis96] for a survey of extractors and their applications). We restate the extractor construction we need in our framework, making use of the observation that if there is a balanced set system of size $f(n)$ and $g(n) \leq f(n)$, then there is one of size $g(n)$.

Theorem 3 ([Zuc97]). For all polynomial-time computable functions $g : \mathbb{N} \rightarrow \mathbb{N}$ with $g(n) = n^{O(1)}$, there is a polynomial-time constructible family of balanced set systems of size $g(n)$.

2.2 Hitting Sets for Combinatorial Rectangles

For a set S (such as the set of committees described above), our collective sampling protocol for S associates elements of S with members of a sparse “hitting set” for combinatorial rectangles, defined below.

Definition 5. A *combinatorial rectangle* R in $[a]^d$ is a cross product $R = R_1 \times \dots \times R_d$, with each $R_i \subset [a]$. The *volume* of such a set is $\text{vol}(R) = a^{-d} \cdot \prod_i |R_i|$.

Definition 6. An (a, d, δ) -*hitting set* is a set $H \subset [a]^d$ which intersects every combinatorial rectangle of volume at least δ . When the universe is understood, such a set will be referred to as an δ -*hitting set*.

An easy probabilistic proof shows that there exist (a, d, δ) -hitting sets of size $\lceil \ln(2)ad/\delta \rceil$. A constructive solution is offered by Linial, et. al., who prove the following theorem:

Theorem 4 ([LLSZ97]). There exists an (a, d, δ) -*hitting set* of cardinality $\text{poly}(\log(d)a/\delta)$ constructible in time $\text{poly}(ad/\delta)$.

2.3 A Lemma about Poly-logarithmic Decay

In order to avoid logarithms of negative numbers, we define iterated logarithms as follows. For $n \geq 1$ and $k \in \mathbb{N}$,

$$\log^{(k)} n = \begin{cases} 1 & \text{if } \log^{(k-1)} n < 2, \\ \log(\log^{(k-1)} n) & \text{otherwise,} \end{cases}$$

with $\log^{(0)} n = n$. Then, for $n \geq 1$, define $\log^*(n)$ to be the smallest natural number k for which $\log^{(k)} n = 1$. We will need the following lemma:

Lemma 5. Let $T : \mathbb{N} \rightarrow \mathbb{N}$ be a function given by the recurrence relation:

$$T(n) = \begin{cases} t_0 & \text{for } n \leq n_T, \\ 1 + T(f(n)) & \text{for } n > n_T, \end{cases}$$

for a function $f = (\log n)^{O(1)}$ and constants t_0 and n_T . Then $T(n) < \log^* n + O(1)$.

Proof. Choose c so that $f(n) < \lfloor (\log n)^c \rfloor$ for all sufficiently large n . For convenience, assume that $c > 2$. Then, defining $S(n)$ as

$$S(n) = \begin{cases} s_0 & \text{for } n \leq n_S, \\ 1 + S(\lfloor (\log n)^c \rfloor) & \text{for } n > n_S, \end{cases}$$

there is an appropriate choice of the constants n_S and s_0 so that S is well defined and, for all $n \in \mathbb{N}$, $T(n) \leq S(n)$. For convenience assume that $n_S > c^{4c}$. Now, for $n \geq 1$ and $k \in \mathbb{N}$ define $L^{(k)}(n)$ so that

$$L^{(k)}(n) = \begin{cases} 1 & \text{if } L^{(k-1)}(n) < 2, \\ \lfloor (\log L^{(k-1)}(n))^c \rfloor & \text{otherwise,} \end{cases}$$

with $L^{(0)}(n) = n$. Then $S(n) = s_0 + L^*(n)$, where $L^*(n)$ is the smallest k for which $L^{(k)}(n) \leq n_S$. We prove by induction on k that $L^{(k)}(n) \leq (c^4 \log^{(k)} n)^c$. The base case $k = 0$ is immediate. Assuming the inequality for $L^{(k)}(n)$, we have

$$\begin{aligned} L^{(k+1)}(n) &= \left\lfloor \left(\log L^{(k)}(n) \right)^c \right\rfloor \leq \left[c \log \left(c^4 \log^{(k)} n \right) \right]^c \\ &= \left(4c \log c + c \log^{(k+1)} n \right)^c \leq \left(c^4 \log^{(k+1)} n \right)^c, \end{aligned}$$

since $c > 2$. Recalling that $n_S > c^{4c}$, the lemma follows. □

3 A One Round Collective Sampling Protocol

We now turn our attention to the collective sampling problem. The sampling protocol below is the combinatorial core of the leader election protocol of Section 4.

