Collective Coin Flipping

Michael Ben-Or
and
Nathan Linial

Institute of Mathematics
and Computer Science
The Hebrew University
Jerusalem 91904
Israel

Randomized algorithms play an important role in parallel and distributed processing. The use of such algorithms assumes that each processor is able to generate random bits during the computation. In some applications the algorithm requires that the ***same*** random bit be generated by a set of processors. This task is easy if we assume that no faults may occur. We have one of the processors flip a coin and announce the outcome. If, however, the processor assigned to flip the coin happens to be faulty this may ruin the probabilistic requirement for our randomized algorithm. Suppose, then, that each processor is equipped with a fair coin, how can they generate a global coin flip which is only slightly biased despite failure of some of the processors?

This problem was considered before, mostly in the framework of the Byzantine Generals Problem ([ACGM,BE,BD,Br,BR,Ra,Ya]). These past solutions are all based on the assumption that information may be communicated so that only some of the parties can read it. This is achieved either by choosing an appropriate model of communication, or by resorting to cryptography. We want to avoid such assumptions. Technically this means that we deal only with games of complete information.

The most obvious approach to solve this question is via what we call a one-round coin flipping scheme: Say that there are n processors involved. Fix a Boolean function $f:\{0,1\}^n \rightarrow \{0,1\}$. Whenever a random bit is needed, instruct each processor P_i to flip his coin and announce the outcome x_i. The global coin flip is taken as $f(x_1,\ldots,x_n)$. How sensitive is this approach to the possible existence of faults? If the processors act simultaneously, the situation is very favorable. If f is the parity function, then clearly, even if only one processor is in order, this yields a perfect coin flip. In a distributed environment, where one cannot assume perfect simultaneity, the parity function is not very useful. A single faulty processor which announces his bit already knowing the bits announced by all the rest, has complete control over the global bit.

We are thus looking for Boolean functions on which every variable has only a small influence. The discussion above hints at the features we should expect from measures of influence. Indeed there are a number of inequivalent such quantities. We present the simplest and probably the most natural one among them: Let $S \subseteq \{1,\ldots,n\}$ be a subset of the variables of the Boolean function f. Randomly set the variables outside S by independent perfect coin flips. This partial setting may already determine the value of f regardless of the values the variables in S. A measure for the influence of S on f is the probability that this does not happen and the variables in S "have the last word" in determining f. Our goal is to find functions f for which this measure of influence is as small as possible for all subsets $S \subseteq \{1,\ldots,n\}$.

The assumption that the variables in S are set after those in $\{1,\ldots,n\}\backslash S$ is obviated by the following observation (Lemma 2.1): For every Boolean function f there is a monotone function g such that

$$Pr(f=0) = Pr(g=0)$$

and every $S \subseteq \{1,\ldots,n\}$ influences f at least as much as it influences g. So for our purposes g is at least as good as f. This lemma thus says that it suffices to consider monotone f's. This observation removes the need for S to be set only after the other variables. All the information we need in this case is embodied in the probability of $f=0$ when all variables in S are set zero (resp. one).

Some very basic questions regarding Boolean functions thus arise. They also turn out to be rather natural questions in game theory. A Boolean function is for the game theorist a *simple game*. Variables are *players* and sets of variables are *coalitions*. The measure of influence of a player (the case $|S|=1$) is a quantity already considered in game theory, under the name of *Banzhaf − Coleman power index*. Much attention has been given in game theory to measures of influence of players and sets of players in games [Ow]. The most important among these is the Shapley value. For us the Shapley value does not seem to be the most appropriate.

Consider the Boolean function $f(x_1,\ldots,x_n)=x_k$, which in game theory is the dictatorship of player k. Here, the influence (Banzhaf index) of the dictator k is 1 and is 0 for all the other players. The influence of each variable on the majority function is $\Theta(1/\sqrt{n})$. It may be thought that this is best but this is not the case. We present a Boolean function for which the influence of each variable is only $O(\log n/n)$. On the other hand it follows from known facts in either game theory [*Ha*] or combinatorics [*H*] that for any function the average influence of a variable is at least $\Omega(1/n)$. This gap between the upper and lower

bound is very intriguing and our conjecture is that the upper bound is closer to the truth. Put informally, there is always some player who affects the outcome of the game in a disproportionate manner.

Until this point we restricted our discussion to single-round schemes. The reader can certainly think of more elaborate schemes for collective coin flipping based on more complicated protocols than just an application of a Boolean function . We make the observation (Proposition 5.2) that the most general coin flipping scheme can be described as follows: There is a rooted binary tree T whose leaves are labeled by zeros and ones. Each internal node of T is labeled by the name of one of the players. We start at the root. Whenever an internal node is reached, the player at this node flips a coin and announces the outcome. According to this outcome the game proceeds to either the right or the left son of that node in the tree. When a leaf is reached the game terminates and the outcome of the game is the $0/1$ label of the leaf.

Such a labeled tree is called a *multi − round coin flipping scheme*. This is a more complicated setup and we need to consider at least two quantities which measure influence. Let $S \subseteq \{1, \ldots, n\}$ be a set of players. Assume that the players not in S do follow the rules and flip perfect coins. Then there is a best strategy for S to maximize the probability of a zero outcome. The difference between this probability and the probability of zero if members of S, also, play randomly measures the influence of S towards zero. A similar quantity is defined for an outcome of one.

Our main relevant result (Theorem 5) says that in every multiround game there is a player with an $\Omega(1/n)$ influence towards one. Of course the same holds towards zero. Unfortunately, we have no similar result to guarantee the existence of a player with a substantial influence both towards zero and one. Unlike the one-round situation this $\Omega(1/n)$ bound is known to be best possible. We exhibit games showing the tightness of this bound.

More general questions regard the influence of sets of players in both single and many round games. Particularly interesting is the following notion of $\varepsilon - control$. For $\varepsilon > 0$ a coalition S has ε-control over a game G if S's influence on the outcome of G is $\geq \varepsilon$. (As mentioned before we deal with a number of measures for influence and the definition of ε-control clearly depends on the measure at hand.) The goal is of course to find games which are not ε-controlled by small sets.

