SIAM J. COMPUT.
Vol. 31, No. 6, pp. 1645–1662
© 2002 Society for Industrial and Applied Mathematics

LOWER BOUNDS FOR LEADER ELECTION AND COLLECTIVE COIN-FLIPPING IN THE PERFECT INFORMATION MODEL*

ALEXANDER RUSSELL†, MICHAEL SAKS‡, AND DAVID ZUCKERMAN§

Abstract. Collective coin-flipping is the problem of producing common random bits in a distributed computing environment with adversarial faults. We consider the *perfect information* model: all communication is by broadcast and corrupt players are computationally unbounded. Protocols in this model may involve many asynchronous rounds. We assume that honest players communicate only uniformly random bits. We demonstrate that any n-player coin-flipping protocol that is resilient against corrupt coalitions of linear size must use either at least $[1/2 - o(1)] \log^* n$ communication rounds or at least $[\log^{(2k-1)} n]^{1-o(1)}$ communication bits in the kth round, where $\log^{(j)}$ denotes the logarithm iterated j times. In particular, protocols using one bit per round require $[1/2-o(1)] \log^* n$ rounds. These bounds also apply to the leader election problem. The primary component of this result is a new bound on the influence of random sets of variables on Boolean functions. Finally, in the one-round case, using other methods we prove a new bound on the influence of sets of variables of size βn for $\beta > 1/3$.

Key words. perfect information model, collective coin-flipping, leader election

AMS subject classifications. 68Q17, 91A15, 05D40

PII. S0097539700376007

1. Introduction.

Collective coin-flipping is the problem of producing a common random bit in a distributed computing environment with adversarial faults. We consider the *perfect information* model introduced by Ben-Or and Linial [5], which can be informally described as follows. A protocol in this model consists of a sequence of rounds. In each round, each player privately generates a uniformly random string of bits of some specified length (possibly 0) and broadcasts the string. Each broadcast is received by all players and the identity of the sender is known with certainty. The round ends after all broadcasts are received. After the completion of all rounds, the outcome of the protocol is computed separately by each player as a prespecified function of all the values broadcast during the protocol; for the coin-flipping problem the outcome is a single bit. A protocol Π is said to be an (n, r, ℓ)-protocol if n is the number of players, r is the number of rounds, and each player broadcasts at most ℓ bits in each round.

Faults are modeled by the presence of an unknown set of b corrupt players who collude in order to bias the outcome. Players are assumed to be computationally unbounded. In addition, the system is not able to enforce perfect synchrony within a

---

*Received by the editors July 28, 2000; accepted for publication (in revised form) April 18, 2002; published electronically September 5, 2002. A preliminary version of this paper appeared in *Proceedings of the* 31*st ACM Symposium on Theory of Computing*, 1999, pp. 339–347.

http://www.siam.org/journals/sicomp/31-6/37600.html

†Department of Computer Science and Engineering, University of Connecticut, Storrs, CT 06269 (acr@cse.uconn.edu). This research was done while the author was a postdoctoral fellow at the University of Texas at Austin. The research of the first author was supported by NSF NYI grant CCR-9457799 and a David and Lucile Packard Fellowship for Science and Engineering.

‡Department of Mathematics, Rutgers University, Piscataway, NJ 08854 (saks@math.rutgers.edu). The research of the second author was supported by NSF grant CCR-9700239. This work was done in part while on sabbatical at University of Washington.

§Department of Computer Science, University of Texas, Austin, TX 78712 (diz@cs.utexas.edu). The research of the third author was supported in part by NSF NYI grant CCR-9457799, a David and Lucile Packard Fellowship for Science and Engineering, and an Alfred P. Sloan Research Fellowship.

round; thus in each round, the corrupt players may wait to see the broadcasts of the other players before selecting their strings.

While not necessary for previous upper bounds, we strengthen our lower bounds by assuming that corrupt players may cheat only in ways that are undetectable to the other players. This means that when the protocol specifies that such a player broadcast a bit string of a given length, he must do so; however, he may cheat by broadcasting a string that he chooses rather than a random string.

The simplest protocol is one that designates a single player to flip a coin, the value of which is the outcome of the protocol; of course, this is unsatisfactory if that player happens to be faulty. More generally, an $(n, 1, 1)$-protocol is defined by a Boolean function $f : \{0, 1\}^n \to \{0, 1\}$; each player i broadcasts a bit r_i and the outcome is $f(r_1, \dots, r_n)$. Throughout the paper we use the terms Boolean function and $(n, 1, 1)$-protocol interchangeably.

The primary goal in designing a protocol is to ensure that it can tolerate as many cheaters as possible.

DEFINITION 1.

1 *Let Π be a coin-flipping protocol for n players, and let $\gamma \in (0, 1/2]$.*
   (a) *For $B \subseteq [n]$, Π is* (B, γ)-resilient *if for any strategy of the players in B,*

$$\gamma \leq \Pr[\Pi \textit{ has outcome } 1] \leq 1 - \gamma,$$

   where the probability is taken with respect to the random bits generated by the players outside of B.
   (b) *Π is (b, γ)-resilient for an integer $b \leq n$ if it is (B, γ)-resilient for all B with $|B| \leq b$.*

2. *Let $\Pi = (\Pi_n : n \geq 1)$ be a sequence where Π_n is an n-player protocol, and let $b(n)$ be a function mapping n to an integer $b(n) \leq n$. We say that Π is* $b(n)$-resilient *if there exists $\gamma > 0$ (independent of n) such that for all n, Π_n is $(b(n), \gamma)$-resilient.*

For example, in the case of $(n, 1, 1)$-protocols, the PARITY function, $\sum_{i=1}^n r_i \bmod 2$, is not even 1-resilient, while the MAJORITY function is $c\sqrt{n}$-resilient for any positive c. Ajtai and Linial [1] constructed a Boolean function that is $\Omega(n/\log^2 n)$-resilient. Kahn, Kalai, and Linial, in a 1988 tour de force, proved an upper bound on the resilience of Boolean functions.

THEOREM 2 (see [13]). *If $b(n) = \omega(\frac{n}{\log n})$, then no sequence $(f_n : n \geq 1)$ is $b(n)$-resilient.*

We emphasize that this bound applies only to $(n, 1, 1)$-protocols. Indeed, Alon and Naor [2] showed that there are protocols using n rounds that are $\Omega(n)$-resilient, and this was followed by a sequence of papers giving more efficient protocols with linear resilience. In what follows $\log^{(k)}(n)$ denotes the maximum of 1 and the kth iterated base 2 logarithm, and $\log^* n$ is the least integer k such that $\log^{(k)}(n) = 1$. The most efficient known protocol is that of [16], requiring $\log^* n + O(1)$ rounds; players send messages of length $O(\log^{(k)} n)$ during the kth round. The protocol achieves βn-resilience for any $\beta < 1/2$. (As noted in [17] no protocol can be $n/2$-resilient.) This protocol can be modified to yield a one bit per round protocol with $\left[1 + o(1)\right] \log n$ rounds. Subsequently, Feige [9] gave a simpler protocol with similar properties.

Despite rapid progress in our understanding of *protocols* for the problem, very little beyond Theorem 2 was known on the negative side. The major contribution of this article is an extension of Theorem 2 to protocols with many rounds. We will prove the following theorem.

THEOREM 3. *Let* $\Pi = (\Pi_n : n \geq 1)$ *be a sequence of protocols, where* Π_n *is an* $(n, r(n), 1)$*-protocol for* $r(n) \leq \frac{1}{2}\log^* n - \log^* \log^* n$*. Then*

1. Π_n *is not* $\Omega(n)$*-resilient and*
2. *if*

$$b(n) = \omega\left(\frac{(r(n))^2}{\log^{(2r(n)-1)} n} \cdot n\right),$$

then Π *is not* $b(n)$*-resilient.*

For instance, when $r(n) = 1$ this reduces to Theorem 2, and when $r(n) = 2$ it implies that no $(n, 2, 1)$-protocol can be $\omega(n/\log\log\log n)$-resilient.

We extend the notation above to describe protocols with variable communication complexity: a protocol Π is said to be an $(n, r, \vec{\ell})$-protocol if n is the number of players, r is the number of rounds, and no more than ℓ_k bits are broadcast by any player in the kth round, where $\vec{\ell} = (\ell_1, \dots, \ell_r)$. We will prove that the conclusion of Theorem 3 holds even if we relax the requirement that each player sends only one bit per round.