Theorem 6. There is a constant $c > 0$ such that for any S of size at least n^c there is a one round collective sampling protocol for S so that for all $T \subset S$,

$$\Pr[s \in T] \leq |T| |S|^{-(1-\mu(B))/c}.$$

Furthermore, this protocol runs in time polynomial in $|S|$ and n . When $|S| \geq 2^n$ a naive protocol can achieve this bound, with $c \leq 2$, in time linear in $\log |S|$.

Proof. First we describe the naive protocol for large $|S|$. Suppose $|S| = 2^{sn}$ for some integer s . Then associate S with $\{0, 1\}^{sn}$, have each player output s random bits, and concatenate the bits of the players. It is easy to check that this achieves the desired bound with $c = 1$. In case $2^{sn} < |S| < 2^{(s+1)n}$ we may have some players flip s bits and others flip $s + 1$; this achieves the bound for $c = 1 + \mu(B)/s$.

We now turn to the more difficult case of smaller S . Assume that $|S| < 2^n$. Our starting idea is due to [ORV94]: each player eliminates a random $\Theta((\log |S|)/n)$ fraction of S . The lexicographically least element (say) that remains is the selected element. This protocol ensures that with high probability no element of T remains. Unfortunately, this would allow the bad players to eliminate every element of S .

Our key idea is to restrict the possible subsets of S that a player may eliminate. Below we give a method for this which prevents the players from eliminating all of S .

We shall associate S with the elements of a δ -hitting set H in $[a]^n$, for appropriately selected $\delta > 0$ and a . For an element $s \in S$, we let $\vec{h}(s) \in H$ denote the element of H associated with s . With such an association, the protocol proceeds as follows. Each player i broadcasts a random $r_i \in [a]$, which removes from consideration all elements $s \in S$ for which $h(s)_i = r_i$. The lexicographically least element in

$$R = \{s : \forall i, h(s)_i \neq r_i\}$$

is then the element selected from S .

Fixing a subset T of S , we must then insure that

1. if $|T|$ is small enough, then the probability that $T \cap R \neq \emptyset$ is small, and
2. R is non-empty.

Observe that if

$$\delta \leq \left(1 - \frac{1}{a}\right)^n = \text{vol}(\{\vec{v} \in [a]^n : \forall i, v_i \neq r_i\}),$$

then H contains an element of any set of form $\{\vec{v} \in [a]^n : \forall i, v_i \neq r_i\}$, so that item 2 is guaranteed. Focusing now on item 1, notice that for any $t \in T$,

$$\Pr[t \in R] \leq \left(1 - \frac{1}{a}\right)^{(1-\mu(B))n}$$

since the honest players select their r_i uniformly in $[a]$.

The statement of the theorem now follows by judicious selection of the parameters a and δ . Specifically, we shall be interested in the case when $\delta < 1/n$ and $a < n$, so that the association of S with H requires that $|S| \geq \delta^{-c} \geq \text{poly}(a\delta^{-1} \log n)$ for a constant c determined by Theorem 4. Assume that $c \geq 2$. Satisfaction of item 2 above demands that $\delta \leq \left(1 - \frac{1}{a}\right)^n$. So assign

$$|S|^{-\frac{1}{c}} = \delta = \left(1 - \frac{1}{a}\right)^n. \tag{1}$$

Observe now that

$$\Pr[\exists t \in T \cap R] \leq |T| \left(1 - \frac{1}{a}\right)^{(1-\mu(B))n} \leq |T| |S|^{\frac{-(1-\mu(B))}{c}},$$