To illustrate this notion we remark that majority is ε-controlled by $O(\varepsilon \sqrt{n})$ players. We construct single-round games in which $\Omega(\varepsilon\, n^{0.63\ldots})$ players are needed to achieve ε-control. Recently Saks [S] analyzed a multistage game where $\Omega(n/\log n)$ players are required to achieve ε-control. We conjecture that Saks' construction is essentially the best possible. This means that there is always a negligible minority of the players which almost determines the outcome assuming all the others play randomly. Even the one-round version of this question is open and very intriguing: Given a Boolean $f: \{0,1\}^n \to \{0,1\}$ with $Pr(f = 0) = 1/2$, is there a set of $o(n)$ variables such that even if the assignment of values to all other variables is known there is a constant probability that this does not yet determine f ?

In the results described here we always assume that the set of faulty processors is fixed during the game. A completely different question arises if the faulty processors are determined by an adversary during the course of the game. In the single-round context it can be shown that the most robust scheme against such an adversary is the majority voting scheme. This is an easy consequence of the isoperimetric inequality in the cube [H]. In the multistage case this problem is left open. A special case of this question was recently settled in a paper of Lichtenstein, Linial and Saks [LLS].

As explained in the introduction a one-round coin flipping scheme is nothing but a Boolean function $f: \{0,1\}^n \to \{0,1\}$. Unless otherwise stated f is usually assumed to satisfy: $|f^{-1}(0)| = |f^{-1}(1)| = 2^{n-1}$ (or equivalently $Pr(f = 0) = 1/2$). With a slight change of notation f may be thought of as a function from the power set of $\{1, \ldots, n\}$ into $\{0,1\}$. Consequently we speak of $f(A)$ for $A \subseteq \{1, \ldots, n\}$. We now set to formally define our notions of influence for a set of variables of f.

Let

$$Q_0 = Q_0(f, S) = \{\ A \mid A \cap S = \emptyset \text{ and for every } B \subseteq S\ ,\ f(A \cup B) = 0\}$$
$$Q_1 = Q_1(f, S) = \{\ A \mid A \cap S = \emptyset \text{ and for every } B \subseteq S\ ,\ f(A \cup B) = 1\ \}\ .$$

Also

$$q_0 = q_0(f,S) = |\,Q_0(f,S)\,| \cdot 2^{|S|-n}$$

$$q_1 = q_1(f,S) = |\,Q_1(f,S)\,| \cdot 2^{|S|-n}$$

$$q_2 = q_2(f,S) = 1 - q_0 - q_1.$$

It is seen that q_0 (q_1) is the probability that the values assigned to variables outside S already set f to zero (one). The probability that this partial assingment does not determine f is q_2. Accordingly, the *influence of S on f* is defined as

$$I_f(S) := q_2(f,S).$$

The *influence of S on f towards zero* (*one*) are defined to be

$$I_f^0(S) := q_0(f,S) + q_2(f,S) - Pr(f=0),$$

$$I_f^1(S) := q_1(f,S) + q_2(f,S) - Pr(f=1),$$

respectively. Note that $q_0(f,S) + q_2(f,S)$ is the probability that $f=0$ assuming the players in S try to set f to zero. The influence towards zero is defined as the excess of this probability over $Pr(f=0)$. For $1 \le r \le n$ we let

$$I_f(r) := \max\{I_f(S) \mid |S| = r\}$$

and similarly for ${I_f}^0(r)$, ${I_f}^1(r)$. Also $I_f := I_f(1)$.

If $S=\{x\}$ is a singleton $I_f(\{x\})$ is a quantity which was studied in game theory. A monotone Boolean function $v:\{0,1\}^n \to \{0,1\}$ is called a *monotone simple game* (N,v) with $N=\{1,\ldots,n\}$. Variables are called *players* and sets of players are *coalitions* . In this context $I_f(\{x\})$ is called the *Banzhaf − Coleman power index* of the player x, see [Ow]. We freely interchange between these two equivalent terminologies, according to the context. In game theoretic terms one of our main problems is thus to find simple games where all Banzhaf indices are small. Our first observation about this problem is that monotone Boolean functions are as good as any other:

Lemma 2.1: Let $f:\{0,1\}^n \to \{0,1\}$ be a Boolean function. Then there exist a monotone function $g:\{0,1\}^n \to \{0,1\}$ such that

$$Pr(f=0) = Pr(g=0)$$

and for every $S \subseteq \{1,\ldots,n\}$

$$\begin{aligned} I_f^0(S) &\ge I_g^0(S), \\ I_f^1(S) &\ge I_g^1(S), \\ I_f(S) &\ge I_g(S). \end{aligned}$$

Proof: The proof is based on a standard technique in the extremal theory of finite sets viz. compression. The function f is made more and more monotone until eventually g is reached. Let us pick an $1 \le x \le n$ and consider $\tilde{f}:\{0,1\}^n \to \{0,1\}$ which is obtained from f as follows: Suppose that for some $A \subseteq \{1,\ldots,n\}$,

$$f(A)=1 \quad \textit{and} \quad f(A \cup \{x\}) = 0.$$

Then we set

$$\tilde{f}(A)=0 \quad \textit{and} \quad \tilde{f}(A \cup \{x\}) = 1.$$

Otherwise $\tilde{f}$ equals f . We keep doing the same thing with $\tilde{f}$, but with respect to a different $1 \le y \le n$. It is easily seen that as long as f is not monotone there is an x for which $\tilde{f} \ne f$. Also after a finite number of such transformations the function becomes monotone. It is also clear that

$Pr(f = 0) = Pr(\tilde{f} = 0)$.

