Making the Long Code shorter, with applications to the Unique Games Conjecture

Boaz Barak – MSR New England

Joint work with
Parikshit Gopalan (MSR-Silicon Valley)
Johan Håstad (KTH)
Raghu Meka (IAS)
Prasad Raghavendra (GA Tech)
David Steurer (MSR New England)

Unique Games Conjecture [Khot'02]

- Hardness of certain constraint satisfaction problem.
- If true has many implications

Fastest algorithm for UG: $\exp(n^{\epsilon})$ time. [Arora-B-Steurer'10]

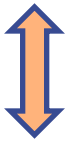

Huge Gap! (even bigger than it seems)

Best evidence for hardness: $n^{\log\log^{\delta} n}$ ... $n^{\log^{\delta} n}$ lower bounds in certain computational models.

[Khot-Vishnoi'04,Khot-Saket'09,Raghavendra-Steurer'09]

(capturing the ABS algorithm [B-Raghavendra-Steurer'11,Guruswami-Sinop'11])

Source of quantitative weakness:
Inefficiency of long-code/noise graph.

Unique Games Conjecture [Khot'02]

- Hardness of certain constraint satisfaction problem.
- If true has many implications

Fastest algorithm for UG: $\exp(n^{\epsilon})$ time. [Arora-B-Steurer'10]

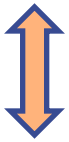

Huge Gap!

Best evidence for hardness: $n^{\log\log^{\delta} n}$ ... $n^{\log^{\delta} n}$ lower bounds in certain computational models.

[Khot-Vishnoi'04,Khot-Saket'09,Raghavendra-Steurer'09]
(capturing the ABS algorithm [B-Raghavendra-Steurer'11,Guruswami-Sinop'11])

Source of quantitative weakness:
Inefficiency of long-code/noise graph.

Fastest algorithm for UG: $\exp(n^{\epsilon})$ time.

Best evidence for hardness: $n^{\log\log^{\delta} n}$ ... $n^{\log^{\delta} n}$ restricted model lower bounds

Source of quantitative weakness: Inefficiency of long-code/noise graph.

Our main result: New exponentially more efficient replacement to long-code/noise-graph gadget.

Some Applications:

- Existence of small-set expander with $\exp(\log^{\epsilon} n)$ eigenvalues close to 1.

  Demostrates ABS algorithm can take $\exp(poly^{*}(n))$ time.

- Alphabet reduction / inner PCP gadget* for unique games with quasipolynomial blowup in alphabet size.

  Improves on exponential blowup of long-code gadget.

- Max Cut integrality gap for $poly^{*}(\log n)$ ... $poly^{*}(n)$ rounds of semi definite hierarchies.

  Exponential improvement over prior bounds.

Talk Overview

1. Review of long-code + noisy Boolean cube
2. Shorter long-code via "succinct cube":
   - Connection to locally testable codes.
   - Construction via Reed-Muller code.
3. Small set expander with many large eigenvalues.
4. Sketch of other applications:
   - "Majority is Stablest" for succinct cubes.
   - Sketch of integrality gap construction.
5. Conclusions and open problems

The Long Code and the Noisy Cube

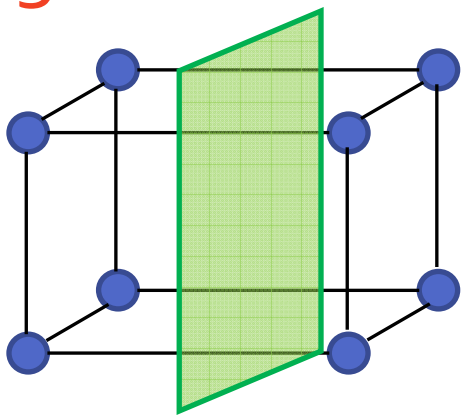

Noisy Cube Graph $T_{1-\epsilon}$: Vertices $\{0,1\}^N$, Edges: $x \sim y$ if $\Delta(x,y) = \epsilon N$