THEOREM 4. *There is a function* $\eta : \mathbb{N} \to [0, 1]$ *with* $\eta(n) = o(1)$ *so that for any sequence* $\Pi = (\Pi_n : n \geq 1)$ *of protocols, where* Π_n *is an* $(n, r(n), \vec{\ell})$*-protocol with*

$$r(n) \leq \frac{1}{2}\log^* n - \log^* \log^* n \qquad \text{and} \qquad \ell_k(n) \leq (\log^{(2k-1)} n)^{1-\eta(n)},$$

Π_n *is not* $\Omega(n)$*-resilient.*

Recall that current upper bounds provide $(n, \log^* n + O(1), \vec{\ell})$-protocols which are linearly resilient, where $\ell_k = O(\log^{(k)} n)$.

The *leader election* problem is that of selecting a "leader" among n players so that the probability that any coalition (of appropriate size) can elect one of its own members is at most $1 - \epsilon$ for a constant $\epsilon > 0$ independent of n. Adopting the above model, the notion of resilience may be extended to this scenario. Collective coin-flipping may be reduced to leader election at the cost of an extra round: the leader may flip a fair coin. Our bounds shall then naturally apply to this problem as well. For a more detailed discussion of coin-flipping, leader election, and the perfect information model, see [7, 14].

Section 2 gives definitions, notation, and preliminary facts. The two main theorems are proved in section 3 and section 4. In section 5 an observation is made about the behavior of large linear sized coalitions. We conclude with some open questions.

2. Preliminaries.

2.1. General notation. Throughout, $\ln x$ denotes the natural logarithm and $\log x$ the logarithm base 2. To avoid logarithms of negative numbers, iterated logarithms are defined inductively as follows: for $x \geq 1$, $\log^{(0)}(x) = x$, and for $k \geq 1$,

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

For $x \geq 1$, define $\log^*(x)$ to be the smallest natural number k for which $\log^{(k)} x = 1$.

For a positive real number y and integer k the *tower* function $\mathrm{T}(k; y)$ is defined by

$$\mathrm{T}(0; y) = y, \text{ and}$$
$$\mathrm{T}(k; y) = 2^{\mathrm{T}(k-1;y)} \text{ for } k > 0.$$

Observe that for any $y \geq 1$, $k \leq \ell$, $\log^{(k)}(\mathrm{T}(\ell; y)) = \mathrm{T}(\ell - k; y)$.

For an integer n, we denote the set $\{1, \dots, n\}$ by $[n]$. For $J \subseteq [n]$, a finite set X, and $\alpha \in X^J$, $C(\alpha)$ denotes the set of all points $x \in X^n$ such that $x_j = \alpha_j$ for all $j \in J$. If $\alpha \in X^J$ and $\beta \in X^{[n]\setminus J}$, then $[\alpha : \beta]$ denotes the unique point of $\{0,1\}^n$ belonging to $C(\alpha) \cap C(\beta)$.

If S is a set, the notation $x \in_U S$ indicates that x is selected uniformly at random from S.

2.2. Coin-flipping protocols and influence. We want to formalize the definition of protocol given in the introduction. We below define $(n, r, \vec{\ell})$-protocols and a number of related notions; (n, r, ℓ)-protocols, where communication is constant across rounds, are covered as a special case. For an $(n, r, 1)$-protocol Π we suppress the third index and simply say that Π is an (n, r)-protocol.

Formally, an $(n, r, \vec{\ell})$-*protocol* is a function

$$\Pi : (\{0,1\}^{\ell_1})^n \times \cdots \times (\{0,1\}^{\ell_r})^n \to \{0,1\}.$$

Such a protocol is executed in r rounds. In the presence of a set $B \subset [n]$ of bad players, the protocol operates as follows. In round i, the players in $[n] \setminus B$ select $\alpha^i \in (\{0,1\}^{\ell_i})^{[n]\setminus B}$ uniformly at random. Then, depending on $\alpha^1, \dots, \alpha^i$, the players in B choose their values. Formally, an $(n, r, \vec{\ell})$-*strategy* for B is a sequence $S = (S_1, S_2, \dots, S_r)$ of functions where

$$S_i : \left(\{0,1\}^{\ell_1}\right)^{[n]\setminus B} \times \cdots \times \left(\{0,1\}^{\ell_i}\right)^{[n]\setminus B} \to \left(\{0,1\}^{\ell_i}\right)^B.$$

The function S_i specifies the choices of the bad players in round i as a function of the choices of the good players in the first i rounds. The outcome of protocol Π, with bad player set B playing strategy S, is a function of the sequence

$$\vec{\alpha} = (\alpha^1, \dots, \alpha^r) \in \left(\{0,1\}^{\ell_1}\right)^{[n]\setminus B} \times \cdots \times \left(\{0,1\}^{\ell_r}\right)^{[n]\setminus B}$$

of the random coins of the good players, which is denoted $\Pi(\vec{\alpha}; S)$ and is defined to be

$$\Pi\big([\alpha^1 : S_1(\alpha^1)], \dots, [\alpha^r : S_r(\alpha^1, \dots, \alpha^r)]\big).$$

Definition 5. *For a protocol* Π, $B \subseteq [n]$, *and strategy* S,

- $p^1_\Pi(B; S)$ *denotes the probability that* $\Pi(\vec{\alpha}; S) = 1$ *if*

$$\vec{\alpha} \in_U \left(\{0,1\}^{\ell_1}\right)^{[n]\setminus B} \times \cdots \times \left(\{0,1\}^{\ell_r}\right)^{[n]\setminus B};$$

- $p^1_\Pi(B)$ *is the maximum of* $p^1_\Pi(B; S)$ *over all strategies* S;
- $p^1_\Pi = p^1_\Pi(\emptyset)$, *the* natural probability of Π, *is the probability that the outcome is* 1 *if there are no bad players;*
- $I^1_\Pi(B)$, *the* influence of B towards 1, *is defined to be* $p^1_\Pi(B) - p^1_\Pi$;
- $p^0_\Pi(B; S)$, $p^0_\Pi(B)$, p^0_Π, *and* $I^0_\Pi(B)$ *are defined analogously;*
- $I_\Pi(B)$, *the* influence of B on Π, *is defined to be*

$$I^1_\Pi(B) + I^0_\Pi(B).$$

An $(n, 1)$-protocol corresponds to a Boolean function $f : \{0, 1\}^n \to \{0, 1\}$, and we typically use the letter f (instead of Π) for such a protocol. It is not hard to see that $p_f^1(B)$ is the probability, with respect to $\alpha \in_U \{0, 1\}^{[n]\setminus B}$, that $1 \in f(C(\alpha))$ and that $I_f(B)$ is the probability that f is not constant on $C(\alpha)$. Furthermore, if $|B| = 1$, then $I_f^1(B) = I_f^0(B)$.

The following result, observed in [6] (cf. Proposition 2.2 of [13]), implies that the most resilient one-round protocols are given by Boolean functions that are monotone.

PROPOSITION 6. *For any Boolean function f, there exists a monotone Boolean function g on the same set of variables for which*

1. $p_f^1 = p_g^1$ *and*
2. *for all* $B \subset [n]$, $I_f^1(B) \geq I_g^1(B)$ *and* $I_f^0(B) \geq I_g^0(B)$.

Finally, we need a variant of a fact first noted in [10] and based on a result in [13], which asserts that if no variable in a Boolean function has large influence, then the average influence of a variable cannot be too small. For completeness, we include a proof.

LEMMA 7. *Let $\gamma \in (0, \frac{1}{2})$ and $\theta \in (0, \frac{1}{8})$. Let $f : \{0, 1\}^n \to \{0, 1\}$ be a Boolean function with $p_f^1 \in (\gamma, 1 - \gamma)$. If $I_f(\{i\}) \leq \theta$ for each $i \in [n]$, then*

$$\sum_{i\in[n]} I_f(\{i\}) \geq \frac{\gamma \log(\frac{1}{\theta})}{20}.$$

Proof. Let $\bar{v} \in \mathbb{R}^n$ denote the vector with $v_i = I_f(\{i\})$. For $p > 0$, the l_p norm of $\bar{v}$, denoted $\|\bar{v}\|_p$, is defined to be $(S_p)^{1/p}$, where $S_p = \sum |v_i|^p$.