We want to show that $q_0(f, S) \leq q_0(\tilde{f}, S)$ for all S. Let us start with the case $x \in S$. In this case we show that

$$Q_0(f, S) \subseteq Q_0(\tilde{f}, S).$$

Suppose to the contrary that there is a set $A \in Q_0(f, S) \setminus Q_0(\tilde{f}, S)$. Since $A \notin Q_0(S, \tilde{f})$ there is a set $B \subseteq S$ with $\tilde{f}(A \cup B) = 1$. But $A \in Q_0(f, S)$ and so $f(A \cup B) = 0$. The definition of $\tilde{f}$ implies that $f(A \cup B \setminus \{x\}) = 1$ and again the fact that $A \in Q_0(f, S)$ implies that $x \in A$. But this contradicts $x \in S$.

Now we have to prove our claim for $x \notin S$. Here we show that if $A \in Q_0(f, S) \setminus Q_0(\tilde{f}, S)$ then $A \setminus \{x\} \in Q_0(\tilde{f}, S) \setminus Q_0(f, S)$. Consequently the desired inequality holds. We repeat the previous arguments and conclude that $x \in A$ and $f(A \setminus \{x\} \cup B) = 1$ whence $A \setminus \{x\} \notin Q_0(f, S)$. Now suppose that $\tilde{f}(A \setminus \{x\} \cup C) = 1$. But the definition of $\tilde{f}$ implies that also $f(A \setminus \{x\} \cup C) = f(A \cup C) = 1$. This, however, contradicts $A \in Q_0(f, S)$ and completes our argument.

The proof for q_1 is identical. The conclusion for I^0, I^1 and I is then immediate. □

We prove the following easy lemma:

Lemma 2.2 (a) Let $f: \{0,1\}^n \to \{0,1\}$ be a Boolean function and let $x \in \{1, \ldots, n\}$. Then

$$I_f(\{x\}) = 2^{1-n} \sum_{x \notin S} | f(S \cup \{x\}) - f(S) | .$$

(b) If f is monotone, let $d(x) = |\{ S \subseteq \{1, \ldots, n\} | \ x \in S \text{ and } f(S) = 0\}|$ then

$$I_f(\{x\}) = 2Pr\ (f = 0) - 2^{-n}\ d(x).$$

Proof:
(a) This is just a restatment of the definition.
(b) Let us denote

$$C_0 = \{ S \mid x \notin S\ ,\ f(S) = 0 \}\ ,\ C_1 = \{ S \mid x \in S\ ,\ f(S) = 0 \}.$$

From (a) and the monotonicity of f it follows that

$$I_f\ (\{x\}) = (\ | C_0 | - | C_1 |\ )\ 2^{-n+1}.$$

But $| C_0 | + | C_1 | = 2^n Pr(f = 0)$. Divide by 2^{n-1} and subtract to deduce (b). □

3. One Round Games - a Construction.

This section presents a one round coin flipping scheme among n players in which the influence of any particular player is only $O(\frac{\log n}{n})$.

Theorem 3: There are one round coin flipping schemes $G = G_n$ where

$$I_G = O(\frac{\log n}{n})$$

Proof: To describe the idea of this construction let us ignore for a while issues of integrality and divisibility. For given n let b be the (unique) solution of the equation

$$(2^b - 1)^{1/b} = 2^{1-1/n}.$$

We will later show that this b satisfies

$$b = \log n - \log\log n + O(1).$$

Now decompose $[n]$ with n/b blocks of size b and consider the ideal J of those subsets of $[n]$ which contain no block.

$$|J| = (2^b - 1)^{n/b} = 2^{n-1}.$$

So we consider the coin flipping scheme where $v(A) = 0$ if $A \in J$ and $v(A) = 1$ if $A \notin J$. Simply stated, the overall decision is 1 if and only if a whole block unanimously votes one and is 0 otherwise. Now let us compute the influence of an individual player. Using Lemma 2.1 we have to find $d(x)$ which is the same for every $x \in [n]$. According to the definition of J we have

$$d(x) = (2^{b-1} - 1)(2^b - 1)^{\frac{n}{b}-1} = \frac{2^{b-1} - 1}{2^b - 1}\,(2^b - 1)^{n/b} = \frac{2^{b-1} - 1}{2^b - 1}\,|J|.$$

Hence

$$I_G = I_G(\{x\}) = (|J| - 2d(x)) \cdot 2^{-n+1} = 1 - \frac{2^b - 2}{2^b - 1} = \frac{1}{2^b - 1} = O(\frac{\log n}{n}).$$

To show the last equality go back to the relation between b, n:

$$(2^b - 1)^{1/b} = 2^{1-1/n}$$

or $n = -b/\log(1 - (\frac{1}{2})^b)$. We use

$$-\ln 2 \cdot (\frac{1}{2})^b \geq \log(1 - (\frac{1}{2})^b) \geq -(\frac{1}{2})^{b-1}$$

whence $b \cdot 2^b \cdot \log e \geq n \geq b \cdot 2^{b-1}$. These functions of b are increasing and therefore by evaluating them at the appropriate values of b the following bound on b results

$$|\, b - (\log n - \log\log n)\,| < 2.$$

It follows that

$$I_G = (2^b - 1)^{-1} = O(\frac{\log n}{n}).$$

To overcome the difficulties involved with b not being an integer we do as follows: We select any integer b and define α to be the real number for which

$$(2^b - 1)^{\alpha} = 2^{\alpha b - 1}.$$

Next we define a to be the integer nearest to α, say $a = \alpha + \varepsilon$, $|\varepsilon| \leq \frac{1}{2}$, and set $n = ab$. The ideal J is defined as before and has size

$$|J| = (2^b - 1)^a = (2^b - 1)^{\alpha+\varepsilon} = 2^{\alpha b - 1}\ \ (2^b - 1)^{\varepsilon}.$$

While $2^{n-1} = 2^{(\alpha+\varepsilon)b-1}$. So

$$\frac{|J|}{2^{n-1}} = (\frac{2^b - 1}{2^b})^{\varepsilon} = 1 - \frac{\varepsilon}{2^b} + O(\frac{\varepsilon^2}{4^b}).$$

Therefore, by adding $2^{n-b-1} \cdot \varepsilon$ sets to J, still maintaining J being an ideal, the influence of every player can rise to at most $(1 + \varepsilon)/2^b = O(\frac{\log n}{n})$. □

