Penn State University
CSE 564 Complexity of Combinatorial Problems

Instructor: Pei Wu

Module 7

Probabilistically checkable proofs (II): the BLR linearity test

The previous lecture introduced probabilistically checkable proofs (PCPs) and explained why the PCP theorem has strong consequences for hardness of approximation. A recurring technical theme in PCP constructions is that the verifier wants to check that a long proof string behaves like a codeword of some error-correcting code. This motivates the study of *property testing* problems in which the verifier has oracle access to a function and must distinguish the case that the function has a desired algebraic structure from the case that it is "far" from every structured function.

This lecture focuses on the simplest and most fundamental example: testing whether a Boolean function is linear over $\mathbb{F}_2$. This is exactly the consistency check needed for the Hadamard code. We present the celebrated Blum–Luby–Rubinfeld (BLR) *linearity test* and give two soundness proofs. The first uses Fourier analysis on $\{0,1\}^n$ and explains the test in terms of the Fourier spectrum. The second is a combinatorial "majority decoding" argument that also hints at the local self-correction property of the Hadamard code.

7.1 The BLR linearity test

We view $\{0,1\}^n$ as the vector space $\mathbb{F}_2^n$. Throughout, addition $x + y$ denotes bitwise XOR.

Definition 7.1 (linear functions). A function $\ell : \{0,1\}^n \to \{0,1\}$ is *linear* if $\ell(0) = 0$ and $\ell(x) + \ell(y) = \ell(x + y)$ for all $x, y \in \{0,1\}^n$, where addition on the right-hand side is in $\mathbb{F}_2$. Equivalently, there exists $a \in \{0,1\}^n$ such that $\ell(x) = \langle a, x \rangle \pmod 2$, where $\langle a, x \rangle = \sum_i a_i x_i$ is the $\mathbb{F}_2$ inner product.

Definition 7.2 (distance and "far"). For functions $f, g : \{0,1\}^n \to \{0,1\}$, define the (relative) Hamming distance

$$\Delta(f, g) := \Pr_{x \in \{0,1\}^n}[f(x) \neq g(x)].$$

Let LIN denote the set of linear functions $\{0,1\}^n \to \{0,1\}$. We say that f is ε-*far* from linear if $\Delta(f, \mathrm{LIN}) := \min_{\ell \in \mathrm{LIN}} \Delta(f, \ell) \geq \varepsilon$.

DEFINITION 7.3 (BLR test). Given oracle access to $f : \{0,1\}^n \to \{0,1\}$, the *BLR linearity test* is:

1. Sample $x, y \in \{0,1\}^n$ independently and uniformly at random.
2. Query $f(x), f(y), f(x+y)$.
3. **Accept** iff $f(x) + f(y) = f(x+y)$ in $\mathbb{F}_2$.

If f is linear, the test always accepts (perfect completeness). The main content is soundness: if f is far from linear, then the test rejects with noticeable probability.

THEOREM 7.4 (BLR soundness). *If f is ε-far from linear, then the BLR test rejects with probability $\Omega(\varepsilon)$. Equivalently, if the BLR test accepts with probability at least $1 - \varepsilon$, then there exists a linear function ℓ with $\Delta(f, \ell) \le \varepsilon$.*

We next prove this theorem in two different ways.

7.2 Fourier analysis preliminaries

For Fourier analysis it is convenient to work with $\{\pm 1\}$-valued functions. Given $f : \{0,1\}^n \to \{0,1\}$, define its ± 1 encoding by

$$F(x) := (-1)^{f(x)} \in \{\pm 1\}.$$

Under this encoding, the BLR acceptance condition $f(x) + f(y) = f(x+y)$ becomes $F(x)F(y) = F(x+y)$, or equivalently $F(x)F(y)F(x+y) = 1$.

For each subset $S \subseteq [n]$, define the Fourier character

$$\chi_S(x) := (-1)^{\sum_{i \in S} x_i} = \prod_{i \in S} (-1)^{x_i}.$$

If we identify S with its indicator vector $a \in \{0,1\}^n$, this is also $\chi_S(x) = (-1)^{\langle a, x \rangle}$.