Affinity of subset of vertices $f: \{0,1\}^N \to \{0,1\}$ is

$$\varphi(f) = \Pr_{x \in f, y \sim x}[y \in f] = \langle f, T_{1-\epsilon} f \rangle / \| f \|^2 = \mathbb{E}_{x \sim y}[f(x) f(y)] / \mu(f)$$

Balanced sets w/ max affinity $(= 1 - \epsilon)$ are dictatorships: $f(x) = x_i$

Every set with large affinity is "close" to dictatorship:

f in span of e-vectors of value $\geq (1-2\epsilon)^d$ iff f is degree-d poly.

Stronger results known

[Kahn-Kalai-Linial'88, Friedgut'98, Bourgain'99,Mossel-O'Donnel-Oleszkiewicz'05],..

Corollary: $T_{1-\epsilon}$ is a small set expander: $\mu(f) \to 0 \Rightarrow \varphi(f) \to 0$

The Long Code and the Noisy Cube

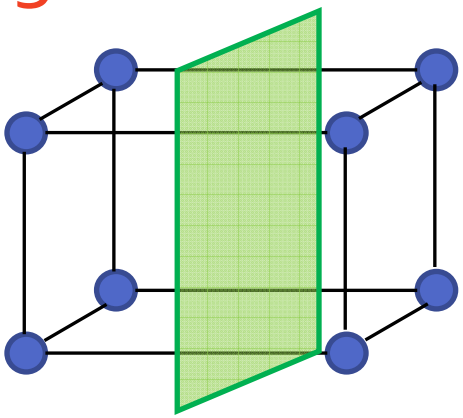

Noisy Cube Graph $T_{1-\epsilon}$: Vertices $\{0,1\}^N$, Edges: $x \sim y$ if $\Delta(x,y) = \epsilon N$

Affinity of subset of vertices $f: \{0,1\}^N \to \{0,1\}$ is

$$\varphi(f) = \langle f, T_{1-\epsilon} f \rangle / \| f \|^2 = \mathbb{E}_{x \sim y}[f(x)f(y)]/\mu(f) = \Pr_{x \in f, y \sim x}[y \in f]$$

Balanced sets w/ max affinity $(= 1 - \epsilon)$ are dictatorships: $f(x) = x_i$

Every set with large affinity is "close" to dictatorship:

f in span of e-vectors of value $\geq (1-2\epsilon)^d$ iff f is degree-d poly.

Stronger results known

[Kahn-Kalai-Linial'88, Friedgut'98, Bourgain'99,Mossel-O'Donnel-Oleszkiewicz'05],..

Corollary: $T_{1-\epsilon}$ is a small set expander: $\mu(f) \to 0 \Rightarrow \varphi(f) \to 0$

The Long Code and the Noisy Cube

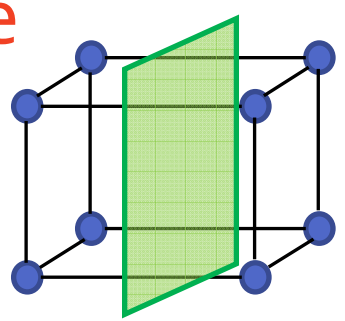

Noisy cube: Graph $T_{1-\epsilon}$ Vertices $\{0,1\}^N$, edges: $x \sim y$ if $\Delta(x,y) = \epsilon N$

Affinity of $f:\{0,1\}^N \to \{0,1\}$ is $\varphi(f) = \Pr_{x \in f, y \sim x}[y \in f]$

Max affinity balanced sets are dictatorships: $f(x) = x_i$

Any large affinity set is "close" to dictatorships, small sets have small affinity.

Use of Noisy Cube in unique-games applications:

- Replace variable taking values in $[N]$ with 2^N vertices of the noisy cube.
- Intention: encode $i \in [N]$ with the dictatorship $f(x) = x_i$
- Place edges according to the noisy cube.
- Any f with high value can be decoded to (list of) dictatorship.