For completeness sake we add the following proposition:

Proposition 3: Let f be a Boolean function on n variables, with $Pr(f = 0) = 1/2$. Then

$$\sum_{i=1}^{n} I_f(x_i) \geq 1.$$

This bound is tight as shown by $f(x_1, \ldots, x_n) = x_1$. This result is known in game theory [H]. The reader may verify it by noticing that given any set of 2^{n-1} vertices in the n-dimensional cube there are at least 2^{n-1} edges in the associated cut. Theorem 5 below gives a more general result. Meanwhile we state:

Conjecture 3: For every Boolean function f on n variables with $Pr(f = 0) = 1/2$ there is a variable x such that

$$I_f(x) = \Omega(\frac{\log n}{n}).$$

Proposition 3 implies the existence of a variable with influence at least $1/n$. Noga Alon showed, using eigenvalue methods the existence of a variable with an influence of at least $(2 - o(1))/n$.

4. Symmetric Games

The following symmetry condition is commonly imposed in the context of human voting games: If every player reverses his vote the collective decision has to change as well. If our coin flipping scheme is described by a simple game the condition is that for every coalition $S \subseteq N$

$$v(S) + v(N \setminus S) = 1.$$

It also turns out that symmetric games with small influences are useful building blocks for robust coin flipping schemes that are not necessarily symmetric (see Section 6).

We would like to consider our general problem in the context of symmetric coin flipping games:

(1) Find symmetric games which minimize $I_G(r)$. Notice that the symmetry condition implies that for all $S \subseteq N$, $I_G^1(S) = I_G^0(S)$.

(2) Find symmetric games for which the least number of players which ε-control the game is as large as possible.

We have the following result for the symmetric case:

Theorem 4: (a) There exist symmetric games $G_n = G$ for which the influence function satisfies

$$I_G(r) \leq \frac{r}{n^\alpha} \quad for \quad 0 \leq r \leq n^\alpha$$

where $\alpha = \log_3 2 \ = 0.6309\ldots$.

(b) There exist symmetric games $G_n = G$ for which the individual influence function satisfies

$$I_G \leq \frac{1}{n^\beta}$$

where $\beta = \log_7(32/9) \ = 0.6518\ldots$.

Proof: We first introduce the notion of the *composition of games*: Let $G = ([n], v)$ and $G_i = (P_i, w_i)$ be simple games, $(i = 1, \ldots, n)$, where the sets of players P_i are mutually disjoint. The G-composition of $\{G_i\}$ is the game

$$G = (\bigcup_1^n P_i, w)$$

where

$$w(S) = 1 \quad \textit{iff} \quad v(\{i \mid w_i(S \cap P_i) = 1\}) = 1.$$

Intuitively this means that the set of players is composed of n committees where the internal voting in the i-th committee is by the w_i rule, and the committees votes are combined by v.

Next, we introduce some more definitions. A hypergraph H is *intersecting* if every two edges have a nonempty intersection. H is *two-colorable* if there is a partition of the vertices $V = V_1 \cup V_2$ such that no edge is contained in either V_1 or V_2. Finally, we say that a game $G = (N, v)$ is *transitive* if there is a transitive group acting on N under which v is invariant.

Proposition:

(a) Let $H = (V, E)$ be an intersecting, non 2-colorable hypergraph, then the coin flipping scheme $G = (V, v)$ given by

$$v(S) = 1 \quad \textit{iff} \quad \text{there is an } A \in E \quad \text{such that} \quad A \subseteq S,$$

is symmetric.

(b) If $G = ([n], v)$ is a symmetric coin flipping scheme and so are $G_1, \ldots, G_n$ then the G composition of $\{G_i\}$ is also symmetric.

(c) Let $G = ([n], v)$ and $H = H_1 = H_2 = \cdots = H_n$ be symmetric transitive games and let K be the G composition of $\{H_i\}$, then K also transitive and

$$I_K = I_G \cdot I_H \quad .$$

Proof: (a),(c) and the transitivity of K in (c) are simple. Since K is transitive $I_K = I_K(\{x\})$ for any player x in K. We may therefore assume that x is some player in $H_1 = H$. Now $I_K(\{x\})$ is the probability that the game is not determined by the votes of all other players. This probability is exactly the probability that H_1 is not determined by the the votes of other players in H_1, that is I_H, times the probability that K is not determined by the outcome of the games $H_2, \ldots, H_n$. Since the players in $H_2, \ldots, H_n$ flip coins to set their votes and $Pr(H = 0) = Pr(H = 1) = 1/2$ the probability that K is not determined by the "votes" of $H_2, \ldots, H_n$ is I_G, and thus

$$I_K = I_H \cdot I_{G.} \quad \square$$

Consider the following two examples:

(1) Let $n = 2t - 1$ and consider the hypergraph of all t-sets of $[n]$ - It is clearly intersecting, non 2-colorable and the game it determines (i.e. majority voting) is transitive.

(2) The hypergraph of lines in the Fano plane. It has 7 vertices and the edges are {(1,2,4), (2,3,5), (3,4,6), (4,5,7), (1,5,6), (2,6,7), (1,3,7)}. It is easy to check directly that it is intersecting, non 2-colorable and transitive. For readers familiar with projective geometry only non 2-colorability needs elaboration. This is true because any set of 4 points that does not contain a line is the complement of a line in this plane.

Using these examples we can now prove our theorem.

Part (a): Let H_1 be the majority game of 3 players. Define H_k recursively as the H_1 composition of three copies of H_{k-1}. Denote by $n = 3^k$ the number of players in H_k and let $J(n, r) = I^1_{H_k}(r)$ be the largest influence towards 1 that r players can have in H_k, that is, the probability of outcome 1 if all the r players vote 1 minus $Pr(H_k = 1) = 1/2$. Clearly there is a set of $2^k = n^\alpha$ players that completely control H_k so only $r \leq n^\alpha$ is of interest. We prove by induction on k that for such r

$$J(n,r) \le \frac{r}{2n^{\alpha}}\,.$$

This is true for $k = 1$ as can easily be verifi ed.

