Analysis of Boolean Functions **(CMU 18-859S, Spring 2007)**

Lecture 4: Locally testing Dictatorship with NAE; explicit PCPPs

Jan. 25, 2005

Lecturer: Ryan O'Donnell *Scribe: Aaron Roth*

1 A Local Test for Dictatorship

The Marquis de Condorcet was a French mathematician and early political scientist. In addition to publishing influential works on integral calculus, in 1785 he published his *Essay on the Application of Analysis to the Probability of Majority Decisions*, in which he outlined a method of aggregating voter preferences into a ranking of candidates.

Definition 1.1 ***Condorcet Method For Ranking 3 Candidates:*** *In an election with n voters and 3 candidates A, B, and C, each voter submits 3 bits representing his preferences:*

(1) $A > B$

(2) $B > C$

(3) $C > A$.

These preferences are aggregated into 3 strings $s_1, s_2, s_3 \in \{-1, 1\}^n$. A boolean function $f : \{-1, 1\}^n \to \{-1, 1\}$ is applied to all 3 strings, and the aggregate preference is represented by the triple $(f(s_1), f(s_2), f(s_3))$.

Problem: "The Condorcet Paradox." If f is the majority function, we can have an irrational outcome, in which all 3 aggregate bits are 1 (or -1), representing a cyclic preference $A < B < C < A$ (or $A > B > C > A$).

Definition 1.2 *A triple $(a, b, c) \in \{-1, 1\}$ is rational if it corresponds to a non-cyclic ordering. Equivalently, it is rational if not all three bits are equal. We define the function $\mathrm{NAE} : \{-1, 1\}^3 \to \{-1, 1\}$ to be true if and only if its three input bits are not all equal.*

Theorem 1.3 ***Arrow's Impossibility Theorem:*** *The only monotone functions f that never give irrational outcomes are dictators. (Allowing non-monotone functions only adds anti-dictators.)*

Note: Arrow's Theorem is actually more general than this, but in particular it implies the above result. This result suggests a 3-query test for dictators: Arrow's Theorem tells us that dictators are the only functions that will always pass a test for rationality.

Remark 1.4 *While Condorcet, who was also a proponent of black and women's rights, was jailed and poisoned after the French revolution, Ken Arrow was awarded a Nobel prize in economics.*

Definition 1.5 ***"NAE Test"***

- *Choose $x, y, z \in \{-1,1\}^n$ by choosing the triples (x_i, y_i, z_i) uniformly from the set of 6 assignments such that all three bits are not equal, and independently across i's.*
- *Accept if* $\mathrm{NAE}(f(x), f(y), f(z))$*, and otherwise reject.*

Lemma 1.6

$$\Pr_{\boldsymbol{x},\boldsymbol{y},\boldsymbol{z}}[\textit{NAE}(f)\textit{ passes}] = \frac{3}{4} - \frac{3}{4}\sum_{S\subseteq[n]}\left(-\frac{1}{3}\right)^{|S|}\hat{f}(S)^2.$$

Sanity check: Is this lemma what we want?

- Dictators pass with probability $3/4 - 3/4((-1/3)^1 \cdot 1^2) = 3/4 + 1/4 = 1$.
- Constant functions pass with probability $3/4 - 3/4((-1/3)^0 \cdot 1^2) = 3/4 - 3/4 = 0$.
- Parities on 2 bits pass with probability $3/4 - 3/4((-1/3)^2 \cdot 1^2) = 3/4 - 1/12 = 2/3$.

So it looks like this lemma may indeed help us get a local test for dictatorship.

Proof: The following is easily seen to be a 0-1 indicator for the NAE predicate:

$$\mathbf{1}_{\mathrm{NAE}(a_1,a_2,a_3)} = \frac{3}{4} - \frac{1}{4}a_1a_2 - \frac{1}{4}a_1a_3 - \frac{1}{4}a_2a_3.$$

Remark 1.7 *When analyzing a general predicate, the indicator expression to use is nothing more than the Fourier expansion of the predicate,* $\{-1,1\}^q \to \{0,1\}$.