By complementing if necessary, we may assume that $p_f^1 \leq 1/2$. Since the function f is Boolean and $p_f^1 \geq \gamma$, [13, eq. (3.4.1)] asserts that for any $\delta \in (0, 1)$ and $t \geq 1$,

$$\delta^{-t} S_{\frac{2}{1+\delta}} + t^{-1} S_1 \geq \frac{\gamma}{2}. \tag{2.1}$$

Inequality (2.10.1) of Hardy, Littlewood, and Pólya [11] asserts that

$$S_r \leq (S_q)^{\frac{s-r}{s-q}} (S_s)^{\frac{r-q}{s-q}}$$

for $0 < q < r < s$, which is equivalent to

$$(S_r)^{\frac{s-q}{s}} \leq (S_q)^{\frac{s-r}{s}} (\|v\|_s)^{r-q}.$$

Setting $q = 1$ and $r = \frac{2}{1+\delta}$ and letting s tend to ∞, we obtain

$$S_{\frac{2}{1+\delta}} \leq S_1 \theta^{\frac{1-\delta}{1+\delta}}.$$

Substituting this into the inequality (2.1) and setting $\delta = 1/2$ we get

$$\left(2^t \theta^{\frac{1}{3}} + \frac{1}{t}\right) S_1 \geq \frac{\gamma}{2}.$$

Choose t such that $\theta = 2^{-3t}/t^3$ (noting that $t > 1$ since $\theta < 1/8$). Then the previous inequality implies $S_1 \geq t\gamma/4$. Since $2^{-3t}/t^3 \geq 2^{-5t}$ for $t \geq 1$, we have

$$t \geq \frac{1}{5} \log\left(\frac{1}{\theta}\right),$$

and therefore

$$S_1 \geq \frac{\gamma \log \frac{1}{\theta}}{20}. \qquad \square$$

2.3. A tail bound for submartingales. Our main theorems are proved by considering a certain stochastic process which, for a Boolean function, selects a set of variables likely to have large influence. Our analysis of this stochastic process involves a tail bound for submartingales, which we record below.

DEFINITION 8. *A* submartingale *is a sequence of real valued random variables* $Z_0, Z_1, \ldots$ *for which* $\mathbb{E}[Z_i \mid Z_{i-1}] \geq Z_{i-1}$.

We were unable to find the exact form of the following tail bound in the literature, so we have included a proof. The basic method is developed in [8, 4, 12]. Our treatment follows [3, 15].

LEMMA 9. *Let* $(Z_i : i \in \{0, \ldots, n\})$ *form a submartingale with* $Z_0 = 0$. *Define* $X_i = Z_i - Z_{i-1}$ *for* $i \in \{1, \ldots, n\}$ *and assume that* $X_i \in [0,1]$ *and* $\mathbb{E}[X_i \mid Z_{i-1}] \geq \mu_i$. *Setting* $\mu = \sum_i \mu_i$ *and* $Z = Z_n$,

$$\Pr[Z < (1-\delta)\mu] < e^{-\frac{\delta^2 \mu}{2}}$$

for all $\delta > 0$.

Proof. Observe that for any $\alpha > 0$,

$$\Pr[Z < (1-\delta)\mu] = \Pr\left[e^{-\alpha Z} > e^{-\alpha(1-\delta)\mu}\right] < \frac{\mathbb{E}\left[e^{-\alpha Z}\right]}{e^{-\alpha(1-\delta)\mu}}.$$

Letting $\ell(x) = 1 + x(e^{-\alpha} - 1)$, we have $e^{-\alpha x} \leq \ell(x)$ for all $x \in [0,1]$ because the exponential function is convex. For any $[0,1]$ valued random variable Y,

$$\mathbb{E}\left[e^{-\alpha Y}\right] \leq \mathbb{E}\left[\ell(Y)\right] = 1 + \mathbb{E}[Y](e^{-\alpha} - 1).$$

By induction, we compute

$$\begin{aligned}
\mathbb{E}\left[e^{-\alpha Z_k}\right] &= \mathbb{E}\left[e^{-\alpha Z_{k-1}} \cdot e^{-\alpha X_k}\right] \\
&= \mathbb{E}\left[\left(e^{-\alpha Z_{k-1}}\right) \mathbb{E}\left[e^{-\alpha X_k} \,\middle|\, Z_{k-1}\right]\right] \\
&\leq \mathbb{E}\left[\left(e^{-\alpha Z_{k-1}}\right)\left(1 + \mathbb{E}\left[X_k \,\middle|\, Z_{k-1}\right](e^{-\alpha} - 1)\right)\right] \\
&\leq \prod_i (1 + \mu_i(e^{-\alpha} - 1)) < e^{\sum_i^k \mu_i (e^{-\alpha} - 1)}.
\end{aligned}$$

Hence

$$\Pr\Big[Z < (1-\delta)\mu\Big] < \frac{e^{\mu(e^{-\alpha} - 1)}}{e^{-\alpha(1-\delta)\mu}}.$$

Setting $\alpha = \ln(\frac{1}{1-\delta})$, we have

$$\Pr\Big[Z < (1-\delta)\mu\Big] < \left(\frac{e^{-\delta}}{(1-\delta)^{(1-\delta)}}\right)^{\mu} \leq e^{-\frac{\delta^2 \mu}{2}},$$

since $(1-\delta)^{(1-\delta)} > e^{-\delta + \frac{\delta^2}{2}}$. ◻

3. Proof of Theorem 3. We begin by considering (n, r)-protocols; each round consists of a single bit broadcast by each player. Fix the integer r. We say that a protocol Π is α-*nontrivial* if the natural probability of Π is at least α, i.e., $p_\Pi^1 \geq \alpha$, terminology that we apply also in the multibit case. By complementing the output if necessary, we may assume that the protocol is 1/2-nontrivial. We want to show that if Π is an (n, r)-protocol, then for n sufficiently large there is a set B of $b \ll n$ players so that B can almost always force the outcome to 1. For $r = 1$ this follows from Theorem 2.

We illustrate the ideas for $r > 1$ by looking at the two round case. By separating the inputs associated with each round, a two round protocol may be viewed as a function to functions:

$$\Pi : \{0,1\}^n \to \{g : \{0,1\}^n \to \{0,1\}\} .$$

As Π is 1/2-nontrivial, many g's will be 1/4-nontrivial (any constant less than 1/2 would do) and by Theorem 2, for each such g there is a sublinear set of players B_2 that can force this g to be 1 with high probability. Also by Theorem 2, there is a sublinear set of players B_1 that can likewise force the output of Π to be one of these g's. A natural strategy is to choose $B = B_1 \cup B_2$.

The problem with this plan is that B_2 depends on g; we really need one B_2 that works for many g's. We show this by proving that a random B_2 will work with significant probability for any 1/4-nontrivial g. It follows that a random B_2 will work for many g's. For general r, we will proceed by induction, with our inductive assumption being that a random sublinear set of players can control the protocol with significant probability.

To make these ideas rigorous, we begin with some definitions. For $\beta \in [0, 1]$, we say that a subset B is *β-powerful in* Π if $p^1_\Pi(B) \geq 1 - \beta$.

DEFINITION 10. *Let $C_n(r; \alpha, \beta)$ (written $C_n(r; \gamma)$ when $\alpha = \beta = \gamma$) denote the collection of pairs $\langle \delta, b \rangle$ so that for any (n, r)-protocol Π that is α-nontrivial, at least a δ fraction of sets $B \subset [n]$ of size b are β-powerful in Π.*

In this notation, we are aiming to show that for some $\delta > 0$ and $b \ll n$, $(\delta, b) \in C_n(r; 1/2, o(1))$ for sufficiently large n. We prove the somewhat stronger statement that $(\delta, b) \in C_n(r; o(1))$.

The basis case of the induction on r is provided by the following result for one-round protocols.

LEMMA 11. *Let $n \in \mathbb{N}$ and $\gamma \in (0, \frac{1}{2})$ and $b \leq n$, and assume $\gamma b \geq 400n/\log n$. Then $\langle \delta(n, b, \gamma), b \rangle \in C_n(1; \gamma)$, where*

$$\delta(n, b, \gamma) = \frac{1}{2} \left(\frac{b}{4n} \right)^{2^{\frac{80n}{b\gamma}}} .$$

The induction step is provided by the following lemma.

LEMMA 12. *Fix n. If $\langle \delta_1, b_1 \rangle \in C_n(r_1; \gamma_1)$ and $\langle \delta_2, b_2 \rangle \in C_n(r_2; \gamma_2)$, then*

$$\left\langle \frac{\delta_1 \delta_2}{2}, b_1 + b_2 \right\rangle \in C_n \left(r_1 + r_2; \frac{2\gamma_1}{\delta_2} + \gamma_2 \right) .$$

These two lemmas are combined to prove the following lemma.