To proceed we consider how these r players are split among the three H_{k-1} component of H_k. Say there are r_i of them in the i-th component $i = 1, 2, 3$. We are allowed to assume that for all i

$$0 \le r_i \le (\frac{n}{3})^{\alpha} = \frac{1}{2}\, n^{\alpha},$$

since $(\frac{n}{3})^{\alpha}$ players can have complete control over H_{k-1}. The condition $\Sigma r_i = r$ must clearly hold, too.

The best strategy for the r players in order to achieve an outcome of 1 in the game is for the r_i players in the i-th component to play towards 1 in their component. The probability for the game to end with a 1 under such strategy is the probability that at least two of the components end with 1. The probability is, therefore,

$$[\frac{1}{2} + J(n/3, r_1)][\frac{1}{2} + J(n/3, r_2)][\frac{1}{2} + J(n/3, r_3)] +$$

$$[\frac{1}{2} + J(n/3, r_1)][\frac{1}{2} + J(n/3, r_2)][\frac{1}{2} - J(n/3, r_3)] +$$

$$[\frac{1}{2} + J(n/3, r_1)][\frac{1}{2} - J(n/3, r_2)][\frac{1}{2} + J(n/3, r_3)] +$$

$$[\frac{1}{2} - J(n/3, r_1)][\frac{1}{2} + J(n/3, r_2)][\frac{1}{2} + J(n/3, r_3)] =$$

$$= \frac{1}{2} + \frac{1}{2}\, \Sigma\, J(n/3, r_i) - 2\, \Pi\, J(n/3, r_i) \le$$

$$\le \frac{1}{2} + \frac{1}{2}\, \Sigma\, J(n/3, r_i).$$

Therefore

$$J(n,r) \le \frac{1}{2} \max\, \Sigma\, J(n/3, r_i)$$

where the maximum is over all choices of r_1, r_2, r_3 with $0 \le r_i \le (\frac{n}{3})^{\alpha} = \frac{1}{2}\, n^{\alpha},\ \Sigma r_i = r$.

By induction

$$J(n/3, r_i) \le \frac{r_i}{2(n/3)^{\alpha}} = \frac{r_i}{n^{\alpha}}$$

and so

$$J(n,r) \le \frac{1}{2}\, \Sigma\, \frac{r_i}{n^{\alpha}} = \frac{r}{2n^{\alpha}}$$

as claimed. The conclusion about ε-control follows by solving $J(n,r) \ge \varepsilon$ for r.

Part (b): The construction here is similar to the construction above with the building block being the Fano game rather than the majority of three game. Let $F = F_1$ be the Fano game on seven players, and inductively define F_k to be the F composition of seven copies of F_{k-1}. Let $n = 7^k$ be the number of players in F_k. It is easy to see that $I_F = 9/32$, and thus

$$I_{F_k} = \frac{9^k}{32^k} = \frac{1}{n^\beta}. \qquad \square$$

5. Multistage Games - Lower Bounds

We have already mentioned in Proposition 3 a limitation of single round coin flipping schemes: There is always a player with an $\Omega(1/n)$ influence on the outcome. One can certainly devise more elaborate schemes for n players to flip a coin and hope to reduce the influence of any of the participants on the outcome. As it turns out, the same limitation still prevails for much more general schemes of complete information. The most general coin flipping scheme we consider here can be described as follows:

Definition: (a) Let $X_1, \ldots, X_n$ be finite probability spaces, and let V be a non empty set. A V-valued random variable $f : \prod X_i \longrightarrow V$ is called a *choice function* for the players $\{1, \ldots, n\}$ on V. We say that the choice function f is *controlled* by player i if f depends only on its i-th coordinate. We say that f is a *coin flip by player* i if f is controlled by player i, $V = \{v_1, v_2\}$ is of size two, and $Pr(f = v_1) = Pr(f = v_2) = 1/2$.

(b) A *general coin flipping scheme* (T, N) is a rooted tree T of finite depth with leaves labeled 0 or 1 and internal nodes labeled by choice functions (for the players N) on the set of their children. To determine the coin flip by this scheme we start at the root. Whenever an internal node is reached use the choice function at the node to select one of its children and move down the tree to that node. Continue in this manner till a leaf is reached. The label at this leaf determines the outcome of the coin flip.

(c) A *restricted coin flipping scheme* is a general coin flipping scheme (T, N) where the choice function attached to each internal node is controlled by one of the players, and a *Boolean coin flipping scheme* is a restricted coin flipping scheme where all the choice functions are just coin flips by one of the players.

We now define the notion of influence in this general context. Let (T, N) be a general coin flipping scheme and let $S \subseteq N$. Denote by $Pr(T = 0)$ the probability that the outcome of T is zero, assuming that at each internal node v the choice function f_v is used to select the next node, and that all the players pick their assignments to f_v according to the prescribed probability distributions. Now assume that at each node the players outside of S first select their assignments using the given probability distributions and then, given these assignments, the players in S can select their assignments to the choice function at the node according to their best strategy to maximize the probability of outcome 1. The probability of S failing and T ending with 0 despite the effort by S is denoted by

$$q_0 = q_0(T, S).$$

The definition of q_1 is symmetric. The reader can verify that if T has a single internal node labeled by a single round scheme then the q_δ defined here are the same as those defined in section 2 for the single round case. In the same way we define the influence of S in T towards zero (one) to be

$$I_T^0(S) = 1 - q_1(T, S) - Pr(T = 0),$$

$$I_T^1(S) = 1 - q_0(T, S) - Pr(T = 1).$$

Also for $1 \leq r \leq n$, and $\delta \in \{0, 1\}$,

$$I_T^\delta(S) = \max_{|S|=r} I_T^\delta(S).$$

and for $r = 1$

$$I_T^\delta = I_T^\delta(1).$$

Note that $1 - q_1(T, S)$ is the probability that the outcome is 0 when the players in S play their best strategy to reach 0. The difference between this and the probability that the outcome is 0 when all players play according to the scheme T is defined to be the influence of S towards 0.