as desired. Finally, we observe that an acceptable value of a is induced by equation (1):

$$a^{-1} = 1 - \delta^{\frac{1}{n}} = 1 - e^{\frac{\ln \delta}{n}} = 1 - e^{-\frac{\ln |S|}{cn}}$$

so that

$$a^{-1} > \frac{\ln |S|}{cn} - \frac{1}{2} \left(\frac{\ln |S|}{cn}\right)^2 = \omega\left(\frac{1}{n}\right)$$

and hence $a = o(n)$ (recall that $|S| \geq n$); similarly, since $|S| < 2^n$,

$$a^{-1} < \frac{\ln |S|}{cn} \leq \frac{\ln 2}{c}$$

and hence $a > 2$ (recall that $c \geq 2$), as desired. \square

4 The Leader Election Protocol

The protocol we present below is recursive, each step discarding all but a small committee of players. The base case invokes the following result of Boppana and Narayanan:

Theorem 7 ([BN]). For every $\beta < \frac{1}{2}$, there is a leader election protocol resilient against coalitions of size βn .

Although this is non-constructive in general, we need the result only for a specific (constant) value n_0 so that the protocol can, of course, be found by exhaustive search (trying all possible protocols and strategies for the bad players). Feige [Fei] has observed that since n_0 is constant, at this stage one can in fact use a simple one-round protocol in lieu of Theorem 7.

Our protocol selects a committee of size $(\log n)^{O(1)}$ in a single round, so we focus on functions

$$f_n : X^n \rightarrow \left\{ C \subset [n] : |C| \leq (\log n)^{O(1)} \right\}$$

where X is some appropriately selected domain.

Lemma 8. For all $\beta < 1$, there is a polynomial-time computable family of functions

$$f_n : X^n \rightarrow \left\{ C \subset [n] : |C| \leq (\log n)^{O(1)} \right\}$$

so that for any set $B \subset [n]$ of size at most βn , the probability that for a random setting of the variables outside B , some setting of the variables of B produces a B -saturated committee $f(x_1, \dots, x_n)$ is at most $O(1/n)$. The set X can be taken to be $\{0, 1\}^{\log n}$.

Proof. Let c_s be the constant guaranteed by Theorem 6 and set $c > \frac{(2.1)c_s}{1-\beta}$. From Theorem 3, there is a balanced set system \mathfrak{C} of subsets of $[n]$ of size n^c . Applying the one round collective sampling protocol of Theorem 6, the probability that a B -saturated committee is selected is at most

$$O(n^{1.1}) \left(n^{-\frac{c(1-\beta)}{c_s}} \right) = o\left(\frac{1}{n}\right)$$

by our choice of c . \square

Theorem 9. For all $\beta < \frac{1}{2}$, there is a $\log^* n + O(1)$ round leader election protocol resilient against coalitions of size βn .

Proof. We apply Lemma 8 recursively until the resulting number of players is at most n_0 , a suitable constant to be chosen later. We then apply Theorem 7. Lemma 5 shows that this protocol does indeed terminate in $\log^* n + O(1)$ rounds.

Fix $\beta < \frac{1}{2}$. There are two types of error to control. First, there is $\hat{\beta}(n)$, the maximum possible resulting fraction of bad players when the protocol begins with n players (βn of which are corrupt), assuming only unsaturated committees were chosen at each step. Then

$$\hat{\beta}(n) \leq \hat{\beta}((\log n)^{O(1)}) + 1/\log n.$$

By choosing n_0 large enough, we can ensure that $\hat{\beta}(n)$ is bounded away from $\frac{1}{2}$ for all n , which is what we need to apply Theorem 7.