LEMMA 13. *Let $b \leq n \in \mathbb{N}$ and $\gamma \in (0, \frac{1}{2})$. Define $\lambda_0 = \frac{1}{2}$ and for $r \geq 1$ define*

$$\lambda_r = 4\lambda_{r-1} \left(\frac{4n}{b} \right)^{2^{\frac{160n}{b\gamma}\lambda_{r-1}}} .$$

Then for all $r \geq 1$ such that $\gamma b \geq 800\lambda_{r-1} n/\log n$,

$$\left\langle \frac{1}{\lambda_r}, rb \right\rangle \in C_n(r; r\gamma).$$

An immediate consequence of this lemma is the following corollary.

COROLLARY 14. *Let n, b, γ, and λ_i be as in Lemma* 13 *and suppose that r is an integer such that*

$$\frac{\gamma b \log n}{800n} \geq \lambda_{r-1}.$$

Then if Π is an (n,r)-protocol that is $r\gamma$-nontrivial, there exists at least one subset B of size rb that is $r\gamma$-powerful.

We first deduce the main theorem from this corollary.

Proof of Theorem 3. We prove the second part of the theorem first.

Let $r(n)$ be an integer valued function with $r(n) \leq (1/2 - \epsilon)\log^* n$ and let $\Pi = (\Pi_n : n \geq 1)$ be a sequence where Π_n is an $(n, r(n))$-protocol. Let

$$b(n) = \frac{(r(n))^2 n}{\log^{(2r(n)-1)} n} a(n),$$

where $a(n)$ is any function tending to infinity. Let n be sufficiently large, and suppose for contradiction that for some $\gamma > 0$, Π is $(b(n), \gamma)$-resilient. Without loss of generality we may assume that $p^1_{\Pi_n} \geq 1/2$. Let $b'(n) = b(n)/r(n)$ and

$$\gamma' = \gamma'(n) = \frac{\gamma}{2r(n)}.$$

By the previous corollary applied to b' and γ', if

$$\frac{\gamma a(n) \log n}{1600 \log^{(2r-1)} n} \geq \lambda_{r-1}, \tag{3.1}$$

then there is at least one subset of size $r(n)b'(n) = b(n)$ that is $(r(n)\gamma'(n) = \gamma/2)$-powerful, which would contradict our assumption. So it suffices to show that inequality (3.1) hold When $r = 1$ inequality (3.1) holds, for large enough n, by inspection. Otherwise, taking $\log^{(2r-2)}$ of both sides, the left-hand side is at least $\frac{1}{2}\log^{(2r-1)} n$ for large enough n and so it suffices to show that this is an upper bound on $\log^{(2r-2)} \lambda_{r-1}$. In the following proposition, T denotes the tower function, as defined in section 2.1.

PROPOSITION 15. *Let $b \leq n$ and $\gamma \in (0,1)$. For all integers $r \geq 0$, $\lambda_r \leq \kappa_r$ where*

$$\kappa_r = \frac{b\gamma}{320n} \mathrm{T}\left(2r; \frac{640n}{b\gamma}\right).$$

Proof. κ_r satisfies the recurrence

$$\kappa_0 = 2,$$
$$\kappa_r = \left(\frac{b\gamma}{320n}\right) 2^{2^{\frac{320n}{b\gamma}\kappa_{r-1}}},$$

so it suffices to show that

$$\lambda_r \leq \left(\frac{b\gamma}{320n}\right) 2^{2^{\frac{320n}{b\gamma}\lambda_{r-1}}},$$

which follows from

$$\begin{aligned}
\lambda_r &= 4\lambda_{r-1}\left(\frac{4n}{b}\right)^{2^{\frac{160n}{b\gamma}\lambda_{r-1}}} \\
&= \left(\frac{b\gamma}{320n}\right)\left(\frac{1280n\lambda_{r-1}}{b\gamma}\right)\left(\frac{4n}{b}\right)^{2^{\frac{160n}{b\gamma}\lambda_{r-1}}} \\
&\le \left(\frac{b\gamma}{320n}\right)\left(\frac{5120n^2\lambda_{r-1}}{b^2\gamma}\right)^{2^{\frac{160n}{b\gamma}\lambda_{r-1}}} \\
&= \left(\frac{b\gamma}{320n}\right)2^{\log\left(\frac{5120n^2\lambda_{r-1}}{b^2\gamma}\right)2^{\frac{160n}{b\gamma}\lambda_{r-1}}} \\
&\le \left(\frac{b\gamma}{320n}\right)2^{2^{\frac{320n}{b\gamma}\lambda_{r-1}}}. \qquad \square
\end{aligned}$$

Using the proposition, and the assumption about b, for n sufficiently large we have

$$\log^{(2r-2)}\lambda_{r-1} \le \frac{640n}{b\gamma} \le \frac{640}{a(n)\gamma}\log^{(2r(n)-1)} n < \frac{1}{2}\log^{(2r(n)-1)} n,$$

as required to complete the proof of the second part of the theorem.

For the first part of the theorem, it suffices to note that if $r(n) \le \frac{1}{2}\log^* n - \Delta$ for $\Delta = \log^*\log^* n$, then $r(n)^2 = o(\log^{(2r(n)-1)} n)$. This follows by taking $\log^{(\Delta)}$ of both sides: $\log^{(\Delta)}(r(n)^2) \le 2$ while $\log^{(\Delta)}(\log^{(2r(n)-1)} n) \ge \mathrm{T}(\Delta; 2)$. Hence we can choose $b(n) = o(n)$ so that it satisfies the hypothesis and, hence, the conclusion of the second part of the theorem. $\square$

So it remains to prove Lemmas 11, 12, and 13.

3.1. Proof of Lemma 11. Let f be a γ-nontrivial function on n variables. We want to show that for b in the given range, a "large" fraction of the sets of size b are γ-powerful. In light of Proposition 6, we may assume that f is monotone.

Fix $\gamma \in (0, 1/2)$. We first describe a stochastic process for selecting a sequence of variables $v_1, v_2, \ldots, v_d$ for an integer d to be specified, and show that with probability at least 1/2, the process produces a set of variables that is γ-powerful. The process depends on a parameter s, which we will also specify later. Having selected the first k of these variables $v_1, \ldots, v_k$, let f_k denote the (monotone) Boolean function on $n - k$ variables obtained by setting each v_i to 1. We then select v_{k+1} as follows:

1. If there is a variable v whose influence in f_k is at least 2^{-s}, let v_{k+1} be such a variable of lowest index.
2. Otherwise, choose v_{k+1} uniformly at random from among the remaining $n-k$ variables.

We will establish the following claim.

Claim A. *Let n be sufficiently large and let $d \in [n]$ and $\gamma \in (0, 1/2)$, and suppose that $\gamma d \ge \frac{160n}{\log n}$. Let s, the parameter of the process, be $\frac{80n}{\gamma d}$. Then*

$$\Pr\Big[\{v_1, \ldots, v_d\} \textit{ is } \gamma\textit{-powerful in } f\Big] \ge 1/2.$$

Define random variables X_k and Z_k, for $i = 0, \ldots, d$, by

$$X_k = \begin{cases} 1 & \text{if } p^1_{f_{k-1}} \geq 1 - \gamma, \\ I^1_{f_{k-1}}(v_k) & \text{otherwise,} \end{cases}$$

$$Z_k = \sum_{i=1}^{k} X_i.$$

Claim A is easily deduced from the following two claims. In both claims, n, d, γ, and s are as in Claim A.

Claim B. *If* $Z_d \geq 1 - 2\gamma$, *then* $\{v_1, \ldots, v_d\}$ *is* γ*-powerful in* f.

Claim C. *Suppose* $3 \leq s \leq \log(20n) - \log\log(20n)$. *For each* $k = 1, \ldots, d$,

$$\mathbb{E}\Big[X_k \,\Big|\, X_0, \ldots, X_{k-1}\Big] \geq \frac{s\gamma}{20n}.$$

Assume Claims B and C. Let $s = 80n/\gamma d$. Since $\gamma d \geq 160n/\log n$, and $d \leq n$, s satisfies the hypothesis of Claim C and therefore

$$\mathbb{E}[Z_d] \geq \frac{ds\gamma}{20n} \geq 4.$$

Applying Lemma 9 with $\mu = 4$ and $\delta = 3/4$ gives

$$\Pr[Z_d < 1] \leq e^{-9/8};$$

now applying Claim B yields the conclusion of Claim A.