DEFINITION 7.5 (Fourier coefficients). For a function $h : \{0,1\}^n \to \mathbb{R}$, its Fourier coefficient at S is

$$\widehat{h}(S) := 2^{-n} \sum_{x \in \{0,1\}^n} h(x) \chi_S(x).$$

The collection $\{\widehat{h}(S)\}_{S \subseteq [n]}$ is the *Fourier spectrum*.

PROPOSITION 7.6 (basic Fourier facts). *Let $\langle f, g \rangle := \sum_{x \in \{0,1\}^n} f(x) g(x)$ denote the unnormalized inner product. Then:*

1. ***Orthogonality:*** *$\langle \chi_S, \chi_T \rangle = 0$ for $S \neq T$, and $\langle \chi_S, \chi_S \rangle = 2^n$.*
2. ***Coefficient formula:*** *$\widehat{h}(S) = 2^{-n} \langle \chi_S, h \rangle$.*
3. ***Parseval/Plancherel:*** *$2^{-n} \langle f, g \rangle = \sum_{S \subseteq [n]} \widehat{f}(S)\, \widehat{g}(S)$. In particular, if $f : \{0,1\}^n \to \{\pm 1\}$ then $\sum_S \widehat{f}(S)^2 = 1$.*

EXAMPLE 7.7 (Fourier expansion of a simple Boolean function). Consider the OR function on $\{\pm 1\}^n$ defined by $\mathrm{OR}(\widetilde{x}) = 1$ iff $\widetilde{x} = 1^n$ and $\mathrm{OR}(\widetilde{x}) = -1$ otherwise. As a polynomial over the reals,

$$\mathrm{OR}(\widetilde{x}) = 2\prod_{i=1}^{n}\frac{1+\widetilde{x}_i}{2} - 1 = 1 + \sum_{\emptyset \neq S \subseteq [n]} 2^{1-|S|}\prod_{i\in S}\widetilde{x}_i.$$

If $\widetilde{x}_i = (-1)^{x_i}$ for $x \in \{0,1\}^n$, then $\prod_{i\in S}\widetilde{x}_i = \chi_S(x)$. So this is exactly a Fourier expansion.

Fourier analysis makes linear functions easy to recognize. If $\ell(x) = \langle a, x\rangle + b \pmod 2$ is an affine function, then its ± 1 encoding is $L(x) = (-1)^{\langle a,x\rangle + b}$. The next calculation (from the notes) shows that its Fourier spectrum is supported on a single character.

FACT 7.8 (Fourier spectrum of a linear/affine function). *Let $L(x) = (-1)^{\langle a,x\rangle+b}$. Let $S \subseteq [n]$ be the set corresponding to a. Then $\widehat{L}(S) = (-1)^b$ and $\widehat{L}(T) = 0$ for all $T \neq S$.*

Proof. Using the coefficient formula,

$$\widehat{L}(S) = 2^{-n}\sum_x (-1)^{\langle a,x\rangle+b}\chi_S(x) = 2^{-n}\sum_x (-1)^{\langle a,x\rangle+b}(-1)^{\langle a,x\rangle} = 2^{-n}\sum_x (-1)^b = (-1)^b.$$

For $T \neq S$, orthogonality gives $\widehat{L}(T) = 2^{-n}\langle L, \chi_T\rangle = 2^{-n}(-1)^b\langle \chi_S, \chi_T\rangle = 0$. ∎

7.3 Fourier-analytic proof of BLR

Fix $F : \{0,1\}^n \to \{\pm 1\}$. Suppose the BLR test accepts with probability at least $1-\varepsilon$, i.e.