Hence:

$$\mathbf{Pr}[\text{NAE passes}] = \frac{3}{4} - \frac{1}{4}\mathbf{E}[f(\boldsymbol{x})f(\boldsymbol{y})] - \frac{1}{4}\mathbf{E}[f(\boldsymbol{x})f(\boldsymbol{z})] - \frac{1}{4}\mathbf{E}[f(\boldsymbol{y})f(\boldsymbol{z})].$$

We note that $(\boldsymbol{x}, \boldsymbol{y})$ has the same distribution as $(\boldsymbol{x}, \boldsymbol{z})$ and $(\boldsymbol{y}, \boldsymbol{z})$; in particular,

$$\mathbf{Pr}[(-1,-1)] = \mathbf{Pr}[(1,1)] = 1/6, \qquad \mathbf{Pr}[(-1,1)] = \mathbf{Pr}[(1,-1)] = 2/6. \tag{1}$$

Therefore, we have:

$$\mathbf{Pr}[\text{NAE passes}] = \frac{3}{4} - \frac{3}{4}\mathbf{E}[f(\boldsymbol{x})f(\boldsymbol{y})]$$

It therefore remains to show that $\mathbf{E}[f(\boldsymbol{x})f(\boldsymbol{y})] = \sum_{S\subseteq[n]}(-1/3)^{|S|}\hat{f}(S)^2$.

Expanding f in terms of its Fourier coefficients:

$$\mathbf{E}[f(\boldsymbol{x})f(\boldsymbol{y})] = \mathbf{E}[(\sum_{S\subseteq[n]}\hat{f}(S)\boldsymbol{x}_S)\cdot(\sum_{T\subseteq[n]}\hat{f}(T)\boldsymbol{y}_T)] = \sum_{S,T\subseteq[n]}\hat{f}(S)\hat{f}(T)\mathop{\mathbf{E}}_{\boldsymbol{x},\boldsymbol{y}}[\boldsymbol{x}_S\boldsymbol{y}_T]$$

We now write

$$\mathop{\mathbf{E}}_{\boldsymbol{x},\boldsymbol{y}}[\boldsymbol{x}_S \boldsymbol{y}_T] = \mathop{\mathbf{E}}_{\boldsymbol{x},\boldsymbol{y}}\left[\prod_{i\in S} \boldsymbol{x}_i \cdot \prod_{i\in T} \boldsymbol{y}_i\right] = \mathop{\mathbf{E}}_{\boldsymbol{x},\boldsymbol{y}}\left[\prod_{i=1}^{n} \boldsymbol{x}_i^{\mathbf{1}_{i\in S}} \boldsymbol{y}_i^{\mathbf{1}_{i\in T}}\right] = \prod_{i=1}^{n} \mathop{\mathbf{E}}_{\boldsymbol{x},\boldsymbol{y}}[\boldsymbol{x}_i^{\mathbf{1}_{i\in S}} \boldsymbol{y}_i^{\mathbf{1}_{i\in T}}]$$

The last equality holds by the mutual independence of the n random variable pairs (x_i, y_i). Writing things out using a product across all i and using indicators is a useful trick in general. Now note that:

If $i \notin S, i \notin T$, then $\mathbf{E}_{\boldsymbol{x},\boldsymbol{y}}[\boldsymbol{x}_i^{\mathbf{1}_{i\in S}} \boldsymbol{y}_i^{\mathbf{1}_{i\in T}}] = E[1] = 1$.
If $i \in S, i \notin T$, then $\mathbf{E}_{\boldsymbol{x},\boldsymbol{y}}[\boldsymbol{x}_i^{\mathbf{1}_{i\in S}} \boldsymbol{y}_i^{\mathbf{1}_{i\in T}}] = E[\boldsymbol{x}_i] = 0$.
If $i \notin S, i \in T$, then $\mathbf{E}_{\boldsymbol{x},\boldsymbol{y}}[\boldsymbol{x}_i^{\mathbf{1}_{i\in S}} \boldsymbol{y}_i^{\mathbf{1}_{i\in T}}] = E[\boldsymbol{y}_i] = 0$.
If $i \in S, i \in T$, then $\mathbf{E}_{\boldsymbol{x},\boldsymbol{y}}[\boldsymbol{x}_i^{\mathbf{1}_{i\in S}} \boldsymbol{y}_i^{\mathbf{1}_{i\in T}}] = \mathbf{E}[\boldsymbol{x}_i \boldsymbol{y}_i] = 2/6 - 4/6 = -1/3$, using (1).