To prove Claim B, assume that $Z_d \geq 1 - 2\gamma$. It suffices to show that $p^1_{f_d} \geq 1 - \gamma$, since this is equivalent to $\{v_1, \ldots, v_d\}$ being γ-powerful. Since $p^1_{f_k}$ is nondecreasing in k, we may assume that $p^1_{f_k} < 1 - \gamma$ for $k < d$. Then, recalling the definition of Z_d,

$$Z_d = \sum_{k=1}^{d} I_{f_{k-1}(v_k)}.$$

Now for each $k \geq 1$, $p^1_{f_k} = p^1_{f_{k-1}} + I^1_{f_{k-1}}(v_k)$, and hence $p^1_{f_d} = p^1_f + Z_d$. Since $p^1_f \geq \gamma$ by hypothesis, $p^1_{f_d} \geq 1 - \gamma$, as required for Claim B.

Since $v_1, v_2, \ldots, v_{k-1}$ determine $X_0, \ldots, X_{k-1}$, Claim C follows if we show

$$\mathbb{E}[X_k | v_1, \ldots, v_{k-1}] \geq \frac{s\gamma}{20n}.$$

If $p_{f_{k-1}} \geq 1 - \gamma$, then X_k is identically 1. Otherwise, $X_k = I_{f_{k-1}}(v_k)$. If v_k was selected by rule 1, then $X_k \geq 2^{-s}$, which is at least $s\gamma/(20n)$ for $s \leq \log(20n) - \log\log(20n)$. If rule 2 was used to select v_k, Lemma 7 gives the desired conclusion. This establishes Claim C and thus Claim A.

We now complete the proof of Lemma 11. The idea is that there are few variables chosen by rule 1, so with nonnegligible probability a random set of variables will contain them all.

More specifically, the hypothesis of the lemma implies that if we set $d = \lfloor b/2 \rfloor$, then d and γ satisfy the hypothesis of Claim A. We will use Claim A to show that a random subset of size b is γ-powerful with the required probability. Choose $s = 80n/(\gamma d)$ in accordance with Claim A, and observe that $s \leq \frac{1}{2}\log n$.

First, we reformulate the selection process for $v_1, v_2, \ldots, v_d$ in such a way that all random selections are made at the beginning of the process. We select a pair (S, σ), where S is a set of d variables chosen uniformly at random and σ is a bijection from $[d]$ to S chosen uniformly at random from the $d!$ such maps. Then $\vec{v}(S, \sigma) = (v_1, \ldots, v_d)$ is selected as above, except that rule 2 is replaced by "let i be the least integer such that $\sigma(i)$ is not a member of $\{v_1, \ldots, v_{k-1}\}$ and set $v_k = \sigma(i)$." It is easy to see that this process generates the same distribution over sequences $v_1, \ldots, v_d$ as the original process. Let $R_1(S, \sigma)$ be the set of v_i's chosen according to rule 1, and let $R_2(S, \sigma)$ be the set of those v_i's chosen according to rule 2. Obviously, $R_2(S, \sigma) \subseteq S$. Also, $|R_1(S, \sigma)| \leq 2^s$, since $I_{f_{k-1}}(v_k) \geq 2^{-s}$ if $v_k \in R_1(S, \sigma)$ and $\sum_{k=1}^{d} I_{f_{k-1}}(v_k) \leq p^1_{f_d} \leq 1$.

Now, consider a randomly chosen set B of size b. We view the probability space for B as consisting of triples (S, σ, T), where S and σ are as above and T is a random subset of size $b - d$ of $[n] \setminus S$. The set B is $S \cup T$. For B to be γ-powerful, it suffices that (i) $\vec{v}(S, \sigma)$ is γ-powerful and (ii) $R_1(S, \sigma) - S \subseteq T$, since then B contains the γ-powerful set $R_1(S, \sigma) \cup R_2(S, \sigma)$. By Claim A, event (i) occurs with probability at least $1/2$. Now, for any $S_0 \subset [n]$ of size b, we may estimate the probability of event (ii) conditioned on $S = S_0$: since $R_1(S, \sigma) - S$ has size at most $2^s = 2^{80n/\gamma d} \leq b/4$, and T is a random subset of $[n] \setminus S$ of size $b - d \geq b/2$, the probability that T contains $R_1(S, \sigma) - S$ is at least

$$\left(\frac{b-d}{n-d}\right)\left(\frac{b-d-1}{n-d-1}\right)\cdots\left(\frac{b-d-\lfloor 2^s \rfloor+1}{n-d-\lfloor 2^s \rfloor+1}\right) \geq \left(\frac{b}{4n}\right)^{2^s} \geq \left(\frac{b}{4n}\right)^{2^{80n/\gamma d}}.$$

Then $\Pr[\text{event (ii)} \mid \text{event (i)}] \geq \left(\frac{b}{4n}\right)^{2^{80n/\gamma d}}$, from which follows the statement of the lemma.

This completes the proof of Lemma 11.

3.2. Proof of Lemma 12. We first give a modified (and slightly more general) formulation of the lemma which will make the exposition a bit clearer.

LEMMA 16. *Fix n. If $\langle \delta_1, b_1 \rangle \in C_n(r_1; \frac{\alpha_1 \delta_2}{2}, \beta_1)$ and $\langle \delta_2, b_2 \rangle \in C_n(r_2; \alpha_2, \beta_2)$, then*

$$\left\langle \frac{\delta_1 \delta_2}{2}, b_1 + b_2 \right\rangle \in C_n(r_1 + r_2; \alpha_1 + \alpha_2, \beta_1 + \beta_2).$$

To deduce Lemma 12 from this, suppose that δ_1, b_1, r_1, γ_1, δ_2, b_2, r_2, and γ_2 are given satisfying the hypotheses of Lemma 12. Apply the above lemma with the same δ_i, b_i, and r_i, and with $\alpha_1 = 2\gamma_1/\delta_2$, $\beta_1 = \gamma_1$, and $\alpha_2 = \beta_2 = \gamma_2$.

So we prove Lemma 16.

Proof. Let $\Pi : (\{0,1\}^n)^{r_1 + r_2} \to \{0,1\}$ be an $(n, r_1 + r_2)$-protocol with $p^1_\Pi \geq \alpha_1 + \alpha_2$. We want to lower bound the probability that a (uniformly) random subset B of $[n]$ of size $b_1 + b_2$ is $\beta_1 + \beta_2$-powerful in Π.

A random subset B of size $b_1 + b_2$ can be selected by selecting subsets B_1, B_2, C, where B_1 is a uniformly random subset of size b_1, B_2 is a uniformly random subset of size b_2, and C is a uniformly random subset of $n - (B_1 \cup B_2)$ of size $b_1 + b_2 - |B_1 \cup B_2|$. Clearly, the probability that $B = B_1 \cup B_2 \cup C$ is $\beta_1 + \beta_2$-powerful is at least the probability that $B_1 \cup B_2$ is $\beta_1 + \beta_2$-powerful, so we lower bound this latter probability.

To do this, we define an event V that implies that $B_1 \cup B_2$ is $\beta_1 + \beta_2$-powerful and such that $\Pr_{B_1, B_2}[V]$ can be analyzed.

The input to Π is a vector in $(\{0,1\}^n)^{r_1 + r_2}$. Fixing the outcome of the first r_1 rounds to $\vec{\sigma} \in (\{0,1\}^n)^{r_1}$ gives rise to an (n, r_2)-protocol $\Pi[\vec{\sigma}] : (\{0,1\}^n)^{r_2} \to \{0,1\}$

by assigning

$$\Pi[\vec{\sigma}](\vec{\tau}) = \Pi(\sigma^1, \ldots, \sigma^{r_1}, \tau^1, \ldots, \tau^{r_2}).$$

Then $p^1_{\Pi[\vec{\sigma}]}$ can be viewed as a function of $\vec{\sigma}$. Let $\mathcal{E}$ be the set of those $\vec{\sigma}$ for which $p^1_{\Pi[\vec{\sigma}]} \geq \alpha_2$.