$$\Pr_{x,y}[F(x)F(y) = F(x+y)] \geq 1-\varepsilon.$$

Equivalently,

$$\Pr_{x,y}[F(x)F(y)F(x+y) = 1] \geq 1-\varepsilon.$$

Since the random variable $F(x)F(y)F(x+y)$ takes values in $\{\pm 1\}$, the above implies

$$\mathbb{E}_{x,y}[F(x)F(y)F(x+y)] \geq 1-2\varepsilon. \tag{7.1}$$

Define the (normalized) convolution of two functions $g, f : \{0,1\}^n \to \mathbb{R}$ by

$$(g * f)(x) := \mathbb{E}_{y\in\{0,1\}^n}\big[g(y)f(x+y)\big].$$

Then the left-hand side of (7.1) can be rewritten as

$$\mathbb{E}_x\Big[F(x)\,\mathbb{E}_y[F(y)F(x+y)]\Big] = \mathbb{E}_x\big[F(x)\,(F*F)(x)\big].$$

FACT 7.9 (Fourier of convolution). *For all g, f and all $S \subseteq [n]$, $\widehat{g*f}(S) = \widehat{g}(S)\,\widehat{f}(S)$.*

Proof from the notes. Expand g and f in the Fourier basis: $g(y) = \sum_S \widehat{g}(S)\chi_S(y)$ and $f(x+y) = \sum_T \widehat{f}(T)\chi_T(x+y)$. Using $\chi_T(x+y) = \chi_T(x)\chi_T(y)$, we obtain

$$\begin{aligned}(g * f)(x) &= \mathbb{E}_y\Big[\sum_S \widehat{g}(S)\chi_S(y)\ \sum_T \widehat{f}(T)\chi_T(x)\chi_T(y)\Big] \\ &= \sum_{S,T} \widehat{g}(S)\widehat{f}(T)\chi_T(x)\,\mathbb{E}_y[\chi_{S\triangle T}(y)].\end{aligned}$$

The expectation $\mathbb{E}_y[\chi_{S\triangle T}(y)]$ equals 0 unless $S = T$, in which case it equals 1. Hence $(g*f)(x) = \sum_S \widehat{g}(S)\widehat{f}(S)\chi_S(x)$, so $\widehat{g*f}(S) = \widehat{g}(S)\widehat{f}(S)$. ∎

Apply Parseval with $f = F$ and $g = F * F$:

$$\mathbb{E}_x[F(x)(F*F)(x)] = \sum_S \widehat{F}(S)\,\widehat{F*F}(S) = \sum_S \widehat{F}(S)^3,$$

where the last equality uses the convolution fact with $g = f = F$. Combining with (7.1),

$$\sum_{S\subseteq[n]} \widehat{F}(S)^3 \geq 1 - 2\varepsilon. \tag{7.2}$$

Since $\sum_S \widehat{F}(S)^2 = 1$, we can upper bound the left-hand side of (7.2) by

$$\sum_S \widehat{F}(S)^3 \leq \Big(\max_T \widehat{F}(T)\Big)\sum_S \widehat{F}(S)^2 = \max_T \widehat{F}(T).$$

Therefore there exists S^* such that $\widehat{F}(S^*) \geq 1 - 2\varepsilon$.

Finally, note that $\widehat{F}(S) = \mathbb{E}_x[F(x)\chi_S(x)]$. If F and χ_{S^*} are ±1-valued, then

$$\Pr_x[F(x) = \chi_{S^*}(x)] = \frac{1 + \mathbb{E}_x[F(x)\chi_{S^*}(x)]}{2} = \frac{1 + \widehat{F}(S^*)}{2} \geq 1 - \varepsilon.$$

Thus F agrees with a character on all but an ε fraction of inputs. Translating back to $\{0,1\}$-valued functions, f is ε-close to the corresponding linear function. This proves BLR soundness.

7.4 Combinatorial proof and majority decoding

We now sketch the combinatorial proof from the notes. Here it is more convenient to use $\{0,1\}$-valued functions and write all additions in $\mathbb{F}_2$. Assume the BLR test accepts f with probability $1-\varepsilon$, i.e. $\Pr_{x,y}[f(x)+f(y)=f(x+y)] = 1-\varepsilon$.