Therefore we have that if $S \neq T$, $\mathbf{E}_{x,y}[\boldsymbol{x}_S \boldsymbol{y}_T] = 0$, and otherwise, $\mathbf{E}_{\boldsymbol{x},\boldsymbol{y}}[\boldsymbol{x}_S \boldsymbol{y}_T] = (-1/3)^{|S|}$. Thus we have shown that:

$$\mathop{\mathbf{E}}_{\boldsymbol{x},\boldsymbol{y}}[f(\boldsymbol{x}) f(\boldsymbol{y})] = \sum_{S,T} \hat{f}(S)\hat{f}(T) \mathop{\mathbf{E}}_{\boldsymbol{x},\boldsymbol{y}}[\boldsymbol{x}_S \boldsymbol{y}_T] = \sum_{S} \left(-\frac{1}{3}\right)^{|S|} \hat{f}(S)^2$$

which completes the proof. □

The following definition will simplify our notation:

Definition 1.8 *Given $f : \{-1,1\}^n \to \mathbb{R}$ we say the* weight at level k *is:*

$$W_k = W_k(f) = \sum_{|S|=k} \hat{f}(S)^2.$$

We note that for all k, $W_k \geq 0$, and for a function $f : \{-1,1\}^n \to \{-1,1\}$, $\sum_{k=0}^{n} W_k(f) = 1$.

Corollary 1.9

$$\mathbf{Pr}[\mathrm{NAE}(f) \text{ accepts}] = \frac{3}{4} - \frac{3}{4}W_0 + \frac{1}{4}W_1 - \frac{1}{12}W_2 + \frac{1}{36}W_3 - \ldots. \tag{2}$$

Hence if f passes with probability $\geq 1 - \epsilon$, then $W_1 \geq 1 - \frac{9}{2}\epsilon$.

Proof: Suppose to the contrary that $W_1 < 1 - \frac{9}{2}\epsilon$. Since $\sum_{k=0}^{n} W_k = 1$, it is easy to see that the most (2) could possibly then be is what one would get if the remaining weight occurred at level 3 (since these two levels contribute the largest positive weight towards the acceptance probability). But this would only achieve

$$\mathbf{Pr}[\mathrm{NAE}(f) \text{ accepts}] < \frac{3}{4} + \frac{1}{4}\left(1 - \frac{9}{2}\epsilon\right) + \frac{1}{36} \cdot \frac{9}{2}\epsilon = 1 - \frac{9}{8}\epsilon + \frac{1}{8}\epsilon = 1 - \epsilon,$$

a contradiction. □

One may now ask: Given that f has almost all its "Fourier weight" at level 1 — i.e., $\sum_{|S|=1} \hat{f}(S)^2 \geq 1 - O(\epsilon)$ — must f be close to a dictatorship or anti-dictatorship, as Arrow's Theorem suggests? It's not immediately clear. In fact, later in this course we will show that this *does* hold. This is a theorem of Friedgut, Kalai, and Naor:

Theorem 1.10 *If $f : \{-1,1\}^n \to \{-1,1\}$ has $W_1(f) \geq 1 - \epsilon$, then f is $O(\epsilon)$-close to either a dictator or an anti-dictator.*

We don't really need this theorem, though, to get a local test for dictatorship. We already have all the tools we need:

Theorem 1.11 *Dictators are locally testable with 6 queries.*

Proof: We run both the BLR linearity test and the NAE test on f, and accept if and only if both tests accept.