Second, there is the error $E(n)$ that, with n starting players, a B -saturated committee is chosen somewhere in the recursion. This error satisfies

$$E(n) \leq E((\log n)^{O(1)}) + O(1/n).$$

This can be made arbitrarily small by choosing n_0 large enough. In fact, we only need it to be less than 1, since we ensure that conditional on reaching n_0 players with all unsaturated committees, there is a constant probability that the protocol given by Theorem 7 will select a good leader. \square

5 Constant Round Protocols

The requirement that the fraction of corrupt players, β , be (a constant) less than $1/2$ manifests itself only in the base case of the above leader election protocol. Indeed, the recursive committee selection process (i.e. Lemma 8) is well behaved for any $\beta < 1$. Returning momentarily to the collective sampling problem, this observation induces a k -round collective sampling protocol, for $k = O(1)$, with

$$\Pr [s \in T] \leq |T||S|^{-\frac{(1-\beta-o(1))}{c}}$$

assuming that $|S| \geq (\log^{(k-1)} n)^{c'}$ for an appropriate constant $c' > 0$ (recall that the protocol of Section 3 required that $|S| \geq n^c$). Sampling in a set of this size is achieved by selecting, in $k - 1$ rounds, a committee C of players for which with high probability

- $\frac{|C \cap B|}{|C|} < \mu(B) + O\left(\frac{1}{\log n}\right)$, and
- $|C| \leq |S|^{\frac{1}{c}}$, where c is the constant of Theorem 6,

and then applying the protocol of Section 3. The error in this protocol is dominated by the error in the last round.

In similar fashion, coupling Lemma 8 with the $\epsilon n / (\log n)^2$ -resilient functions of Ajtai-Linial (see Theorem 1, above), we now present k -round leader election protocols which, for small enough $\epsilon_k > 0$, are $\epsilon_k n / (\log^{(k)} n)^3$ -resilient.

5.1 The Functions of Ajtai-Linial and Sub-linear Coalitions

Definition 7. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be a boolean function on variables $\{x_1, \dots, x_n\}$. The *influence* of a set $S \subseteq \{x_1, \dots, x_n\}$ on f , written $I_f(S)$, is the probability that the function is undetermined by a random setting of the variables outside S .

Ajtai and Linial [AL93] have shown the existence of a family of functions for which the influence of any set of $\epsilon n / (\log n)^2$ variables is $O(\epsilon)$. As the base case of our constant round constructions, we need a family of functions for which the influence of any set of $\epsilon n / (\log n)^3$ variables is $O(\epsilon / \log n)$. A simple adaptation of the proof in [AL93] shows that the functions they construct also enjoy this property. For completeness, we briefly outline their construction adapted to the case we need. We also provide a streamlined proof of one portion of their result.

Theorem 10 (Adapted from [AL93]). There is a sequence of boolean functions f_n on $n = 1, 2, \dots$ variables, having expectation $\frac{1}{2} + o(1)$, such that for any $c > 2$ and $\epsilon > 0$, for any large enough n , the influence of any set of $\epsilon n / (\log n)^c$ variables is $O(\epsilon / (\log n)^{c-2})$. The time to construct such a function deterministically is $n^{O(n^2)}$.

Proof. For a positive integer b , let n be the smallest multiple of b for which $(1 - 2^{-b})^{\frac{n}{b}} \leq \frac{\ln 2}{n}$. Then $b = \log n - 2 \log \log n + o(1)$ and $(1 - 2^{-b})^{\frac{n}{b}} \geq \frac{\ln 2}{n} \left(1 - \frac{(\ln n)^2}{n}\right)$. For such a pair b, n , let \mathcal{P} be the collection of all partitions of $\{1, \dots, n\}$ into classes of size b . The collection of sequences $\mathbf{P} = (P^1, \dots, P^n)$ with each $P^i \in \mathcal{P}$ is denoted

$\vec{\mathcal{P}}$. Defining \mathcal{M} to be the collection of all mappings $m : \{1, \dots, n\} \rightarrow \{0, 1\}$, the collection of all sequences $\mathbf{m} = (m_1, \dots, m_n)$ with each $m_i \in \mathcal{M}$ is denoted $\vec{\mathcal{M}}$. Finally, for $\mathbf{P} \in \vec{\mathcal{P}}$ and $\mathbf{m} \in \vec{\mathcal{M}}$, let $f = f_{\mathbf{P}, \mathbf{m}}$ be the function

$$f(x_1, \dots, x_n) = \bigwedge_{1 \leq i \leq n} \bigvee_{1 \leq j \leq n/b} \bigwedge_{k \in P_j^i} (x_k = m_i(k)),$$

where P_j^i denotes the j th class of the partition P^i . For convenience, let

$$f^i(x_1, \dots, x_n) = \bigvee_{1 \leq j \leq n/b} \bigwedge_{k \in P_j^i} (x_k = m_i(k)).$$

