

Simple and Efficient Leader Election in the Full Information Model*

Rafail Ostrovsky[†] Sridhar Rajagopalan[‡] Umesh Vazirani[§]

Abstract

In this paper, we study the leader election problem in the full information model. We show two results in this context. First, we exhibit a constructive $O(\log N)$ round protocol that is resilient against linear size coalitions. That is, our protocol is resilient against any coalition of size less than βN for some constant (but small) value of β . Second, we provide an easy, non-constructive probabilistic argument that shows the existence of $O(\log N)$ round protocol in which β can be made as large as $\frac{1}{2} - \epsilon$ for any positive ϵ . Our protocols are extremely simple.

*Preliminary version appeared in STOC-94

[†]Work done while at Computer Science Division, University of California at Berkeley, and International Computer Science Institute, Berkeley, CA 94720. Email: rafail@bellcore.com. Supported by an NSF Postdoctoral Fellowship and ICSI.

[‡]IBM, Almaden. Work done while at Computer Science Division, University of California at Berkeley, CA 94720. Email: sridhar@cs.Berkeley.EDU. Supported by NSF grant IRI 91-20074 and CCR-9310214.

[§]Computer Science Division, University of California at Berkeley, CA 94720. Email: vazirani@cs.Berkeley.EDU. Supported by NSF grant IRI 91-20074 and CCR-9310214.

1 Introduction

A central problem in distributed computing is that of electing a leader. We consider this problem in the setting where there is a fixed coalition of malicious players who are trying to maximize the chances of one of them being elected as a leader. A protocol, on the other hand must guarantee that there is a reasonable chance that an honest player will be elected regardless of any devious plot hatched by the coalition. This problem received much attention in the literature. (For example, as a good introduction, see surveys of Ben-Or, Linial and Saks [BLS-87], Chor, Dwork [CD-89] and Linial [L-92].) This problem is studied in principally two settings: the cryptographic setting and the information theoretic setting. In the cryptographic setting some cryptographic assumptions are made and the computational power of all players is limited. In the information theoretic setting, no such assumptions are made and faulty players are assumed to be capable of arbitrarily complicated computation.

Under cryptographic assumptions [GMW-87] or assuming private channels between players [BGW-88, CCD-88], it is possible to elect a leader in a constant number of rounds if less than $\frac{1}{2}$ (assuming a broadcast channel) (or less than $\frac{1}{3}$ without a broadcast [FM-88]) of the players are dishonest. Since this setting is not the focus of this paper, we will not review the rest of the extensive literature in this area. The interested reader is referred to the survey in Chor and Dwork [CD-89].

In this paper, we consider this question in the full-information model, introduced by Ben-or and Linial. This model is characterized by the following two “worst case” assumptions. First, that all communication is public and made via broadcast channels. That is, no two players share a private communication channel. The second, is that dishonest players are capable of arbitrarily complicated computations. Honest players however, are required to work in polynomial time. There are two parameters of interest in this setting: *resilience* and *efficiency*. Resilience deals with the number of dishonest players that a protocol can tolerate. Let us consider a system with N players. A protocol is β -immune if it can tolerate a coalition of size smaller than βN dishonest players, where dishonest players have been chosen before the protocol begins in a way which maximizes the probability of one of them being elected as a leader. Clearly, the objective here is to make β as large as possible. A great deal of effort has been invested in designing protocols which can tolerate a constant value of β . We measure *efficiency* in terms of the number of *rounds* of communication. In each round, every player can send a single message. Players can send messages during the same round in an asynchronous manner. That is, the notion of a round models the effects of distributed asynchrony, and the beginning of each round represents a synchronization point in the protocol. Clearly, minimizing the number of rounds is crucial since we minimize the number of times the system needs to be synchronized. (For more complete definitions see section 2.)