Define a new function $g : \{0,1\}^n \to \{0,1\}$ by *majority decoding*: for each x, let $g(x)$ be the plurality value of $f(y) + f(y+x)$ over random y. Equivalently,

$$g(x) \in \arg\max_{v\in\{0,1\}} \big|\{y \in \{0,1\}^n : f(y) + f(y+x) = v\}\big|.$$

Claim 7.10. $\Pr[\textit{BLR rejects } f] \geq \frac{1}{2}\,\Delta(f,g)$.

Proof. Fix x. The BLR test (with this x) accepts iff $f(y)+f(x+y)=f(x)$. If $f(x)\neq g(x)$, then by definition of $g(x)$ the event $f(y)+f(x+y)=g(x)$ occurs with probability at least $1/2$ over random y. On that event, the test must reject because $g(x)\neq f(x)$. Therefore

$$\Pr[\text{reject}] \geq \Pr_x[f(x)\neq g(x)]\cdot\frac{1}{2} = \frac{1}{2}\,\Delta(f,g).$$

□

The next claim says that for any fixed direction x, the "derivative" $f(y)+f(y+x)$ is usually consistent across different base points y.

Claim 7.11. *For every fixed $x\in\{0,1\}^n$,*

$$\Pr_{y,z}\big[f(y)+f(y+x)=f(z)+f(z+x)\big]\geq 1-2\varepsilon.$$

Proof from the notes. Fix x. Define the set T of pairs (y,z) such that both of the following BLR constraints hold:

$$f(y)+f(y+z)=f(z), \qquad f(x+y)+f(y+z)=f(x+z).$$

Each displayed equality is of the BLR form $f(a)+f(b)=f(a+b)$ for a uniformly random pair (a,b): namely $(a,b)=(y,y+z)$ in the first and $(a,b)=(x+y,y+z)$ in the second. Therefore each fails with probability at most ε, and by the union bound $\Pr_{y,z}[(y,z)\in T]\geq 1-2\varepsilon$.

If $(y,z)\in T$, then adding the two equalities gives

$$f(y)+f(x+y)=f(z)+f(x+z),$$

which is exactly the desired conclusion. □

Now fix x and write $p:=\Pr_y[f(y)+f(y+x)=g(x)]$. By definition of majority, $p\geq 1/2$. Since the quantity $f(y)+f(y+x)$ is a single bit, we have

$$\Pr_{y,z}[f(y)+f(y+x)=f(z)+f(z+x)] = p^2+(1-p)^2.$$

Combining with the previous claim yields $p^2+(1-p)^2\geq 1-2\varepsilon$. Rearranging gives $p(1-p)\leq\varepsilon$. Since $p\geq 1/2$, this implies $1-p\leq 2\varepsilon$, i.e.

$$\Pr_y[f(y)+f(y+x)=g(x)]\geq 1-2\varepsilon \qquad \text{for all } x. \tag{7.3}$$

Claim 7.12. *If $\varepsilon<\frac{1}{6}$, then g is linear.*

Proof from the notes. Fix $x,z\in\{0,1\}^n$. Consider the following three events over a random y:

$$\begin{aligned} &f(y)+f(y+x)=g(x),\\ &f(y)+f(y+z)=g(z),\\ &f(y+x)+f(y+z)=g(x+z). \end{aligned}$$

Each event has probability at least $1-2\varepsilon$ by (7.3) (the third event is (7.3) for shift $x+z$ and base point $y'=y+x$, which is uniform when y is). By the union bound, all three events

hold simultaneously with probability at least $1 - 6\varepsilon$. If $\varepsilon < 1/6$, there exists a y for which all three equalities hold. Adding them in $\mathbb{F}_2$ cancels all $f(\cdot)$ terms (each appears twice), leaving $g(x) + g(z) + g(x + z) = 0$, i.e. $g(x) + g(z) = g(x + z)$. Since this holds for all x, z, g is linear. □

Combining linearity of g with the first claim gives a distance bound. If the BLR test rejects with probability ε, then $\Delta(f, g) \leq 2\varepsilon$, and (for $\varepsilon < 1/6$) this shows that f is $O(\varepsilon)$-close to a linear function. Taking contrapositives yields the BLR soundness statement.