Quantitative issue: 2^N blowup.

The Long Code and the Noisy Cube

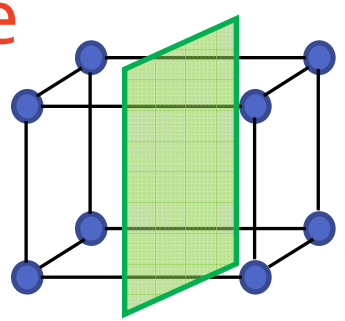

Noisy cube: Graph $T_{1-\epsilon}$ Vertices $\{0,1\}^N$, edges: $x \sim y$ if $\Delta(x,y) = \epsilon N$

Affinity of $f:\{0,1\}^N \to \{0,1\}$ is $\varphi(f) = \Pr_{x \in f, y \sim x}[y \in f]$

Max affinity balanced sets are dictatorships: $f(x) = x_i$

Any large affinity set is "close" to dictatorships, small sets have small affinity.

Use of Noisy Cube in unique-games applications:

- Replace variable taking values in $[N]$ with 2^N vertices of the noisy cube.
- Intention: encode $i \in [N]$ with the dictatorship $f(x) = x_i$
- Place edges according to the noisy cube.
- Any f with high value can be decoded to (list of) dictatorship.

Quantitative issue: 2^N blowup.

The Long Code and the Noisy Cube

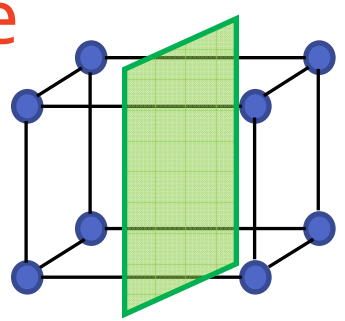

Noisy cube: Graph $T_{1-\epsilon}$ Vertices $\{0,1\}^N$, edges: $x \sim y$ if $\Delta(x,y) = \epsilon N$

Affinity of $f:\{0,1\}^N \to \{0,1\}$ is $\varphi(f) = \Pr_{x\in f, y\sim x}[y \in f]$

Max affinity balanced sets are dictatorships: $f(x) = (1-x_i)/2$

Any large affinity set is "close" to dictatorships, small sets have small affinity.

Quantitative issue: 2^N blowup.

Another manifestation: Noisy cube is small set expander with $\log|V|$ eigenvalues close to 1.

[Arora-B-Steurer'10]: Every small set expander has at most $|V|^{\epsilon}$ eigenvalues larger than $1-\epsilon^3$
$\Rightarrow \exp(n^{\epsilon})$-time algorithm for small set expansion (and unique games)

Question: Is there small set expander with $\log^{\omega(1)}|V|$ large e-vals?

This work: Construction with $\exp(\log|V|^{\delta})$ large e-vals.

Succinct Cubes

Goal: Construct small set expander with $\log^{\omega(1)} |V|$ large e-vals

Idea: Find induced subgraph G of the cube with $K \ll 2^N$ vertices s.t.

1) N dictatorships are still large eigenvectors

2) Subgraph still small set expander.

More generally: preserve properties of cube

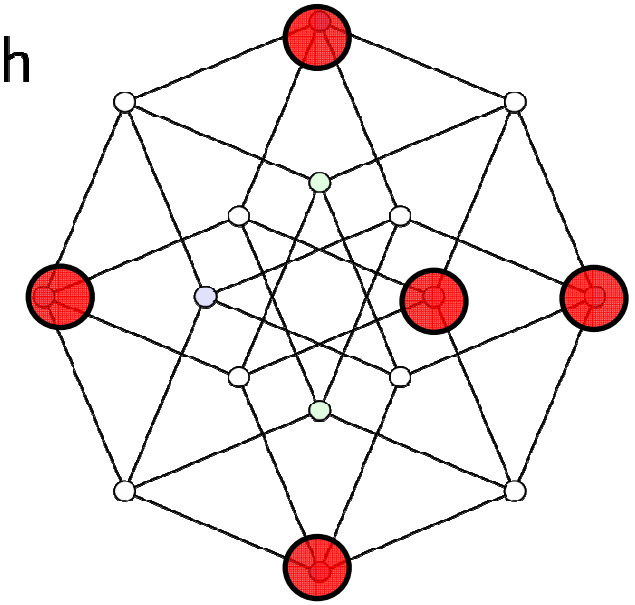

Natural approach: Take random subset.