In the information-theoretic setting the problem was first formulated by Ben-Or and Linial [BL-85, BL-89]. In the first paper, Ben-Or and Linial propose a protocol that is $N^{(\log_3 2)-1}$ -immune. The second paper improved upon this result and exhibited a one-round $1/\log N$ -immune protocol. A widely known result due to Kahn, Kalai and Linial [KKL-88], implies that if each player were restricted to providing one random bit, then $1/\log N$ is the largest possible

value of β for any one round protocol. Since then, protocols were designed that were resilient against larger and larger coalitions [S-89, AL-89, AN-90, BN1-93, BN2-93]. In particular, Saks [S-89], Ajtai and Linial [AL-89] designed so-called baton-passing protocols which are $1/\log N$ -immune. Furthermore, Saks [S-89] also showed that no protocol can be immune against $\lceil \frac{N}{2} \rceil$ coalitions. Finally, Alon and Naor showed the existence of a protocol which is β -immune for any $\beta < \frac{1}{3}$. This constituted major progress since it established that a very large number of dishonest players can be tolerated. Further analysis of the protocol of Alon and Naor was done by Boppana and Narayanan where they showed that the protocol of Alon and Naor is β -immune for all $\beta < 0.44$ in [BN1-93] and for all $\beta < \frac{1}{2} - \epsilon$ and positive ϵ in [BN2-93], establishing that the protocol of Alon and Naor [AN-90] has, in fact, optimal resilience. Notice, however, that Alon and Naor solution requires linear number of rounds. A faster (in the number of rounds) solution was constructed by Cooper and Linial [CL-93], where they showed a fairly complicated protocol which requires $O(\log^{17} N)$ rounds and is c -immune for some constant $c \ll 1/2$ (in fact, c must be smaller than $\frac{1}{6e^6}$ in [CL-93] construction.)

In this paper we exhibit simple protocols that improve upon previous results in various ways:

- We first exhibit a constructive protocol which can tolerate a constant fraction βN (for any $\beta < .0045$) of malicious players and runs in $O(\log N)$ rounds. (This compares favorably to $O(\log^{17} N)$ constructive protocol of [CL-93] which also achieves linear resilience.)
- We then show, using an Erdős probabilistic argument, the existence of a protocol that also requires $O(\log N)$ rounds and can tolerate any $(\frac{1}{2} - \epsilon)N$ malicious players for all positive ϵ . (This compares favorably to the non-constructive protocol of [AN-90, BN2-93] which also achieves optimal resilience but runs in $O(N)$ number of rounds.)

It should be noted that in our protocol, we use a notion of a *committee*, which is a subset of players. The notion of a *committee* was introduced by Bracha [B-85] in the context of the Byzantine agreement problem, and has since become an important tool for designing resilient protocols. It has been used in many of the earlier papers on leader election in the full-information model.

Using the algorithm presented in the current paper, together with the novel use of Extractors, David Zuckerman was recently able to exhibit a constructive counterpart of our second (non-constructive) protocol and show how to implement it constructively with optimal $\beta < \frac{1}{2}$ resilience [Z-96].

2 Problem Statement

There are N players. All communication is via broadcast and is public. Every player has a unique name and the name of the sender of any message is explicitly known to all other players. The goal is to choose one of the N players (a leader) in this setting. This seemingly simple matter becomes complicated if some βN of the players are involved in a dishonest coalition.

This means that they do not necessarily adhere to any specified protocol. Further, given any protocol, they will behave in a way as to maximize the chances of one among themselves being elected leader. A protocol is said to be β -immune if regardless of which βN of the players are dishonest, the protocol chooses an honest leader with probability bounded away from 0 as N approaches infinity.

A *round* is a round of communication. It consists of a broadcast message of polynomial length transmitted by each player. In principle, all messages are to be transmitted and received in parallel. However, in a distributed system, there is no guarantee on the order in which messages are sent or delivered within a round. Thus, within any round, dishonest players may receive (or consciously wait until they receive), the messages of the honest players before they compose their own. The players are synchronized between rounds. Thus, all messages in round i are received before any in round $i + 1$ are transmitted. This models the effects of distributed asynchrony in the most pessimistic (and thus, the most difficult) case.

Finally, honest players have to work in polynomial time. There is however no such computational limitation on the dishonest ones. This stipulation rules out the use of cryptographic assumptions.

3 The constructive protocol

The protocol first constructs a collection $\mathcal{C}$ of committees. The committees are constructed by choosing an explicit constant degree d expander graph on $N + o(N)$ vertices (for example, see [GG-79]). Each player is represented by some vertex in the expander graph. Each committee is composed of the players represented by different vertices in a $k \log N$ long walk on the expander graph where k is a fixed constant. A player can be a member of a committee a multiple number of times. It is easily verified that there are only a polynomial number M of such walks¹.