Definition 8. A partition $P \in \mathcal{P}$ and a set $B \subset \{1, \dots, n\}$ are said to *match* if for each $1 \leq k \leq b$, the number of classes P_j of P with $|B \cap P_j| \geq k$ does not exceed

$$2^k \binom{\frac{n}{b} \binom{b}{k}}{\binom{|B|}{n}}.$$

Notice that if the partition P is selected randomly, then the probability that a certain P_j contains more than k elements of B is at most $\binom{b}{k} \left(\frac{|B|}{n}\right)^k$, whence the expected number of such P_j is at most $\frac{n}{b} \binom{b}{k} \left(\frac{|B|}{n}\right)^k$.

The proof proceeds in four steps:

1. For all \mathbf{P} , and almost all \mathbf{m} , the expectation of $f_{\mathbf{P}, \mathbf{m}}$ is $\frac{1}{2} + o(1)$.
2. For almost all \mathbf{P} and every set $B \subset \{1, \dots, n\}$ with $|B| = \epsilon n / (\log n)^c$, the number of partitions P^i in \mathbf{P} failing to match B is less than $n / (\log n)^{\omega(1)}$.
3. There is a constant $\epsilon_0 > 0$ so that for any partition P^i in \mathcal{P} and any $m_i \in \mathcal{M}$, the influence of any set $B \subset \{1, \dots, n\}$ with $|B| \leq \epsilon_0 n / (\log n)^2$ on f^i is at most $\frac{1}{n}$.
4. If P^i and B match then the influence of B on f^i is $O(\epsilon / (n(\log n)^{c-2}))$.

Steps 1 and 3 are exactly Propositions 5.4 and 5.1 of [AL93].

Proof of Step 4. (cf. Proposition 5.2 of [AL93].) Fix f^i , given by P^i and m_i , and a matching set B . Notice that an assignment to the variables outside of B leaves f^i undetermined only when

1. every P_j^i not meeting B contains a variable x_k for which $x_k \neq m_i(k)$, and
2. for some P_j^i , meeting B , the assignment completely agrees with m_i .

These two events are independent. The probability of event 1 is at most

$$(1 - 2^{-b})^{\frac{n}{b} - |B|} \leq \left(\frac{\ln 2}{n}\right)^{1 - \frac{b|B|}{n}} = O\left(\frac{1}{n}\right).$$

Now focus on event 2. For a fixed class P_j^i with $|P_j^i \cap B| = k$, the probability that $x_s = m_i(s)$ for all $s \in P_j^i \setminus B$ is 2^{k-b} . Since P^i matches B , the probability of event 2 is bounded above by

$$\sum_{1 \leq k \leq b} 2^k \frac{n}{b} \binom{b}{k} \left(\frac{|B|}{n}\right)^k 2^{-(b-k)} = \frac{n}{b2^b} \left[\left(1 + \frac{4|B|}{n}\right)^b - 1 \right] \leq \frac{n}{b2^b} \left[\exp\left(\frac{4|B|}{n}\right) - 1 \right].$$

Recalling that $b = \log n - 2 \log \log n + o(1)$, we have $b2^b \geq (1 - o(1)) \frac{n}{\log n}$ so that the above sum is

$$O\left(\frac{\epsilon}{(\log n)^{c-2}}\right).$$

□

Anticipating the proof of step 2, we record Azuma's inequality for discrete martingales.

Definition 9. A *martingale* is a sequence X_1, X_2, \dots, X_n of real valued random variables for which $\mathbb{E}[X_{i+1} | X_i] = X_i$.