Complete Failure: Subgraph empty when $K \ll 2^{\epsilon N}$

We chose vertices from linear code $C \subseteq GF(2)^N$.

Need "bad code": low rate $(\dim(C) \ll N)$, low distance $(dist(C) \leq \epsilon N)$

But not too bad…

Noisy cube: Graph $T_{1-\epsilon}$ Vertices $\{0,1\}^N$, edges: $x \sim y$ if $\Delta(x,y) = \epsilon N$

Main Thm: If subspace $C \subseteq GF(2)^N$ satisfies:

(i) $D := dist(C^{\perp}) = \omega(1)$

(ii) $C^{\perp}$ is ϵN-query locally testable via canonical tester:

> Test $y \in GF(2)^N$:
>
> 1) Choose random $w \in C$ s.t. $|w| \leq \epsilon N$
> 2) Accept iff $\langle y, w \rangle = 0$

i.e., $dist(y, C^{\perp}) > D/100 \Rightarrow \Pr[y\ accepted] = \frac{1}{2} + o(1)$

Then $T_{1-\epsilon} \upharpoonright C$ is small set expander w/ N e-vals $\geq 1 - 2\epsilon$

Instantiation: $\mathrm{N} = 2^{\mathrm{n}}$, C= Reed-Muller code of degree d ,$\epsilon = 2^{-d}$

$C^{\perp}$= Reed-Muller code of degree $\mathrm{n} - d - 1$

$dist(C^{\perp}) = 2^d$,tester works via

[Batthacharya-Kopparty-Schoenbeck-Sudan-Zuckerman'10]

Main Thm: If subspace $C \subseteq GF(2)^N$ satisfies:

(i) $D := dist(C^\perp) = \omega(1)$

(ii) $C^\perp$ is ϵN-query locally testable via canonical tester:

Test $y \in GF(2)^N$:

1) Choose random $w \in C$ s.t. $|w| \leq \epsilon N$

2) Accept iff $\langle y, w \rangle = 0$

i.e., $dist(y, C^\perp)$

Then $T_{1-\epsilon} \uparrow C$ is

Vertices: degree d poly's over $GF(2)^n$

Edges: $p \sim q$ if $p - q = f_1 f_2 \cdots f_d$

Where $f_1, .., f_d$ affine functions.

Instantiation: $N = 2^n$, C= Reed-Muller code of degree d, $\epsilon = 2^{-d}$

$C^\perp$= Reed-Muller code of degree $n - d - 1$

$dist(C^\perp) = 2^d$, tester works via

[Batthacharya-Kopparty-Schoenbeck-Sudan-Zuckerman'10]

Resulting graph has $\sim 2^{n^d}$ vertices and 2^n large e-vals.

Main Thm: If subspace $C \subseteq GF(2)^N$ satisfies:

(i) $D := dist(C^\perp) = \omega(1)$

(ii) $C^\perp$ is ϵN-query locally testable via canonical tester:

> Test $y \in GF(2)^N$:
>
> 1) Choose random $w \in C$ s.t. $|w| \leq \epsilon N$
> 2) Accept iff $\langle y, w \rangle = 0$

i.e., $dist(y, C^\perp) > D/100 \Rightarrow \Pr[y\ accepted] = \frac{1}{2} + o(1)$

Then $T_{1-\epsilon} \upharpoonright C$ is small set expander w/ N e-vals $\geq 1 - 2\epsilon$

Proof outline:

1) Dictatorships survive and correspond to large e-vecs of $T_{1-\epsilon} \upharpoonright C$

2) All large e-vecs of $T_{1-\epsilon} \upharpoonright C$ correspond to low degree polynomials.