We say that a committee is *dangerous* if the fraction of bad players in it exceeds $\beta + \delta$ for some fixed δ . By $\mathcal{B}$ we denote the set of all dangerous committees. Our choice of committees guarantees that the fraction of dangerous committees tends to zero rapidly as we increase the number of committees (by increasing k). There are various proofs of this kind of notion in the recent literature (see for instance [IZ-89]). We choose to state here a derivative of the theorem of Gillman [G-93]. Gillman provides a Chernoff style bound on the occupation time for each subset S of vertices during a random walk on any expander graph.

Lemma 3.0.1 It is possible to choose a collection $\mathcal{C}$ of committees such that:

- $|\mathcal{C}| = M \leq \text{poly}(N)$. The degree of the polynomial depends on k and the out degree of the expander graph.
- $|\mathcal{B}| \leq M^{1-\lambda\delta^2}$ for some fixed constant $\lambda < 1^2$.

¹The number of committees is smaller than $N \cdot d^{k \log N}$. The first factor represents the number of starting points and the second factor the number of walks from each starting point. Each walk is a committee.

²It should be noted that the value of λ depends on the gap between the largest and next largest eigenvalues

The protocol works in two stages. In the first stage, a committee $C \in \mathcal{C}$ will be chosen. In the second stage, a leader is chosen in C . The second stage is accomplished either non-constructively by using the [AN-90, BN2-93] sequential protocol or by using the first stage one more time recursively and then working on problem instance of size $O(\log \log n)$, for which explicit construction of [AN-90] protocol is derived by enumerating all possible [AN-90] protocols for $O(\log \log n)$ players and picking the optimal one. In either case, since $|C|$ is only logarithmic in N after the first stage, the second stage runs in time $O(\log N)$. Henceforth, we will concentrate on the first stage.

The first stage works in a series of rounds. At the beginning, all M committees are marked ‘eligible’. Let $\mathcal{E}_i$ be the collection of eligible committees at the inception of round i . In round i , each player publishes a list of size $\left\lceil \frac{|\mathcal{E}_i|}{N} \right\rceil$. Each entry in this list is the name of a committee in $\mathcal{E}_i$. In the case of honest players each entry is chosen independently and randomly from $\mathcal{E}_i$. Dishonest players may choose their lists in any arbitrary fashion. Every committee in the published lists are marked ‘not eligible’. The protocol terminates when there are fewer than N committees left in the eligible set. At this point player 1 chooses an arbitrary committee C from among the remaining eligible committees. If there is no such committee, the protocol fails. In section 4 we prove the following theorem pertaining to this protocol:

Theorem 4.1: The protocol described here is an $O(\log N)$ round leader election protocol that is β -immune for every β smaller than .0045.

4 The proof

In this section we will address the two issues that arise in proving the main theorem in this paper. We first address the easier issue, and show that the protocol runs in $O(\log N)$ rounds. Then, we address the main issue, that of showing that the protocol described here has the required immunity.

4.1 Rounds

Consider an arbitrary committee $C \in \mathcal{C}$. The probability that C is eliminated in an arbitrary round i is at least the probability that C is eliminated by one of the votes cast by the honest players. There are at least $(1 - \beta)|\mathcal{E}_i|$ votes cast by honest players. Each vote eliminates C

for the expander graph. The theorem of Gillman implies that λ can be taken to be roughly $\frac{1}{20}$ th the eigenvalue gap. Expanders that have an eigenvalue gap arbitrarily close to 1 are known. We note that Gillman’s bound is somewhat weaker than the standard Chernoff bound in the sense that λ should ideally be close to 1/2. An improvement in this yields an increase in the value of β in our main theorem (4.1). In [K-94] Nabil Kahale have shown small improvement, increasing β somewhat, but still not achieving optimal 1/2. Optimal bound of β arbitrary close to a 1/2 was recently achieved by David Zuckerman [Z-96].