For $B_2 \subset [n]$, let $\mathcal{E}_{B_2}$ be the set of all $\vec{\sigma} \in \mathcal{E}$ such that B_2 is β_2-powerful with respect to the protocol $\Pi[\vec{\sigma}]$, i.e.,

$$\mathcal{E}_{B_2} = \left\{ \vec{\sigma} \in \mathcal{E} \ : \ p^1_{\Pi[\vec{\sigma}]}(B_2) \geq 1 - \beta_2 \right\}.$$

For each $B_2 \subseteq [n]$ of size b_2, let $\hat{\Pi}_{B_2}$ be the (n, r_1)-protocol $\hat{\Pi} = \hat{\Pi}_{B_2}$ given by

$$\hat{\Pi}_{B_2}(\vec{\sigma}) = \begin{cases} 1 & \text{if } \vec{\sigma} \in \mathcal{E}_{B_2}, \\ 0 & \text{otherwise.} \end{cases}$$

We now define V to be the event (depending on B_1 and B_2) that B_1 is β_1-powerful in $\hat{\Pi}_{B_2}$.

First we show that V implies that $B_1 \cup B_2$ is $\beta_1 + \beta_2$-powerful in Π. Consider the following two-step strategy for $B_1 \cup B_2$: (i) For the first r_1 rounds, B_1 plays so as to maximize the probability that $\vec{\sigma} \in \mathcal{E}_{B_2}$. Assuming this is successful then (ii) during the next r_2 rounds, B_2 tries to force the outcome of Π to be 1. The probability that this strategy fails is at most the sum of the probability that (i) fails and that (ii) fails given that (i) succeeds. The probability that (i) fails is at most β_1 by the definition of V. Assuming that (i) succeeds, the probability that (ii) fails is at most β_2 by the definition of the relation $\mathcal{E}_{B_2}$. Thus, given V, $B_1 \cup B_2$ is $\beta_1 + \beta_2$-powerful.

It remains to show that $\Pr[V] \geq \delta_1\delta_2/2$. To do this we consider, for $\eta > 0$, the event U_η (depending on B_2 alone) that $\Pr_{\vec{\sigma}}[\vec{\sigma} \in \mathcal{E}_{B_2}] \geq \eta$. We will show that when $\eta = \frac{\alpha_1\delta_2}{2}$, $\Pr[U_\eta] \geq \delta_2/2$ and $\Pr[V|U_\eta] \geq \delta_1$, which immediately gives the desired lower bound on $\Pr[V]$.

First we lower bound $\Pr[U_\eta]$. For fixed B_2 we have

$$\Pr_{\vec{\sigma}}[\vec{\sigma} \in \mathcal{E}_{B_2}] = \Pr_{\vec{\sigma}}[\vec{\sigma} \in \mathcal{E}] \times \frac{|\mathcal{E}_{B_2}|}{|\mathcal{E}|}. \tag{3.2}$$

By the definition of $\mathcal{E}$,

$$\mathop{\mathbb{E}}_{\vec{\sigma}}[p^1_{\Pi[\vec{\sigma}]}] \leq \Pr_{\vec{\sigma}}[\vec{\sigma} \in \mathcal{E}] + (1 - \Pr_{\vec{\sigma}}[\vec{\sigma} \in \mathcal{E}])\alpha_2 \leq \Pr_{\vec{\sigma}}[\vec{\sigma} \in \mathcal{E}] + \alpha_2.$$

We also have $\mathbb{E}_{\vec{\sigma}}[p^1_{\Pi[\vec{\sigma}]}] = p^1_\Pi \geq \alpha_1 + \alpha_2$, and thus

$$\Pr_{\vec{\sigma}}[\vec{\sigma} \in \mathcal{E}] \geq \alpha_1. \tag{3.3}$$

Letting $W = W(B_2)$ denote the random variable $|\mathcal{E}_{B_2}|/|\mathcal{E}|$ and combining (3.3) and (3.2), we have

$$\Pr_{B_2}[U_\eta] \geq \Pr_{B_2}[W \geq \eta/\alpha_1].$$

So we lower bound this latter probability. For $\sigma \in \mathcal{E}$, the protocol $\Pi[\vec{\sigma}]$ is an (n, r_2) protocol that is α_2-nontrivial. Thus, by the hypothesis of the lemma, for any $\vec{\sigma} \in \mathcal{E}$,

$$\Pr_{\substack{B_2 \subset [n] \\ |B_2| = b_2}} \left[\vec{\sigma} \in \mathcal{E}_{B_2} \right] \geq \delta_2.$$

Summing over $\sigma \in \mathcal{E}$ and dividing by $|\mathcal{E}|$ we obtain $\mathbb{E}_{B_2}[W] \geq \delta_2$. Since $W \in [0,1]$, we also have $\mathbb{E}_{B_2}[W] \leq \Pr_{B_2}[W \geq \eta/\alpha_1] + \eta/\alpha_1$, which implies $\Pr_{B_2}[W \geq \eta/\alpha_1] \geq \delta_2 - \eta/\alpha_1$. Setting $\eta = \alpha_1\delta_2/2$ we have $\Pr_{B_2}[U_{\alpha_1\delta_2/2}] \geq \Pr_{B_2}[W \geq \delta_2/2] \geq \delta_2/2$ as required.

Finally, we lower bound $\Pr[V|U_\eta]$. V is the event that B_1 is β_1-powerful in $\hat{\Pi}_{B_2}$. The event U_η implies that the protocol $\hat{\Pi}_{B_2}$ is η-nontrivial, and for $\eta = \alpha_1\delta_2/2$, the hypothesis of the lemma implies that the probability that V occurs is at least δ_1. ◻

3.3. Proof of Lemma 13. Fix b, n, and γ as hypothesized. Let $H(r)$ denote the hypothesis $\gamma b \geq 800n\lambda_{r-1}/\log n$, and let $C(r)$ denote the conclusion

$$\left\langle \frac{1}{\lambda_r}, rb \right\rangle \in C_n(r; r\gamma).$$

We want to show that $H(r)$ implies $C(r)$ for all $r \geq 1$. We proceed by induction on r.

The basis case is immediate from Lemma 11. For the induction step, let $r \geq 1$, and suppose that $H(r)$ implies $C(r)$. Assume $H(r+1)$ is true; we want to show $C(r+1)$ holds. Now $H(r+1)$ implies $H(r)$ since $\lambda_r \geq \lambda_{r-1}$, and hence $C(r)$ holds. If $\gamma' \in (0, 1/2)$ is such that $\gamma' b \geq 400n/\log n$, then Lemma 11 implies

$$\langle \delta(n, b, \gamma'), b \rangle \in C_n(1; \gamma').$$

Combining this and $C(r)$ using Lemma 12, and setting $\gamma' = \frac{\gamma}{2\lambda_r}$, gives

$$\left\langle \frac{\delta(n, b, \gamma')}{2\lambda_r}, (r+1)b \right\rangle \in C_n(r+1, (r+1)\gamma),$$

which is equivalent to $C(r+1)$.

This completes the proof of Lemma 13 and the proof of the main theorem.

4. Extensions to protocols with longer messages. We now indicate how to generalize the bounds proven above to protocols which permit players to send longer messages. Recall that for $n, r \in \mathbb{N}$ and $\vec{\ell} = (\ell_1, \ldots, \ell_r) \in \mathbb{N}^r$, we say that Π is a $(n, r, \vec{\ell})$-protocol if n is the number of players, r is the number of rounds, and no more than ℓ_k bits are broadcast by each player in the kth round.

We extend Definition 10 to account for variable message lengths.

DEFINITION 17. *Let $C_n^{\vec{\ell}}(r; \alpha, \beta)$ (written $C_n^{\vec{\ell}}(r; \gamma)$ when $\alpha = \beta = \gamma$) denote the collection of pairs $\langle \delta, b \rangle$ so that for any $(n, r, \vec{\ell})$-protocol Π that is α-nontrivial, at least a δ fraction of sets $B \subset [n]$ of size b are β-powerful in Π.*

We begin by considering a single-round protocol $f : (\{0,1\}^\ell)^n \to \{0,1\}$ in which each player broadcasts ℓ bits. Simply treating f as a function on $n\ell$ Boolean variables and examining the stochastic process of Section 3.1 yields the following version of Claim A.