First note that if f is a dictator then it passes both tests with probability 1 (in particular, it passes BLR because dictators are parities). Now suppose that the overall test passes with probability at least $1 - \epsilon$. Then the BLR test passes with probability at least $1 - \epsilon$, and so there exists a unique set S^* such that $\hat{f}(S^*) \geq 1 - 2\epsilon$. (That S^* is unique follows from Parseval, $\sum \hat{f}(S)^2 = 1$, and because we can assume ϵ to be sufficiently small.) We also have that the NAE test passes with probability at least $1 - \epsilon$, and so we have $\sum_{|S|=1} \hat{f}(S)^2 \geq 1 - \frac{9}{2}\epsilon$. Therefore, it must be that $|S^*| = 1$, because otherwise, $\sum_S \hat{f}(S)^2 \geq 1 - \frac{9}{2}\epsilon + (1 - 2\epsilon)^2 > 1$ (using that ϵ is sufficiently small). Hence, $\langle f, \chi_{S^*} \rangle \geq 1 - 2\epsilon$ and so f is ϵ-close to to a dictator, completing the proof. □

Remark 1.12 *A note on our assumption that ϵ is sufficiently small: We implicitly used here that $\epsilon < 1/8$ to show that if f is ϵ-far from being a dictator then the probability of acceptance is at most $1 - \epsilon$. Now if f is ϵ-far for some $\epsilon \geq 1/8$ then we can't conclude that the acceptance probability is at most $1 - \epsilon$; but, we* can *conclude it's at most $1 - 1/8$ (since f being $> 1/8$-far implies it is $1/8$-far). Hence for all $0 \leq \epsilon \leq 1$, we may say that any f that is ϵ-far from being a dictator passes with probability at most $1 - \epsilon/8$, and this statement is enough to satisfy the definition of local testability.*

In fact, we can do slightly better using a general trick: If we have a test that consists of multiple sub-tests $(T_1, T_2, \ldots, T_k)$, we don't need to do each test in series. We can instead randomly conduct one test T_i for $i \in [k]$ drawn uniformly at random. This allows us to use only the maximum number of queries necessary for any of the sub-tests, rather than their sum.

Theorem 1.13 *Dictatorship is locally testable with only 3 queries.*

Proof: We flip a coin and with probability $1/2$ perform the BLR test, and otherwise perform the NAE test. If this passes with probability $\geq 1 - \epsilon/2$, then both sub-tests must pass with probability $\geq 1 - \epsilon$, and we proceed as before. We lose only a constant to the $\Omega(\cdot)$ in the overall $1 - \Omega(\epsilon)$ rejection probability. □

Corollary 1.14 *For every subset $P \subseteq \{1, \ldots, n\}$, the property $\{\chi_{\{i\}} : i \in P\}$ is locally testable.*

Proof:

(1) Perform the BLR+NAE test, rejecting if it rejects. If it passes with probability at least $1 - \epsilon$, then we know that f is $O(\epsilon)$-close to some dictator.

(2) Construct the string $x \in \{-1, 1\}^n$, where $x_i = -1$ iff $i \in P$. Then locally decode f on x.

(3) Accept if and only if the local decoding yields -1.

It is easy to verify that this accepts dictators in P with probability 1 and rejects functions ϵ-far from being dictators in P with probability at least $\Omega(\epsilon)$. □

Note that we can do this in three queries by using our sub-test trick.

2 Probabilistically Checkable Proofs of Proximity

We will now get to see the great usefulness of locally testing dictators — specifically locally testing subsets of Dictators. We will show that "*every* property is locally testable with 3 queries" — if we are allowed to see a proof.

Definition 2.1 *A* string tester T *is a randomized algorithm with black-box query access to a string $w \in \{-1, 1\}^m$. It can query any coordinate $i \in [m]$ and get w_i. A string tester for a property $P \subseteq \{-1, 1\}^m$ distinguishes whether $w \in P$ or w is ϵ-far from P (i.e., $\Delta(w, v) \geq \epsilon m$ for all $v \in P$).*

String testing is more general than function testing, since you can think of a function as the string that represents its truth table. If $m = 2^n$ for some n then the reverse is also true, as you can identify an m-bit string with a boolean function on n bits.