with probability $\frac{1}{|\mathcal{E}_i|}$. Thus, we obtain

$$\begin{aligned} & Pr[C \text{ is not eliminated in round } i \mid C \in \mathcal{E}_i] \\ & \leq \left(1 - \frac{1}{|\mathcal{E}_i|}\right)^{(1-\beta)|\mathcal{E}_i|} \\ & \approx e^{-(1-\beta)} \end{aligned}$$

Therefore, we have:

$$Pr[C \in \mathcal{E}_{i+1}] \leq e^{-i(1-\beta)}$$

So, we have the following lemma:

Lemma 4.1.1 With probability $1 - \frac{1}{N}$, the protocol completes in fewer than $\frac{\ln M}{1-\beta}$ rounds.

Proof : By substituting $\frac{\ln M}{1-\beta}$ for i in the preceding discussion and applying Markov's inequality. $\square$

4.2 Resilience

The proof of the resilience of this protocol is based on a simple but crucial trade-off that we elaborate upon next. As we increase β , the rate at which $\mathcal{E}$ decays becomes increasingly faster with respect to the rate at which $\mathcal{B}$ decays. The protocol works if β is small enough so that this increased decay rate of $\mathcal{E}$ is still not enough to overcome the initial size disadvantage that $\mathcal{B}$ had with respect to $\mathcal{E}$. In this case, all the bad committees will be eliminated before there are only N eligible committees left. The resilience figure, .0045 represents the value of β where the trade-off between the decay rate and the initial size disadvantage is made.

If we choose δ so that $\beta + \delta < \frac{1}{2}$, then, it suffices to show that the probability that the committee C , is dangerous is small. Recall that C is the committee that is chosen after the first stage of the protocol. This is because of the following reasoning: Let δ be chosen to be such that $\beta + \delta < \frac{1}{2}$. Then, if C is not dangerous, by definition, the fraction of bad players in C is smaller than $\frac{1}{2}$. Thus, we can use the result of Boppana and Narayanan [BN2-93] to complete our proof immediately.

Lemma 4.2.1 If $\ell \geq \frac{\ln(N|\mathcal{B}|)}{(1-\beta)}$, then

$$Pr(\exists \text{ a dangerous committee in } \mathcal{E}_\ell) \leq \frac{1}{N}$$

Proof : Let $C \in \mathcal{B}$ be a dangerous committee. Consider a fixed round i in the first phase. Each vote cast by an honest player in this round eliminates C with probability $\frac{1}{|\mathcal{E}_i|}$. There are $(1 - \beta)|\mathcal{E}_i|$ votes cast by honest players. Thus, the probability that C is not eliminated in

the i^{th} round is at most $\left(1 - \frac{1}{|\mathcal{E}_i|}\right)^{(1-\beta)|\mathcal{E}_i|} \approx e^{-(1-\beta)}$. Therefore, the probability that C is not eliminated in rounds 1 through ℓ is $e^{-(1-\beta)\ell}$. Thus, we get

$$\mathbb{E}(|\mathcal{E}_\ell \cap \mathcal{B}|) \leq e^{-\ell(1-\beta)} |\mathcal{B}|$$

If we apply Markov's inequality and substitute the value of ℓ we obtain the lemma. $\square$

Lemma 4.2.1 states that if the protocol runs for more than $\frac{\ln(N|\mathcal{B}|)}{(1-\beta)}$ rounds, then it is highly unlikely that any dangerous committee survives. In order to complete the proof, it would be sufficient to show that it is very likely that the protocol runs for at least $\frac{\ln(N|\mathcal{B}|)}{(1-\beta)}$ rounds. This is what we shall exhibit now. We begin with a lemma of what is called a Chernoff bound for occupancy, studied in [KMPS-94], where we throw $(1-\beta)t$ balls independently at random into t bins. The bound says that the number of empty bins does not deviate significantly from the expected value.