3) By 2), many properties of hypercube are preserved, including small set expansion.

Main Thm: If subspace $C \subseteq GF(2)^N$ satisfies:

(i) $D := dist(C^\perp) = \omega(1)$

(ii) $C^\perp$ is ϵN-query locally testable via canonical tester:

> Test $y \in GF(2)^N$:
>
> 1) Choose random $w \in C$ s.t. $|w| \leq \epsilon N$
> 2) Accept iff $\langle y, w \rangle = 0$

i.e., $dist(y, C^\perp) > D/100 \Rightarrow \Pr[y\ accepted] = \frac{1}{2} + o(1)$

Then $T_{1-\epsilon} \upharpoonright C$ is small set expander w/ N e-vals $\geq 1 - 2\epsilon$

Proof outline:

1) Dictatorships survive and correspond to large e-vecs of $T_{1-\epsilon} \upharpoonright C$

2) All large e-vecs of $T_{1-\epsilon} \upharpoonright C$ correspond to low degree polynomials.

3) By 2), many properties of hypercube are preserved, including small set expansion.

1) Dictatorships survive and correspond to large e-vecs of $T_{1-\epsilon} \restriction C$

Lemma: If $f: GF(2)^N \to \mathbb{R}$ is e-vec of $T_{1-\epsilon}$, $C \subseteq GF(2)^N$ subspace, then $f \restriction C$ is e-vec of $T_{1-\epsilon} \restriction C$.

Proof: Both $T_{1-\epsilon}$, $T_{1-\epsilon} \restriction C$ are Cayley graphs:

In $T_{1-\epsilon}$, $x \sim y$ iff $x \oplus y \in B_\epsilon$

In $T_{1-\epsilon} \restriction C$, $x \sim y$ iff $x \oplus y \in B_\epsilon \cap C$

f e-vec of $T_{1-\epsilon} \Rightarrow f$ is character $\left(f(x \oplus y) = f(x)f(y)\right) \Rightarrow f \restriction C$ is character.

Note: Every character f has form $f(x) = -1^{\langle x,y \rangle}$ for some $y \in GF(2)^N$.

The corresponding e-val is $\mathbb{E}_{w \in B_\epsilon \cap C}[f(w)] = \mathbb{E}_{w \in B_\epsilon \cap C}\left[-1^{\langle w,y \rangle}\right]$

Lemma: If f is dictatorship character $f(x) = -1^{x_i}$ then e-val of $f \restriction C$ in $T_{1-\epsilon} \restriction C$ is $\geq 1 - 2\epsilon$

Proof: e-val is $\mathbb{E}_{w \in B_\epsilon \cap C}[-1^{w_i}]$, use $|w| \leq \epsilon N$, assume i typical.

Corollary: $T_{1-\epsilon} \restriction C$ has $\geq N$ eigenvalues of value $\geq 1 - 2\epsilon$.

Main Thm: If subspace $C \subseteq GF(2)^N$ satisfies:

(i) $D := dist(C^\perp) = \omega(1)$

(ii) $C^\perp$ is ϵN-query locally testable via canonical tester:

> Test $y \in GF(2)^N$:
>
> 1) Choose random $w \in C$ s.t. $|w| \leq \epsilon N$
> 2) Accept iff $\langle y, w \rangle = 0$

i.e., $dist(y, C^\perp) > D/100 \Rightarrow \Pr[y\ accepted] = \frac{1}{2} + o(1)$

Then $T_{1-\epsilon} \upharpoonright C$ is small set expander w/ N e-vals $\geq 1 - 2\epsilon$

Proof outline:

✔ 1) Dictatorships survive and correspond to large e-vecs of $T_{1-\epsilon} \upharpoonright C$

2) All large e-vecs of $T_{1-\epsilon} \upharpoonright C$ correspond to low degree polynomials.