Claim D (cf. Claim A). *Let $n\ell$ be sufficiently large and let $d \in [n]$ and $\gamma \in (0,1/2)$, and suppose that $\gamma d \geq \frac{160n\ell}{\log(n\ell)}$. Let s, the parameter of the process, be $\frac{80n\ell}{\gamma d}$. Then*

$$\Pr\Big[\{v_1,\ldots,v_d\} \textit{ is } \gamma\textit{-powerful in } f\Big] \geq 1/2.$$

If the Boolean variables $\{v_1,\ldots,v_d\}$ are γ-powerful in $f:\{0,1\}^{n\ell} \to \{0,1\}$, then the $\{0,1\}^\ell$-valued variables $\{x \mid \exists i, v_i \text{ is a component of } x\}$ are γ-powerful in f, again viewed as a function on $(\{0,1\}^\ell)^n$. Observe that applying Claim A in this way does not exploit the fact that each player controls many bits of the function f. The proof of Lemma 11 now yields the following lemma.

Lemma 18 (cf. Lemma 11). *Let $n,\ell \in \mathbb{N}$ and $\gamma \in (0,\frac{1}{2})$ and $b \leq n$, and assume $\gamma b \geq 400n\ell/\log(n\ell)$. Then $\langle \delta, b\rangle \in C_n^{(\ell)}(1;\gamma)$, where*

$$\delta = \delta(n,b,\ell,\gamma) = \frac{1}{2}\left(\frac{b}{4n}\right)^{2^{\frac{80n\ell}{\gamma b}}}.$$

The number of bits broadcast per round is immaterial to the proof of Lemma 12; restating that lemma for multibit protocols yields the following lemma.

Lemma 19 (cf. Lemma 12). *Fix n. If $\langle \delta_1, b_1\rangle \in C_n^{\vec{\ell}}(r_1;\gamma_1)$ and $\langle \delta_2, b_2\rangle \in C_n^{\vec{m}}(r_2;\gamma_2)$, then*

$$\left\langle \frac{\delta_1\delta_2}{2}, b_1+b_2\right\rangle \in C_n^{(\vec{\ell},\vec{m})}\left(r_1+r_2; \frac{2\gamma_1}{\delta_2}+\gamma_2\right),$$

where $(\vec{\ell},\vec{m})$ denotes the vector $(\ell_1,\ldots,\ell_{r_1},m_1,\ldots,m_{r_2})$.

We combine these to prove the following lemma.

Lemma 20 (cf. Lemma 13). *Let $b \leq n$, $\gamma \in (0,1/2)$, and $l_i \in \{1,2,\ldots\}$ for each $i \geq 0$. Define $\lambda_0 = \frac{1}{2}$, and for $r \geq 1$ define*

$$\lambda_r = 4\lambda_{r-1}\left(\frac{4n}{b}\right)^{2^{\frac{160nl_{r-1}}{\gamma b}\lambda_{r-1}}}.$$

Assume that for each $r \geq 1$, $\lambda_r l_r \geq \lambda_{r-1}l_{r-1}$. Then, if $\gamma b \geq 800nl_{r-1}\lambda_{r-1}/\log n$,

$$\left\langle \frac{1}{\lambda_r}, rb\right\rangle \in C_n^{\vec{\ell}}(r;r\gamma),$$

where $\ell_i = l_{r-i}$, so $\vec{\ell} = (\ell_1,\ldots,\ell_r) = (l_{r-1},\ldots,l_0)$.

Proof. Fix b,n,γ, and l_i as hypothesized. Let $H(r)$ denote the hypothesis $\gamma b \geq 800nl_{r-1}\lambda_{r-1}/\log n$, and let $C(r)$ denote the conclusion

$$\left\langle \frac{1}{\lambda_r}, rb\right\rangle \in C_n^{\vec{\ell}}(r;r\gamma),$$

where $\vec{\ell} = (l_{r-1},\ldots,l_0)$. We want to show that $H(r)$ implies $C(r)$ for all $r \geq 1$. We proceed by induction on r.

The basis case is immediate from Lemma 18. For the induction step, let $r \geq 1$, and suppose that $H(r)$ implies $C(r)$. Assume $H(r+1)$ is true; we want to show that

$C(r+1)$ holds. Now $H(r+1)$ implies $H(r)$ since, by assumption, $\lambda_r l_r \geq \lambda_{r-1} l_{r-1}$, and hence $C(r)$ holds. If $\gamma' \in (0, 1/2)$ is such that $\gamma' b \geq 400 n l_r / \log n$, then Lemma 18 implies that

$$\langle \delta(n, b, l_r, \gamma')), b \rangle \in C_n^{(l_r)}(1; \gamma').$$

Combining this and $C(r)$ using Lemma 19, and setting $\gamma' = \frac{\gamma}{2\lambda_r}$, gives

$$\left\langle \frac{\delta(n, b, l_r, \gamma')}{2\lambda_r}, (r+1)b \right\rangle \in C_n^{\vec{\ell}}(r+1, (r+1)\gamma),$$

where $\vec{\ell} = (l_r, \ldots, l_0)$, which is equivalent to $C(r+1)$. $\square$

This may be applied to prove Theorem 4.

Proof of Theorem 4. Fix n. Set $\alpha = \frac{1}{\log^* n}$ and define $\gamma = \alpha^2$, $b = \lceil \alpha^2 n \rceil$, and, for $i \in \{0, \ldots, r-1\}$,

$$l_i = \max\left(1, \left\lfloor \frac{\alpha\gamma b (\log^{(2(r-i)-1)} n)^{1-\alpha}}{800n} \right\rfloor \right). \tag{4.1}$$

Note that

$$\begin{aligned} l_o &\geq \frac{\alpha^5 (\log^{(2r-1)} n)^{1-\alpha}}{800} = \frac{(\log^{(\log^* n - 2\log^* \log^* n - 1)} n)^{1-\alpha}}{800(\log^* n)^5} \\ &= \frac{(T(2\log^* \log^* n - 1; 1))^{1-o(1)}}{800(\log^* n)^5} = (\log^* n)^{\omega(1)} \end{aligned}$$

so that, when n is sufficiently large, $\gamma < \frac{1}{2}$ and $l_{r-1} \geq \cdots \geq l_0 > 1$. In this case, with λ_i defined as in Lemma 20,

$$\begin{aligned} \lambda_i = 4\lambda_{i-1} \left(\frac{4n}{b}\right)^{2^{\frac{160 n l_{i-1}\lambda_{i-1}}{b\gamma}}} &= 2^{\log 4 + \log \lambda_{i-1} + 2^{\left(\frac{160 n l_{i-1}\lambda_{i-1}}{b\gamma} + \log\log \frac{4n}{b}\right)}} \\ &\leq 2^{2^{\frac{160 n l_{i-1}\lambda_{i-1}}{b\gamma} + \log\log 4 + \log\log\lambda_{i-1} + \log\log\frac{4n}{b}}} \end{aligned}$$

and, as $\max(\log\log 4, \log\log\lambda_{i-1}, \log\log(4n/b)) < 160 n l_{i-1}\lambda_{i-1}/\gamma b$,

$$\lambda_i \leq 2^{2^{\frac{640 n l_{i-1}\lambda_{i-1}}{b\gamma}}} \leq 2^{2^{\alpha(\log^{(2(r-i)+1)} n)^{1-\alpha}\lambda_{i-1}}}.$$

As noted above, these l_i are monotonically increasing (in i) and therefore satisfy the hypothesis of Lemma 20. We show that for sufficiently large n, $\lambda_i \leq (\log^{(2(r-i)-1)} n)^\alpha$. Since this is clearly true for λ_0, by induction

$$\lambda_i \leq 2^{2^{\alpha(\log^{(2(r-i)+1)} n)^{1-\alpha}\lambda_{i-1}}} \leq 2^{2^{\alpha(\log^{(2(r-i)+1)} n)}} \leq (\log^{(2(r-i)-1)} n)^\alpha, \tag{4.2}$$

where we have applied the inequality $x^\epsilon \leq \epsilon x$, valid when, for example, $x \geq 4$ and $\epsilon \in [1/\sqrt{x}, 1]$. (We apply the inequality with $\epsilon = \alpha$ and $x = \log^{(2(r-i))} n$; both these requirements are met for sufficiently large n.)