Theorem 11 (Azuma's Inequality, [Hoe63, Azu67]). Let X_1, \dots, X_n be a martingale with $|X_i - X_{i-1}| \leq 1$. Then

$$\Pr [X_n - \mathbb{E}[X_n] > \lambda\sqrt{n}] \leq e^{-\frac{\lambda^2}{2}}.$$

See [AS92, §7] for a general discussion of discrete martingales and a proof of Azuma's inequality.

Proof of Step 2. For convenience fix a specific partition P^i and consider the uniform probability space on subsets B of $\{1, \dots, n\}$ of size $\frac{\epsilon n}{\log^c n}$. Let \mathcal{E}_k be the event that

$$|\{j : |P_j^i \cap B| \geq k\}| \geq \frac{2^k n}{b} \binom{b}{k} \left(\frac{|B|}{n}\right)^k.$$

Then $\Pr[P^i \text{ matches } B] = 1 - \Pr[\cup_k \mathcal{E}_k] \geq 1 - \sum_k \Pr[\mathcal{E}_k]$. As observed earlier, the expected number of P_j containing more than k elements of B is less than $\frac{n}{b} \binom{b}{k} \left(\frac{|B|}{n}\right)^k$. Then, focusing our attention on those $k \geq \sqrt{\log n}$, an application of Markov's inequality shows that $\Pr[\mathcal{E}_k] \leq 2^{-k} = (\log n)^{-\omega(1)}$.

Suppose now that $k \leq \sqrt{\log n}$. Let X_1, \dots, X_n be indicator random variables given by $X_p = 1$ iff $p \in B$. Then define Y_1, \dots, Y_{nb-1} so that

$$Y_j = \begin{cases} 1 & \text{if } |P_j^i \cap B| \geq k, \\ 0 & \text{otherwise.} \end{cases}$$

and set $Y = \sum_j Y_j$. Our goal is to demonstrate strong tail bounds on the random variable Y . Finally, for $0 \leq p \leq n$ define

$$Z_p = \mathbb{E}[Y | X_1, \dots, X_p].$$

Then $Z_0 = \mathbb{E}[Y]$ is a constant random variable and $Z_n = Y$. Notice that, by definition, $\mathbb{E}[Z_{p+1} | Z_p] = Z_p$, so that these Z_p form a martingale. Furthermore, $|Z_{p+1} - Z_p| \leq 1$, the proof for which we defer for a moment. In this case, application of Azuma's inequality yields

$$\begin{aligned} \Pr[\mathcal{E}_k] &\leq \Pr \left[Y - \mathbb{E}[Y] \geq (2^k - 1) \left[\frac{n}{b} \binom{b}{k} \left(\frac{|B|}{n}\right)^k \right] \right] \\ &\leq \Pr \left[Z_n - \mathbb{E}[Z_n] \geq (2^k - 1) n^{1-o(1)} \right] \\ &\leq \exp \left(\frac{2^k - 1}{4} n^{1-o(1)} \right) = \frac{1}{\log^{\omega(1)} n}, \end{aligned}$$

since the quantity $\frac{c2^k n}{b} \binom{b}{k} \left(\frac{|B|}{n}\right)^k$ is at least $n^{1-o(1)}$.

Then $\sum_k \Pr[\mathcal{E}_k] = (\log n)^{-\omega(1)}$, and an application of Markov's inequality shows that with probability $1 - o(1)$ the number of P^i which do not match B is less than $n(\log n)^{-\omega(1)}$.

We return to the proof that $|Z_{p+1} - Z_p| \leq 1$. It suffices to show that for any $\vec{x} = (x_1, \dots, x_p) \in \{0, 1\}^p$,

$$|\mathbb{E}[Y | X_i = x_i (i \leq p), X_{p+1} = 1] - \mathbb{E}[Y | X_i = x_i (i \leq p), X_{p+1} = 0]| \leq 1. \quad (2)$$

The only interesting case is when $\mathbf{wt}(\vec{x}) < |B|$. We establish (2) by observing that it holds under further conditioning. In particular, for both conditioned probability spaces in (2), we think of first choosing a uniformly random set B' of $|B| - \mathbf{wt}(\vec{x}) - 1$ elements from $\{p+2, \dots, n\}$ to add to B . When $X_{p+1} = 1$, this condition determines B ; when $X_{p+1} = 0$, the last element of B is a random element from $\{p+2, \dots, n\} \setminus B'$. Conditioned on any such B' , then, the resulting Y 's can differ by at most 1, as we wanted. □

In [AL93], the above theorem is established for $c = 2$.