Suppose that you are given a long proof that a string s has property P, but: (a) you don't trust the prover, and (b) the prover will only let you see a few bits of the proof. (However, you can convince the prover to write the proof in your favorite format.) We make the following definition, made independently by Ben-Sasson, Goldreich, Harsha, Sudan, and Vadhan and by Dinur and Reingold:

Definition 2.2 *A property $\mathcal{P}$ of m-bit strings has* probabilistically checkable proofs of proximity (PCPPs) of length $l(m)$ *if there is some string tester T making $O(1)$ nonadaptive queries to $w \in \{-1, 1\}^m$ and a purported "proof of $w \in \mathcal{P}$," $\pi \in \{-1, 1\}^{l(m)}$ such that:*

(1) If $w \in \mathcal{P}$ then $\exists \pi_w$ such that $\mathbf{Pr}[T(w, \pi_w) \text{ accepts}] = 1$.

(2) If w is ϵ-far from $\mathcal{P}$ then $\forall \pi \in \{-1, 1\}^{l(m)}$, $\mathbf{Pr}[T(w, \pi) \text{ accepts}] < 1 - \Omega(\epsilon)$.

Theorem 2.3 *Every property $\mathcal{P}$ of m-bit strings has PCPPs of length 2^{2^m}, where the tester makes only 3 queries.*

Remark 2.4 *Although proofs of size 2^{2^m} may seem bad, keep in mind that there are 2^{2^m} properties of m-bit strings.*

Remark 2.5 *PCPPs are also called* assignment testers *or* assisted tests*.*

Proof: Given $w \in \mathcal{P}$, we will generate the "correct proof of $w \in \mathcal{P}$", π_w as follows:

1. We will identify strings in $\{-1,1\}^m$ with the set $\{1,2,\ldots,n\}$, where $n := 2^m$ (via lexicographical order, say).
2. In particular, say that w gets identified with $t \in [n]$.
3. Further, say that the property $\mathcal{P}$ gets identified with the subset $P \subseteq [n]$.
4. We let π_w be the truth table of $\chi_{\{t\}} : \{-1,1\}^n \to \{-1,1\}$, the dictator on the variable associated with w.

We will further identify all purported proofs (which are strings of length 2^{2^m}) with functions $f : \{-1,1\}^n \to \{-1,1\}$ (recall that $n = 2^m$).

Our tester now has to check three things:
(a) The proof f should be a dictator function.
(b) This dictator function is on a coordinate that is in P.
(c) This coordinate actually corresponds to given strings $w \in \{-1,1\}^m$.

To do this, we can perform the following test:

1. Run the BLR+NAE test on f. If this passes with probability $\geq 1-\epsilon$, then f is $O(\epsilon)$-close to some dictator function $\chi_{\{u\}}$.
2. Run local decoding on f using the string y such that $y_i = -1 \Leftrightarrow i \in P$ to verify that $u \in P$.
3. Choose $i \in [m]$ uniformly at random and create $x \in \{-1,1\}^n$:

$$x = \underbrace{(-1,-1,1,1,-1,\ldots,1)}_{\text{-1 in the } v\text{th coordinate iff } v_i=-1}.$$

   Here we are viewing v in two ways: $v \in [n]$, and $v \in \{-1,1\}^m$. Use local decoding to find $f(x) = \chi_{\{u\}}(x) = u_i$. Then additionally query w_i, and accept iff $w_i = u_i$.

It is straightforward to check that this tester works as desired. In addition, we can make it use 3 queries — sub-test 1 requires 3 queries (using the trick), sub-test 2 requires 2 queries, and sub-test 3 requires 3 queries; we can then use the trick again to make the number of queries $\max(3,2,3) = 3$. □