Unlike the one round schemes that were studied in Game Theory and in the early days of Computer Science (e.g. in the context of Threshold Functions [Wi]), the influence of players in general coin flipping schemes, to the best of our knowledge, has not been studied before. For example, the following basic question has not been answered before: Can we approximate an unbiased coin as well as we wish despite the intervention of one of the players, by a long enough game. Or using our notation, given n and $\varepsilon > 0$, is there a general coin flipping scheme T for n players with $Pr(T = 1) = Pr(T = 0) = 1/2$ such that $I_T^0, I_T^1 < \varepsilon$.

In the following theorem we answer this question negatively, by showing that in any general coin flipping scheme for n players, for any r, $1 \le r \le n$, there is always a set S of r players that can bias the coin by $\Omega(r/n)$. We wonder whether this natural result is a consequence of some more general principle. In contrast, consider the example at the beginning of Section 6. It shows an election scheme which cannot be biased by any single player.

Theorem 5: Let (N, T) be an n player general coin flipping scheme. Let $\delta \in \{0, 1\}$ and let $p = Pr(T = \delta)$. For any r, $1 \le r \le n$

$$I_T^\delta(r) \ \ge \ \frac{r}{n} \cdot p \cdot \ln \frac{1}{p}.$$

In particular if $Pr(T = 0) = Pr(T = 1) = 1/2$, there are subsets S_0 and S_1 of N of cardinality r with

$$I_T^0(S_0) \ \ge \ c\,\frac{r}{n}$$
$$I_T^1(S_1) \ \ge \ c\,\frac{r}{n}$$

where $c = (ln\, 2)/2 \ = 0.34657\ldots$.

Proof: Our first observation is that for the purpose of lower bounds it is enough to consider only restricted schemes. This follows immediately from

Proposition 5.1: For any general coin flipping scheme (T, N) there is a restricted scheme $(\tilde{T}, N)$ such that $Pr(T = 0) = Pr(\tilde{T} = 0)$ and for every $S \subseteq N$ and $\delta = 0, 1$

$$I_T^\delta(S) \ge I_{\tilde{T}}^\delta(S).$$

Proof: The idea is very simple: Instead of assigning values to the variables of the choice function at each node simultaneously we do this sequentially. This can only reduce the influence of any set of players. Let u be an internal node of T and let $f_u : \prod X_i \longrightarrow V$ be its choice function. We replace this node by a tree T_u of $n + 1$ level as follows: All the nodes of T_u will have the given probability spaces X_i as their attached probability spaces. The root of T_u will be the node u and the set of its children will be the set X_1, The choice function attached to this root will be the function $f_u^1(x_1, \ldots, x_n) = x_1$. In the same manner, the set of children of all nodes at the i-th level of T_u will be the set X_i, with the choice function $f_u^i(x_1, \ldots, x_n) = x_i$. In other words, all the nodes of depth i are controlled by player i, and his action at this level is to select his assignment to f_u. With each leaf of T_u we can associate the n-tuple $(x_1, \ldots, x_n)$ according to the path that leads from the root u to the leaf. At this leaf we attach a copy of the node $v = f_u(x_1, \ldots, x_n) \in V$.

To construct $\tilde{T}$ we begin at the root u of T and replace it by the tree T_u. At each leaf of T_u that is labeled by $v \in V$ we put a copy of the the subtree rooted at the child v of u. We now proceed with each subtree in the same manner. This way $\tilde{T}$ is constructed.

Let $S \subseteq N$ be any subset of the players. In evaluating $I_T^\delta(S)$, at each node of T we set the the variables in S after the other variables have been set at random. In $\tilde{T}$ this may not be possible just because members of S may precede other players outside of S. Thus in scheme $\tilde{T}$ the players in S may have less strategies to chose from than in the scheme T. (In fact for $S = \{n - |S| + 1, \ldots, n\}$ they have exactly the same set of strategies to play.) Since any strategy for the players in $\tilde{T}$ can be used as a strategy for T it is clear that $I_T^\delta(S) \geq I_{\tilde{T}}^\delta(S)$ □

Our next observation is that it is enough to prove the theorem for Boolean schemes:

Proposition 5.2: Let T be a general n player coin flipping scheme and let $\varepsilon > 0$. Then there exists a Boolean scheme $\tilde{T}$ such that $|Pr(T = 0) - Pr(\tilde{T} = 0)| < \varepsilon$ and for every coalition S and $\delta \in \{0, 1\}$, we have

$$I_{\tilde{T}}^\delta(S) \leq I_T^\delta + \varepsilon.$$

Sketch of Proof: W.l.o.g. assume T is a restricted scheme. To construct $\tilde{T}$ simply approximate the random variable at each node of T by a dyadic approximation and in a similar manner to the proof of Proposition 5.1, replace the action of the player at this node by a sequence of coin flips. □

We now return to the proof of Theorem 5. Proposition 5.2 allows us to assume that T is a Boolean scheme. Consider first the influence of one player, (i.e. $r = 1$). For a node in the game tree T we consider the game H of the subtree below it. Let $a_i = 1 - q_0(H, \{i\})$ denote the largest probability that this game ends with a 1 under i's best strategy and let $a = Pr(H = 1)$ be this probability under random play by all the players. We prove