A function satisfying the conditions in the theorem can be found in time $n^{O(n^2)}$. This follows from two observations. First, $|\vec{\mathcal{P}}| \leq (n!)^n$, and $|\vec{\mathcal{M}}| = 2^{n^2}$, so the number of possible functions is $n^{O(n^2)}$. Second, a function can be tested for the desired property in exponential time. \square

5.2 Constant Round Leader Election Protocols

With Theorem 10 in hand, it is not difficult to show that there exist one round leader election protocols resilient against coalitions of size $O(n/(\log n)^3)$:

Lemma 12. There exists a family of functions $g_n : X^n \rightarrow [n]$ so that for any set of variables B of size $\frac{\epsilon n}{(\log n)^3}$, the probability that for a random setting of the variables outside B , there is a completion so that $F(\vec{x}) \in B$ is $O(\epsilon)$. The set X can be taken to be $\{0, 1\}^k$, for $k = O(\log n)$.

Proof. Consider the probability distribution where each $\vec{x}_i = x_{i1} \dots x_{ik}$ is selected independently and uniformly at random in $\{0, 1\}^{8 \log n}$. Set $Y_j = f_n(x_{1j}, x_{2j}, \dots, x_{nj})$ where f_n are the functions of Theorem 10. This is a sequence of independent $\frac{1}{2} + o(1)$ biased bits. To correct the biases, we use von Neumann's trick [vNe51]: collecting them into pairs, $Z_1 = (Y_1, Y_2), Z_2 = (Y_3, Y_4), \dots$, consider the string $N(Z_1)N(Z_2) \dots N(Z_{4 \log n})$ where

$$N(a, b) = \begin{cases} 1 & \text{if } a = 1, b = 0, \\ 0 & \text{if } a = 0, b = 1, \\ \Lambda & \text{if } a \oplus b = 0, \end{cases}$$

and Λ denotes the empty string. Then

$$\Pr[N(Y_i, Y_{i+1}) = 1] = \Pr[N(Y_i, Y_{i+1}) = 0] = \frac{1}{4} + o(1).$$

Applying Chebyshev's inequality, we see that with probability $1 - o(1)$, this sequence of independent and unbiased values has length at least $\lceil \log n \rceil$. When this sequence is long enough, the first $\lceil \log n \rceil$ bits are used to produce a value v in $[n]$, which is the value of g_n on these \vec{x}_i . Otherwise $g_n(\vec{x}_i) = 1$. The mapping ϕ from $\{0, 1\}^{\lceil \log n \rceil}$ to $[n]$ used to induce v can be chosen so that $\forall B \subset [n], |\phi^{-1}(B)| \leq 2|B|$.

Fix a collection of variables B of size at most $\frac{\epsilon n}{(\log n)^3}$. Notice that $\Pr[g_n(\vec{x}_1, \dots, \vec{x}_n) \in B] \leq 2\mu(B) + o(1) = o(1)$. From the bound of Theorem 10, the probability that a random $\vec{x}_i, i \notin B$, results in a function which is non-constant on the variables of B is at most $k \cdot O(\epsilon / \log n) = O(\epsilon)$, which establishes the lemma.

N.b. It is in fact true that $\mathbb{E}[f_n] = \frac{1}{2} + o(\frac{1}{\log n})$, so that one can avoid "correcting" the bias of these Y_i , resulting in functions $g_n : \{0, 1\}^{\log n} \rightarrow [n]$. \square

Using this as a base case, the next result gives a k -round leader election protocol resilient against coalitions of size $\frac{\epsilon n}{(\log^{(k)} n)^3}$.