3) By 2), many properties of hypercube are preserved, including small set expansion.

Main Thm: If subspace $C \subseteq GF(2)^N$ satisfies:

(i) $D := dist(C^\perp) = \omega(1)$

(ii) $C^\perp$ is ϵN-query locally testable via canonical tester:

> Test $y \in GF(2)^N$:
>
> 1) Choose random $w \in C$ s.t. $|w| \leq \epsilon N$
> 2) Accept iff $\langle y, w \rangle = 0$

i.e., $dist(y, C^\perp) > D/100 \Rightarrow \Pr[y\ accepted] = \frac{1}{2} + o(1)$

Then $T_{1-\epsilon} \upharpoonright C$ is small set expander w/ N e-vals $\geq 1 - 2\epsilon$

Proof outline:

1) Dictatorships survive and correspond to large e-vecs of $T_{1-\epsilon} \upharpoonright C$

2) All large e-vecs of $T_{1-\epsilon} \upharpoonright C$ correspond to low degree polynomials.

3) By 2), many properties of hypercube are preserved, including small set expansion.

2) All large e-vecs of $T_{1-\epsilon} \upharpoonright C$ correspond to low degree polynomials.

Recall: (i) Every character f has form $f_y(x) = -1^{\langle x,y \rangle}$ where $y \in GF(2)^N$.

The corresponding e-val is $\mathbb{E}_{w \in B_\epsilon \cap C}[f(w)] = \mathbb{E}_{w \in B_\epsilon \cap C}\left[-1^{\langle w,y \rangle}\right]$

(ii) $C^\perp$ is ϵN-query locally testable via canonical tester:

Test $y \in GF(2)^N$:

1) Choose random $w \in C$ s.t. $|w| \leq \epsilon N$
2) Accept iff $\langle y, w \rangle = 0$

i.e., $dist(y, C^\perp) > dist(C^\perp)/100 \Rightarrow \Pr[y \text{ accepted}] = \frac{1}{2} + o(1)$

Note: $f_y \upharpoonright C \equiv f_{y'} \upharpoonright C$ iff $y \oplus y' \in C^\perp$

Corollary: f_y is equivalent to degree ℓ character (i.e., $f_{y'}$ with $|y'| \leq \ell$) iff $dist(y, C^\perp) \leq \ell$

Corollary: $\forall$ character f of $T_{1-\epsilon} \upharpoonright C$. If f's e-val $\geq \Omega(1)$ then $f = f_y \upharpoonright C$ for some $|y| \leq dist(C^\perp)/100$.

2) All large e-vecs of $T_{1-\epsilon} \upharpoonright C$ correspond to low degree polynomials.

Recall: (i) Every character f has form $f_y(x) = -1^{\langle x,y \rangle}$ where $y \in GF(2)^N$.

The corresponding e-val is $\mathbb{E}_{w \in B_\epsilon \cap C}[f(w)] = \mathbb{E}_{w \in B_\epsilon \cap C}\left[-1^{\langle w,y \rangle}\right]$

(ii) $C^\perp$ is ϵN-query locally testable via canonical tester:

Test $y \in GF(2)^N$:

1) Choose random $w \in C$ s.t. $|w| \leq \epsilon N$
2) Accept iff $\langle y, w \rangle = 0$

i.e., $dist(y, C^\perp) > dist(C^\perp)/100 \Rightarrow \Pr[y \; accepted] = o(1)$

Note: $f_y \upharpoonright C \equiv f_{y'} \upharpoonright C$ iff $y \oplus y' \in C^\perp$

Corollary: f_y is equivalent to degree ℓ character (i.e., $f_{y'}$ with $|y'| \leq \ell$) iff $dist(y, C^\perp) \leq \ell$

Corollary: $\forall$ character f of $T_{1-\epsilon} \upharpoonright C$. If f's e-val $\geq \Omega(1)$ then $f = f_y \upharpoonright C$ for some $|y| \leq dist(C^\perp)/100$.