Lemma 5.3: For every node in a game tree

$$a_1 \cdots a_n \geq a^{n-1}.$$

Proof: We prove this by induction on the height in the tree. In a leaf all a_i and a are either zero or one. Let u be the father of v and w and say w.l.o.g. that u is controlled by player 1. We use b_i, c_i to denote the appropriate quantities at v, w. We have

$$b_1 \cdots b_n \geq b^{n-1},$$

$$c_1 \cdots c_n \geq c^{n-1}.$$

$$a_1 = \max(b_1, c_1), \quad \text{say} \quad a_1 = b_1.$$

$$a_i = \frac{1}{2}(b_i + c_i) \quad n \geq i \geq 2$$

$$a = \frac{1}{2}(b + c).$$

We wish to show

$$a_1 \cdots a_n \geq a^{n-1}.$$

That is

$$b_1\left(\frac{b_2 + c_2}{2}\right) \cdots \left(\frac{b_n + c_n}{2}\right) \geq \left(\frac{b + c}{2}\right)^{n-1}$$

or

$$b_1(b_2+c_2)\cdots(b_n+c_n) \geq (b+c)^{n-1}.$$

Expand the product $\prod_{i=2}^{n}(b_i+c_i)$ and consider the $\binom{n-1}{t}$ terms with t b-factors and $(n-t-1)$ c-factors. This gives

$$\sum_{\substack{A\subseteq\{2..n\}\\ |A|=t}} \prod_{i\in A} b_i \cdot \prod_{j\notin A} c_j$$

By the arithmetic - geometric inequality this is greater or equal to

$$\binom{n-1}{t}\Big[\sum_{\substack{A\subseteq\{2..n\}\\ |A|=t}} \prod_{i\in A} b_i \cdot \prod_{j\notin A} c_j\Big]^{1/\binom{n-1}{t}} =$$

$$= \binom{n-1}{t}\Big[\Big(\prod_{i=2}^{n} b_i\Big)^{\binom{n-2}{t-1}}\Big(\prod_{i=2}^{n} c_i\Big)^{\binom{n-2}{t}}\Big]^{1/\binom{n-1}{t}} =$$

$$= \binom{n-1}{t}\Big(\prod_{2}^{n} b_i\Big)^{t/(n-1)}\Big(\prod_{2}^{n} c_i\Big)^{(n-t-1)/(n-1)}$$

So we have

$$b_1(b_2+c_2)\cdots(b_n+c_n) \geq$$

$$\geq b_1 \sum_{t=0}^{n-1}\binom{n-1}{t}\Big(\prod_{2}^{n} b_i\Big)^{t/(n-1)}\Big(\prod_{2}^{n} c_i\Big)^{(n-t-1)/(n-1)}$$

$$\geq \sum_{t=0}^{n-1}\binom{n-1}{t}\Big(\prod_{1}^{n} b_i\Big)^{t/(n-1)}\Big(\prod_{1}^{n} c_i\Big)^{(n-t-1)/(n-1)} \geq$$

$$\geq \sum_{t=0}^{n-1}\binom{n-1}{t} b^t c^{n-t-1} = (b+c)^{n-1}.$$

Where $\prod b_i \geq b^{n-1}$ and $\prod c_i \geq c^{n-1}$ were used. □

The derivation of the theorem is easy now: We conclude from the lemma that at the root of T

$$\max a_i \geq a^{1-1/n} = a\,(1+\frac{\ln(1/a)}{n}+O(n^{-2})),$$

or

$$I_T^1 \;\geq\; \frac{p}{n}\ln\frac{1}{p}.$$

This completes the proof for $r=1$. The general case is treated in a similar way: Let $a_S = 1-q_0(T,S)$ for all $S\subseteq N$ of size r, and let $a = Pr(T=0)$, then by an argument similar to the one presented above we have

$$\prod_S a_S \geq a^{\binom{n-1}{r}}$$

and thus

$$\max_S a_S \;\geq\; a^{1-r/n} \;>\; a\,(1+\frac{r}{n}\ln(1/a)). \quad \square$$

Remark: Our lower bound shows that there is always a player that can bias the coin by $O(1/n)$ toward the value 1. A similar result holds of course if we are interested in bias towards 0. We note that this lower bound is optimal as for any p_i and p between 0 and 1 satisfying $\prod p_i = p^{n-1}$ and $p_i \geq p$ for all i, we can construct a scheme T such that $p = Pr(T=0)$ and $p_i = p + I_T^1(\{i\})$ as follows: Let $\alpha_i = p/p_i$. Each player i in N flips a biased coin with probability α_i of outcome 1. If all players get 1 the outcome of T is 1 and otherwise 0. Note that in this construction any player can bias the outcome towards 0, but the bias towards 1 is bounded. In the following section we give a construction where in influence of any player (towards 0 or 1) is only $O(1/n)$.

6. Multistage games - Upper Bounds

An *election scheme* (T, N) is defined like a coin flipping scheme, but unlike coin flipping schemes where the leaves of T are labeled by 0 or 1, in an election scheme they are marked by names of players from N. The game proceeds exactly like a coin flipping scheme and when a leaf is reached, the player whose name marks it is elected. We restrict our attention to election games where if everyone plays randomly each player is elected with equal probability. To measure the influence of a coalition S we assume that players outside S play randomly, while those in S play the best strategy to maximize the probability that a member of S is elected. The excess of this probability over $|S|/N$ is defined as the influence of S on the election scheme.

Here is an example of an election scheme $E = (T, N)$ where every player has zero influence. The root is controlled by player 1. Its $n-1$ sons are controlled by $2, \ldots, n$ respectively and 1 has to select between them with equal probability. The children of i's node are $n-1$ leaves and all marks appear there but for i. Player 1 is to be chosen with probability $1/n$ and all the rest with probability $(n-1)/n(n-2)$ each. The reader can easily check that under random play every player is elected with probability $1/n$ and a single player cannot increase his chance of being elected no matter what strategy he plays. The next theorem follows easily now.