Finally, from (4.1) and (4.2) above,

$$\frac{800 n l_{r-1}\lambda_{r-1}}{\log n} \leq \alpha\gamma b \leq \gamma b,$$

so that Lemma 20 applies. This asserts the existence of an $r\gamma = o(1)$-powerful set of $rb = o(n)$ players for any protocol Π under the following assumptions:

- Π is $r\gamma = o(1)$-nontrivial,
- Π lasts for r rounds, with $r \leq \frac{1}{2}\log^* n - \log^* \log^* n$, and
- Π calls for no more than

$$l_{r-k} = \Omega\left(\frac{(\log^{(2k-1)} n)^{(1-\alpha)}}{\mathrm{poly}(\log^* n)}\right) = (\log^{(2k-1)} n)^{(1-O(\alpha))}$$

communication bits in the kth round. $\square$

5. The influence of large coalitions. Applying the results of [13], one can show that, for a Boolean function f with $p_f^1 = 1/2$ and $b(n) = \Theta(n)$, there is always a coalition L of size $b(n)$ for which $p_f^0(L) \geq 1 - 1/n^c$ for some appropriate constant c (depending on b). When $b(n) \geq n/2$, however, the following observation from [17] may be applied.

PROPOSITION 21. *Let X be a finite probability space and $f : X^n \to \{0,1\}$. Let $A_1, A_2 \subset [n]$ be a partition of the variables on which f is defined (so that $A_1 \cup A_2 = [n]$ and $A_1 \cap A_2 = \emptyset$). Then for at least one of these two sets, A_i,*

$$p_f^1(A_i) = 1 \qquad \text{or} \qquad p_f^0(A_i) = 0.$$

Below we observe that near this $\frac{n}{2}$ threshold (specifically, for $b(n) > (1/3 + \epsilon)n$), the above bound of [13] may be improved to $1 - 1/\exp(\Omega(n))$.

In preparation for the lemma, we record a Chernoff bound (see, e.g., [3]).

LEMMA 22. *Let $X_i, i = 1, \ldots, n$, be independent random variables, each uniformly distributed in $\{0,1\}$. Then*

$$\Pr\left[\sum_i X_i - \frac{n}{2} > a\right] < \exp\left(-\frac{a^2}{2n}\right).$$

THEOREM 23. *Let $\gamma > \frac{1}{3}$. Let $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function and let $\mathfrak{B} = \{B \subset [n] \;:\; |B| = \lceil \gamma n \rceil\}$. If $p_f^1(B) < 1$ for all $B \in \mathfrak{B}$, then for all $B \in \mathfrak{B}$, $p_f^0(B) \geq 1 - \epsilon$, where*

$$\epsilon = \exp\left(-\frac{(1-3\gamma)^2}{8(1-\gamma)} n\right).$$

Proof. Assume that f is monotone. Recall that a min-term of a monotone function f is a minimal subset of variables which, if set to 1, ensures that $f = 1$. If f has a min-term of cardinality at most γn, then clearly there is $B \in \mathfrak{B}$ for which $p_f^1(B) = 1$. Otherwise all min-terms have cardinality larger than γn. Fix $B \in \mathfrak{B}$ and consider an input $\vec{x} = x_1 \ldots x_n$, where each x_i, for $i \notin B$, is chosen independently at random in $\{0,1\}$, and $x_i = 0$ for $i \in B$. Then $\mathbb{E}[\sum_i x_i] \leq \frac{1-\gamma}{2} n$, so that by applying the above Chernoff bound,

$$\Pr\left[\sum_i x_i > \gamma n\right] < \exp\left(-\frac{(3\gamma-1)^2}{8(1-\gamma)} n\right).$$

Then $p_f^0(B) > 1 - \exp\left(-\frac{(3\gamma-1)^2}{8(1-\gamma)} n\right)$, as desired. $\square$

6. Open problems. We summarize the known results concerning protocols that are resilient against a linear number of corrupt players:

1. By [16] there is an $(n, \big[1+o(1)\big]\log n, 1)$-protocol which is $\Omega(n)$-resilient. By Theorem 3, there is no $(n, (1/2-\epsilon)\log^* n, 1)$-protocol that is $\Theta(n)$-resilient.
2. By [16], there is an $(n, \log^* n + O(1), \vec{\ell})$-protocol, where $\ell_k = O(\log^{(k)} n)$, that is $\Omega(n)$-resilient. By Theorem 4, there is no $(n, (1/2-o(1))\log^* n, \vec{\ell})$-protocol that is $\Theta(n)$-resilient for some $\ell_k = (\log^{(2k-1)} n)^{1-o(1)}$.
3. It is not difficult to show that Theorem 2 actually implies that there can be no $(n, 1, o(\log n))$-protocol that is $\Theta(n)$-resilient.

These suggest several avenues of investigation:

1. In the case where each player sends a single bit per round (item 1 above), $\big[1+o(1)\big]\log n$ rounds are sufficient to guarantee $\Omega(n)$-resilience, $\big[1/2 - o(1)\big]\log^* n$ rounds are necessary—what is the right answer?
2. In the general case (item 2 above), can Theorem 4 be strengthened to show any $\Omega(n)$-resilient protocol has some round k during which $\Omega(\log^{(k)} n)$ communication occurs?
3. From (3) above, no one-round protocol using $o(\log n)$ bits per player can be $\Omega(n)$-resilient. Even abandoning all constraints on the number of bits sent per round, is there a one (or even constant) round $\Omega(n)$-resilient protocol?
4. We have focused on protocols where honest players flip a fair coin; what can be said when the honest players' coin flips are biased?

Acknowledgments. We thank Nati Linial for pointing out the failure of iterative methods in the multibit case and several illuminating discussions. We also thank Uri Feige for suggesting our last open question and for useful discussions.

REFERENCES

[1] M. Ajtai and N. Linial, *The influence of large coalitions*, Combinatorica, 13 (1993), pp. 129–145.

[2] N. Alon and M. Naor, *Coin-flipping games immune against linear-sized coalitions*, SIAM J. Comput., 22 (1993), pp. 403–417.

[3] N. Alon and J. H. Spencer, *The Probabilistic Method*, John Wiley and Sons, New York, 1992.

[4] K. Azuma, *Weighted sums of certain dependent random variables*, Tôhoku Math. J. (2), 19 (1967), pp. 357–367.

[5] M. Ben-Or and N. Linial, *Collective coin flipping, robust voting schemes and minima of Banzhaf values*, in 26th Annual Symposium on Foundations of Computer Science, Portland, Oregon, 1985, IEEE, pp. 408–416.

[6] M. Ben-Or and N. Linial, *Collective coin flipping*, in Randomness and Computation, S. Micali, ed., Academic Press, New York, 1990, pp. 91–115.

[7] M. Ben-Or, N. Linial, and M. Saks, *Collective coin flipping and other models of imperfect randomness*, in Proceedings of the Seventh Hungarian Colloquium on Combinatorics, Colloq. Math. Soc. János Bolyai 52, A. Hajnal, L. Lovász, and V. Sós, eds., North-Holland, Amsterdam, 1988, pp. 75–112.

[8] H. Chernoff, *A measure of the asymptotic efficiency for tests of a hypothesis based on the sum of observations*, Ann. Math. Statistics, 23 (1952), pp. 493–507.

[9] U. Feige, *Noncryptographic selection protocols*, in 40th Annual Symposium on Foundations of Computer Science, 1999, IEEE, pp. 142–152.

[10] E. Friedgut and G. Kalai, *Every monotone graph property has a sharp threshold*, Proc. Amer. Math. Soc., 124 (1996), pp. 2993–3002.

[11] G. H. Hardy, J. E. Littlewood, and G. Pólya, *Inequalities*, 2nd ed., Cambridge University Press, Cambridge, UK, 1952.

[12] W. Hoeffding, *Probability inequalities for sums of bounded random variables*, J. Amer. Statist. Assoc., 58 (1963), pp. 13–30.

[13] J. Kahn, G. Kalai, and N. Linial, *The influence of variables on Boolean functions (extended abstract)*, in 29th Annual Symposium on Foundations of Computer Science, White Plains, NY, 1988, IEEE, pp. 68–80.

[14] N. Linial, *Game-theoretic aspects of computing*, in Handbook of Game Theory with Economic Applications, Vol. II, R. J. Aumann and S. Hart, eds., North-Holland, Amsterdam, 1994, pp. 1339–1395.

[15] R. Motwani and P. Raghavan, *Randomized Algorithms*, Cambridge University Press, Cambridge, UK, 1995.

[16] A. Russell and D. Zuckerman, *Perfect information leader election in* $\log^* n + O(1)$ *rounds*, J. Comput. System Sci., 63 (2001), pp. 612–626.

[17] M. Saks, *A robust noncryptographic protocol for collective coin flipping*, SIAM J. Discrete Math., 2 (1989), pp. 240–244.