Lemma 4.2.2 Let $X_1, X_2 \dots X_{(1-\beta)t}$ be uniformly and independently chosen from $\{1, 2, \dots, t\}$. Let $Y_i = 1$ if $\exists j$ such that $X_j = i$. Let $Y = \sum_1^t Y_i$. Then, $\mathbb{E}(Y) = t(1 - e^{-(1-\beta)})$ and $Pr(|Y - \mathbb{E}(Y)| \geq \log t \sqrt{t}) \leq o\left(\frac{1}{\text{poly}(t)}\right)$

Proof : [KMPS-94] $\square$

Lemma 4.2.3 Let $Z_i = |\mathcal{E}_{i+1}|$. Let $p = e^{-(1-\beta)} - \beta$. If $\theta_i = \alpha i \log^2 N \sqrt{p^i |\mathcal{C}|}$, for $\alpha = \frac{1}{\sqrt{p}}$, then for any $i \leq N$,

$$Pr[Z_i < p^i |\mathcal{C}| - \theta_i] \leq o\left(\frac{1}{\text{poly}(N)}\right)$$

Consequently, if $\ell \leq -\frac{\ln(\frac{|\mathcal{C}|}{N})}{\ln(e^{-(1-\beta)} - \beta)}$, then

$$Pr[Z_\ell \geq N] \geq 1 - o\left(\frac{1}{\text{poly}(N)}\right)$$

Proof : We will show the following conditional statement for every choice of k . The lemma and its consequence follow by simply taking a union bound over all k thereafter.

$$\begin{aligned} & Pr[Z_k \leq p^k |\mathcal{C}| - \theta_k \mid Z_{k-1} \geq p^{k-1} |\mathcal{C}| - \theta_{k-1}] \\ & \leq o\left(\frac{1}{\text{poly}(N)}\right) \end{aligned}$$

This conditional statement follows immediately with the help of the following consequence of the occupancy bound described above.

$$Pr[Z_k \leq p Z_{k-1} - \log N \sqrt{Z_{k-1}}] \leq o\left(\frac{1}{\text{poly}(N)}\right)$$

Conditioning on the value of Z_{k-1} we obtain the following.

$$\Pr \left[Z_k \leq pZ_{k-1} - \log N \sqrt{Z_{k-1}} \mid Z_{k-1} \geq p^{k-1} |\mathcal{C}| - \theta_{k-1} \right] \leq o \left(\frac{1}{\text{poly}(N)} \right)$$

Plugging in the value of θ_{k-1} and doing a little algebraic manipulation gives the conditional statement and hence, the lemma. $\square$

Theorem 4.1 The protocol described in section 3 is an $O(\log N)$ round leader election protocol that is β -immune for every β smaller than .0045.

Proof : The lemmas 4.2.1 and 4.2.3 together imply that if ℓ is chosen such that $\frac{\ln(N|\mathcal{B}|)}{(1-\beta)} < \ell < \frac{\ln(\frac{|\mathcal{C}|}{N})}{-\ln(e^{-(1-\beta)} - \beta)}$, then after ℓ rounds, it is very likely that all the dangerous committees are gone and that the first phase is not yet complete. This in turn implies that the committee C is unlikely to be dangerous. We know that $|\mathcal{C}| = M$ and $|\mathcal{B}| = M^{1-\lambda\delta^2}$ where λ is 1/20th of eigenvalue gap for the expander graph of our choice. We also take note that there are explicit expander constructions that produce expanders with eigenvalue gap arbitrarily close to 1.

Substituting these values, we notice that it is sufficient if we can find an ℓ satisfying

$$\frac{\ln(NM^{1-\lambda\delta^2})}{(1-\beta)} < \ell < \frac{\ln\left(\frac{M}{N}\right)}{-\ln(e^{-(1-\beta)} - \beta)}$$

Notice further that as we increase the length of the walk on the expander in the construction of the committees by increasing k , and thus increase M relative to N , $\ln(NM^{1-\lambda\delta^2}) \approx (1-\lambda\delta^2)\ln M$ and $\ln(\frac{M}{N}) \approx \ln M$. Thus, we need to find a ℓ satisfying

$$\frac{(1-\lambda\delta^2)\ln(M)}{(1-\beta)} < \ell < \frac{\ln(M)}{-\ln(e^{-(1-\beta)} - \beta)}$$

Finally, we notice that if β is small enough, then such an ℓ exists. The largest possible value of β is easily computed as the solution the equation

$$\frac{(1-\lambda\delta^2)}{(1-\beta)} = \frac{1}{-\ln(e^{-(1-\beta)} - \beta)}$$

where $\delta = \frac{1}{2} - \beta$ and λ is $\frac{1}{20}$. It is easily verified that the solution to the above equation is at least .0045. This completes the proof of the main theorem. $\square$