2) All large e-vecs of $T_{1-\epsilon} \upharpoonright C$ correspond to low degree polynomials.

$\forall$ character f of $T_{1-\epsilon} \upharpoonright C$. If f's e-val $\geq \Omega(1)$ then $f = f_y \upharpoonright C$ for some $|y| \leq dist(C^{\perp})/100$.

3) By 2), many properties of hypercube are preserved, including small set expansion.

Crucial observation: Random variable $x \in_R C$ is $D-1$-wise independent.

Corollary: For every poly f of degree$< D$, $\mathbb{E}_{x \in C}[f(x)] = \mathbb{E}_{x \in GF(2)^N}[f(x)]$

Small set expansion of noisy cube is proven by showing that if f is degree-ℓ polynomial, then $\mathbb{E}_{x \in GF(2)^N}[f^4(x)] \leq 9^{\ell}\mathbb{E}_{x \in GF(2)^N}[f^2(x)]^2$.

[Nash'??,Bonami'68]

Thus proof immediately carries over to our case.

Our main result: New exponentially more efficient replacement to long-code/noise-graph gadget.

Some Applications:

- Existence of small-set expander with $\exp(\log^{\epsilon} n)$ eigenvalues close to 1.

  Demostrates ABS algorithm can take $\exp(poly^*(n))$ time.

- Alphabet reduction / inner PCP gadget* for unique games with quasipolynomial blowup in alphabet size.

  Improves on exponential blowup of long-code gadget.

- Max Cut integrality gap for $poly^*(\log n)$ … $poly^*(n)$ rounds of semi definite hierarchies.

  Exponential improvement over prior bounds.

Invariance Principle and Alphabet Reduction

Invariance Principle: [Mossel-O'Donnel-Oleszkiewicz'05]

$f: \{\pm 1\}^N \to \mathbb{R}$ low-degree polynomial of low max influence, then

$$f(X) \sim f(Y)$$

X:Uniform on $\{\pm 1\}^N$ Y:N independent Gaussians.

Crucial tool in several applications of noisy cube via "Majority is Stablest" [Khot-Kindler-Mossel-O'Donnel'04]

We want: Prove still holds when X is D-wise independent.

Unfortunately, proof doesn't involve just finite moments.

What we achieve: Holds when X is a random Reed-Muller codeword.

Main tool: PRG for polynomial threshold functions. [Meka-Zuckerman'10]

Corollary: $2^{poly(n)}$-size gadget reducing unique-games over $GF(2)^n$ into max-cut / sparsest cut (/ smaller alphabet UG*)

Integrality gaps for Unique Games, Max Cut

- Take our graph – treat as label-extended graph for unique games instance.
- Results in integrality gap with completeness $\sim 1 - 1/\log|V|$ (compare with $1 - 1/\log\log|V|$ of prior works)
- Compose with alphabet-reduction gadget to get Max-Cut instance, analyze via [Raghavendra-Steurer'09]
- n vertex Max-Cut instance resisting:
  - $poly^*(log\ n)$ rounds in SDP+SA hierarchy (compare to $poly(\log\log n)$)
  - $poly^*(n) = \exp(\exp(poly(\log\log n)))$ in local embedability hierarchy (compare to $poly(\log n)$)

Open Problem: Prove the UGC

(Maybe with $n^{\log n}$ or $2^{\sqrt{n}}$ time reduction)

Our gadget yields a candidate "inner PCP".

Need to find the right outer PCP..

Q: Is there a subexponential algorithm for smooth label cover?

Open Problem: Refute the UGC

(Maybe with $2^{2^{\sqrt{\log n}}}$ or $n^{\log n}$ time algorithm)

Seems we need to use a different algorithm..

..or maybe just stronger SDP hierarchies?

Q: Is there Max-Cut/UG instance requiring $\omega(1)$ rounds for weak SDP hierarchies, but $O(1)$ rounds for Lasserre?