Theorem 13. For $k \in \mathbb{N}$, there is $\epsilon_k > 0$ for which there is a k round leader election protocol resilient against coalitions of size $\epsilon_k n / (\log^{(k)} n)^3$. This is constructive for $k \geq 3$.

Proof. The first $k - 1$ rounds are given by the protocol of Lemma 8. The last round is given by the protocol of Lemma 12. The errors are handled as in Theorem 9. \square

6 Open Question

An outstanding open question is whether there exists a constant round leader election protocol resilient against linear-sized coalitions. It is unknown even whether there is such a one round protocol.

References

- [AL93] Miklós Ajtai and Nathan Linial. The influence of large coalitions. *Combinatorica*, 13(2):129–145, 1993.
- [AN93] Noga Alon and Moni Naor. Coin-flipping games immune against linear-sized coalitions. *SIAM Journal of Computing*, 22(2):403–417, April 1993.
- [AS92] Noga Alon and Joel H. Spencer. *The Probabilistic Method*. John Wiley & Sons, Inc., 1992.
- [Azu67] Kazuoki Azuma. Weighted sums of certain dependent random variables. *Tôhoku Math. J. (2)*, 19:357–367, 1967.
- [BL90] Michael Ben-Or and Nathan Linial. Collective coin flipping. In Silvio Micali, editor, *Randomness and Computation*, pages 91–115. Academic Press, New York, 1990.
- [BN] Ravi Boppana and Babu Narayanan. Perfect-information leader election with optimal resilience. To appear, *SIAM Journal on Computing*.
- [CL95] Jason Cooper and Nathan Linial. Fast perfect-information leader-election protocols with linear immunity. *Combinatorica*, 15:319–332, 1995.
- [Fei] Uriel Feige. Noncryptographic selection protocols. Preliminary version, January 15, 1999.
- [FM97] Pesech Feldman and Silvio Micali. An optimal probabilistic protocol for synchronous Byzantine agreement. *SIAM Journal on Computing*, 26(4):873–933, August 1997.
- [GGL91] Oded Goldreich, Shafi Goldwasser, and Nathan Linial. Fault-tolerant computation in the full information model (extended abstract). In *32nd Annual Symposium on Foundations of Computer Science*, pages 447–457, San Juan, Puerto Rico, 1–4 October 1991. IEEE.
- [Hoe63] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.*, 58:13–30, 1963.
- [HS97] Chris Hall and Bruce Schneier. Remote electronic gambling. In *13th Annual Computer Security Applications Conference*, pages 227–230. ACM, December 1997.
- [KKL88] Jeff Kahn, Gil Kalai, and Nathan Linial. The influence of variables on Boolean functions (extended abstract). In *29th Annual Symposium on Foundations of Computer Science*, pages 68–80, White Plains, New York, 24–26 October 1988. IEEE.
- [LLSZ97] Nathan Linial, Michael Luby, Michael Saks, and David Zuckerman. Efficient construction of a small hitting set for combinatorial rectangles in high dimension. *Combinatorica*, 17(2):215–234, 1997.
- [Nis96] Noam Nisan. Extracting randomness: How and why – a survey. In *Proceedings of the 11th Annual IEEE Conference on Computational Complexity*, pages 44–58, 1996.
- [ORV94] Rafail Ostrovsky, Sridhar Rajagopalan, and Umesh Vazirani. Simple and efficient leader election in the full information model. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on the Theory of Computing*, pages 234–242, Montréal, Québec, Canada, 23–25 May 1994.
- [Sak89] Michael Saks. A robust noncryptographic protocol for collective coin flipping. *SIAM Journal on Discrete Mathematics*, 2(2):240–244, May 1989.
- [vNe63] John von Neumann. *Collected Works*. Pergamon Press, 1961–63. Edited by A. H. Traub.
- [vNe51] John von Neumann. Various techniques used in connection with random digits. *Applied Math Series*, 12:36–38, 1951. Notes by G. E. Forsythe. National Bureau of Standards. Also appears in vol. 5 of [vNe63].
- [Zuc97] David Zuckerman. Randomness-optimal oblivious sampling. *Random Structures and Algorithms*, 11:345–367, 1997.