5 Non Uniform protocols

In this section we study non-uniform protocols. We assume, in this case, that we design our protocol probabilistically. It is easy to see that this situation implies the non-uniform case. The part of the protocol that we single out for attention is the construction of committees. Suppose we replace the expander based deterministic construction by a random construction as follows. A collection $\mathcal{C} = \{C_i : 1 \leq i \leq M\}$ of committees is chosen with $|C_i| = k \log N$. Each member of each committee is chosen independently and uniformly at random from $\{1, 2, \dots, N\}$.

Let X , $|X| = \beta N$ be any subset of $\{1, 2, \dots, N\}$. We wish to view X as the set of dishonest players. An immediate consequence of the Chernoff bound is that for any i

$$Pr[|X \cap C_i| \geq (\beta + \delta) |C_i|] \leq 2^{-\frac{\delta^2 k \log N}{4}}$$

Here, the probability is taken over all possible choices of C_i . Thus, if k is chosen to be $\frac{4\mu}{\delta^2}$ where μ is an yet un-chosen constant, then,

$$Pr[|X \cap C_i| \geq (\beta + \delta) |C_i|] \leq N^{-\mu}$$

Let $\text{Bad}(X)$ be the sets in $\mathcal{C}$ such that $|C_i \cap X| \geq (\beta + \delta) |C_i|$. Then, clearly, the expected size of $\text{Bad}(X)$ is at most $N^{-\mu} M$. Additionally, for any α , the probability that the size of $\text{Bad}(X)$ exceeds αM is at smaller than $\binom{M}{\alpha M} N^{-\mu \alpha M}$. Using the fact that $\binom{M}{\alpha M} < M^{\alpha M}$ and choosing $\alpha = N/M$, we obtain,

$$Pr(|\text{Bad}(X)| \geq N) \leq \left(\frac{M}{N^\mu}\right)^N$$

If we choose μ to be a large enough constant so that $M \leq N^{\mu-1}$, we observe that

$$Pr(|\text{Bad}(X)| \geq N) \leq N^{-N}$$

Since there are at most 2^N possible choices for X , we get the following lemma,

Lemma 5.0.4 There exists a collection $\mathcal{C} = \{C_i : 1 \leq i \leq M\}$ of committees such that for any X of size βN , $|\text{Bad}(X)| \leq M^{1-\lambda}$ for any $\lambda < 1$. The size of M is polynomial in N . The degree of the polynomial depends only on λ . And so that the size of any committee C_i is $O(\log N)$. The size of committees depends polynomially on λ^{-1} and $(\frac{1}{2} - \beta)^{-1}$.

Proof : Consider the union bound over the possible choices of X in the previous discussion. $\square$

It is now easily seen that by choosing the value of $\lambda = 4/5$, the root of the equation

$$\frac{(1-\lambda)}{(1-\beta)} = \frac{1}{-\ln(e^{-(1-\beta)} - \beta)}$$

is larger than $\frac{1}{2}$. Thus, showing the existence of a protocol that is $\frac{1}{2}$ -immune and runs in $O(\log N)$ rounds. This completes the proof of the second main theorem:

Theorem 5.1 For every $\epsilon > 0$, there exists an $O(\log N)$ round leader election protocol that is β -immune for $\beta = \frac{1}{2} - \epsilon$.

6 Further Research and Open Problems

We remark that due to the inherent simplicity of our solution it is conceivable that our technique can be parallelized in order to reduce the number of rounds even further. Moreover, it is conceivable that our technique might find other applications. For example, it might be useful in the [GGL-91] setting.

Acknowledgements

We thank Nati Linial for a helpful discussion, Ravi Boppana and Babu Narayanan, for providing us with a manuscript of their paper and Nabil Kahale and David Zuckerman, for telling us about their subsequent work.