Theorem 6.1: There are multistage games $T = T_n$ such that $p(T=0) = 1/2$ and the influence of any player x satisfies

$$I_T^1(\{x\}), I_T^0(\{x\}) \;=\; O(\frac{1}{n}).$$

Proof: The scheme may be described as follows: Run the election scheme E and have the elected player flip a fair coin. By the property of this election scheme a player may bias the coin only if he is elected. Since everyone is elected with probability $1/n$ the influence towards either zero or one are both $O(1/n)$. $\square$

As for the influence of larger sets of players we have the following

Theorem 6.2: There are n players multistage schemes $T = T_n$ such that $p(T=0) = 1/2$ and for all k, $k < n^{\alpha - o(1)}$, where $\alpha = \log_3 2 = 0.63\ldots$, we have

$$I_T^0(k), I_T^1(k) = O(\frac{k}{n}).$$

Proof: We assume $n = 2^r$ and let $k \le n^\alpha/2r$. Let G_0 be the n-player game of Theorem 4(a). Define the scheme $T = T_n$ to be the following: We play the game G_0 for $r = \log n$ times and let b_i be the i-th outcome. The sequence $b = b_1, \cdots, b_r$ identifies one of the n players. This selected player now flips his coin again to set the outcome of T. Note that the only way a set of players can bias the coin is by trying to have one of them selected to flip the final coin. As the influence of the k players on the G_0 game is bounded by $k/2n^\alpha$, the probability that one of these k players will be reached is bounded by

$$k\,(\frac{1}{2} + \frac{k}{2n^\alpha})^r \le \frac{k}{n}\,(1 + \frac{1}{2r})^r < \frac{2k}{n}$$

Thus the probability of outcome 1 (or 0) is at most

$$\frac{2k}{n} + \frac{1}{2}\,(1 - \frac{2k}{n}) = \frac{1}{2} + \frac{k}{n}$$

and so the influence of any k players is bounded by $O(k/n)$. □

7. Final Remarks and Open Questions

Below we list some open problems raised in this paper along with our conjectured answers. We always refer to f as an n variable Boolean function with $Pr(f = 0) = 1/2$ and T is always an n player multiround scheme with $Pr(T = 0) = 1/2$. Most of these questions deal with the minimization of the influence function, They are classified according to the following criteria:

- Single round / multi round scheme.
- The influence of a single player / the gain of ε control.
- The influence function considered. We deal with the following three quantities: I, $\min(I^0, I^1)$ and $\max(I^0, I^1)$.

1) As mentioned in Section 3 we conjecture that I_f is always $\Omega(\frac{\log n}{n})$. In other words, every Boolean function has an influential variable.

2) Given n and $\varepsilon > 0$, what is the least r such that for every f there holds

$$\max\,(I_f^0(r), I_f^1(r)) > \varepsilon?$$

We gave examples showing $r = \Omega(n^{0.63..})$, but we are not sure of the best bound.

3) We conjecture that for every f there is a set S of $O(\frac{n}{\log n})$ variables for which both

$$I_f^0(S), I_f^1(S) = \Omega(1).$$

In other words, every Boolean function has a negligible set of variables with a significant influence. In Section 3 this is shown to be best possible.

4) We showed in Theorem 5 that for every T both I_T^0 and I_T^1 are at least $\Omega(1/n)$ and this is tight. What is the largest $\psi = \psi(n)$ such that in every T there is a player x for which both

$$I_T^0(x), I_T^1(x) \;\ge\; \psi(n)$$

hold?

5) The claim of (3) is conjectured to hold also for multiround schemes. By results of Saks [S] this, if true, is best possible.

6) We want to describe another notion of *adaptive influence* over a scheme. Let T be a given scheme, let k be an integer and start with S an empty set. As we play the scheme T an adversary adds at most k players to S. The players currently in S play their best strategy towards δ and the rest play randomly. Let

$$A_k^{\delta}(T) := Pr(T \text{ ends with } \delta) - Pr(T = \delta).$$

For which T is this quantity minimized? Among the one round games this quantity is minimized for the majority function. This follows from an isoperimetric inequality in the cube [H]. A recent article of Lichtenstein, Linial and Saks [LLS] shows that the majority function remains optimal even after unfolding the one round game to a tree scheme of n levels. It may be that majority is the answer also without extra assumptions.

7) The connection between election games and coin flipping games seems very interesting. Consider coin flipping schemes where we first run an election game and have the elected player flip a coin. Is it true that the best coin flipping schemes have this form? Given a robust coin flipping scheme how can it be used to create a robust election scheme? The Vaziranis' recent results on sampling with slightly random sources [VV] are relevant but do not seem to answer this question.

Acknowledgement: We acknowledge useful discussions with A. Broder, N. Megiddo, A. Neyman, L. Stockmeyer and U. Vazirani.

References

[ACGM] Awerbuch B., Chor B., Goldwasser S. and Micali S., Constructive and Provably Fair Coin Flip in Byzantine Networks, Manuscript, 1984.

[BE] Ben-Or M., Fast Asynchronous Byzantine Agreement, 4th PODC, 1985.

[BD] Broder A. and Dolev D., Flipping Coins in Many Pockets, 25th FOCS, 1984.

[Br] Bracha G., An $O(\log n)$ Expected Rounds Randomized Byzantine Generals Algorithm, 17th STOC 1985.

[BR] Broder A., A Provably Secure Polynomial Approximation Scheme for the Distributed Lottery Problem, 4th PODC, 1985.

[DS] Dubey P. and Shapley L.S., Mathematical Properties of the Banzhaf Power Index, Math. Oper. Res. 4, 1979, 99-131.

[F] Frankel P., On The Trace of Finite Sets, Journal of Combinatorial Theory ser. A, Vol. 34, 1983, 41-45.

[H] Harper L., Optimal Numberings and Isoperimetric Problems on Graphs, Journal of Combinatorial Theory ser. A, 1966, 385-393.

[Ha] Hart S., A Note on the Edges of the n-Cube, Discrete Math., 14, 1976, 157-163.

[LLS] Lichtenstein D., Linial N. and Saks M., Some Extremal Problems Arising from Discrete Control Processes, to appear 19th STOC, 1987.

[Ow] Owen G., Game Theory, 2nd Ed., Academic Press, 1982.

[Ra] Rabin M.O., Randomized Byzantine Generals, 24th FOCS, 1983, 403-409.

[S] Saks M., private communication.

[VV] Vazirani U.V. and Vazirani V.V., Random Polynomial Time is Equal to Slightly Random Polynomial Time, 26th FOCS, 1985, 417-428.

[Ya] Yao A.C., On the Succession Problem for Byzantine Generals, TR, to appear, 1984.

[Wi] Widner R.O., Chow Parameters in Threshold Logic, JACM, Vol. 18, 1971, 265-289.