References

- [AL-89] M. AJTAI, N. LINIAL The influence of large coalitions. *IBM Research Report 7133* (67380), Nov. 1989.
- [AN-90] N. ALON, M. NAOR Coin-flipping games immune against linear-sized coalitions. *FOCS-90* pp. 46-54. Journal version in *SIAM Journal on Computing* 22:403-417, 1993.
- [AR-89] N. ALON, M. RABIN Biased coins and randomized algorithms. *Advances in Computing Research, JAI Press* (Silvio Micali, ed.) 5:499-507, 1989.
- [BGW-88] BEN-OR M., S. GOLDWASSER, A. WIGDERSON Completeness Theorem for Non-cryptographic Fault-tolerant Distributed Computing, *STOC 88*, pp. 1-10
- [BL-85] M. BEN-OR, N. LINIAL Collective coin flipping, robust voting schemes and minima of Banzhaf values. *FOCS-85* pp. 408-416.
- [BL-89] M. BEN-OR, N. LINIAL Collective coin flipping. *Advances in Computing Research, JAI Press* (Silvio Micali, ed.) 5:91-116, 1989.
- [BLS-87] M. BEN-OR, N. LINIAL, M. SAKS Collective coin flipping and other models of imperfect randomness. *Colloq. Math. Soc. Janós Bolyai No., 52 Combinatorics Eger* pp. 75-112, 1987. North-Holland Publ.
- [BN1-93] R. BOPPANA, B. NARAYANAN. The biased coin problem *STOC 93*
- [BN2-93] R. BOPPANA, B. NARAYANAN. Collective Coin Flipping and Leader Election with Optimal Immunity., *manuscript*.

- [B-85] G. BRACHA An $O(\log n)$ expected rounds randomized Byzantine Generals protocol. *STOC-85* pp. 316-326.
- [CCD-88] D. CHAUM, C. CREPEAU AND I. DAMGARD Multiparty Unconditionally Secure Protocols. *STOC 88*, pp. 11-19.
- [CD-89] B. CHOR, C. DWORK Randomization in Byzantine agreement. *Advances in Computing Research, JAI Press* (Silvio Micali, ed.) Vol. 5, 1989.
- [CL-93] J. COOPER, N. LINIAL Fast Perfect-Information Leader-Election Protocol with Linear Immunity. *STOC-93* pp. 662-671.
- [FM-88] P. FELDMAN, S. MICALI Optimal algorithms for Byzantine Agreement. *STOC-88* pp. 148-161.
- [GG-79] O. GABBER, Z. GALIL Explicit construction of linear sized super-concentrators. *FOCS-20* pp. 364,370. Also see the journal version, *JCSS***22** (1981), pp. 407-420.
- [G-93] D. GILLMAN A Chernoff bound for random walks on expanders. *FOCS-93*.
- [GGL-91] O. GOLDBREICH, S. GOLDWASSER, N. LINIAL Fault-tolerant Computation in the Full Information Model. *FOCS 91* pp. 447-457.
- [IZ-89] R. IMPAGLIAZZO, D. ZUCKERMAN How to recycle random bits. *FOCS 89* pp. 248-253.
- [KKL-88] J. KAHN, G. KALAI, N. LINIAL The influence of variables on boolean functions. *FOCS 88*, pp. 68-80.
- [K-94] N. KAHLE. Large deviation bounds for markov chains. DIMACS technical report 94-39, June 1994.
Electronic copy in <http://dimacs.rutgers.edu/TechnicalReports/1994.html>
- [KMPS-94] A. KAMATH, R. MOTWANI, K. PALEM, P. SPIRAKIS Tail bound for Occupancy and the Satisfiability Threshold Conjecture. *FOCS-94*, pp. 592-603.
- [GMW-87] O. GOLDBREICH, S. MICALI, A. WIGDERSON How to Play any Mental Game. *STOC 87*, pp. 218-229.
- [L-92] N. LINIAL Games Computers Play: Game-Theoretic Aspects of Computing *C.S. Tech. Report 92-5, The Hebrew University of Jerusalem*, February 1992.
- [S-89] M. SAKS A robust noncryptographic protocol for collective coin flipping. *SIAM Journal of Discrete Mathematics*, 2:240-244, 1989.
- [Z-96] D. ZUCKERMAN Randomness-Optimal sampling, Extractors, and Constructive Leader Election. *STOC-96* pp. 286-295.