PCPs and Inapproximability: Recent Milestones and Continuing Challenges

Venkatesan Guruswami

Carnegie Mellon University

--- Princeton, June 15, 2011 ---

Intended to be a survey* of some key developments of last ≈ 10 years

Let’s begin with a glimpse of where PCP theory was ≈ a decade back

PCPs circa 2000

PCP theorem

- PCP theorem: [AS, ALMSS] Polynomial size witnesses for NP languages that can be checked by a randomized polytime verifier by just probing 3 bits (with soundness error < 0.999)
- Equivalently, $\text{Gap3SAT}_{1,\alpha}$ is NP-hard for some $\alpha < 1$

- (Surprising?) Connection to approximating Clique (based on FGLSS graph) discovered before connection to approximating Max 3SAT
  - PCP theorem implies factor n^{γ} inapproximability for Clique

PCPs: early years

- PCP theorem implied APX hardness for many problems (via their classic reductions from 3SAT).
  - Generally very weak factors.
  - Quest for better (optimal?) factors followed

- 2-prover 1-round proof systems ("bipartite 2-query PCPs") emerged as the canonical PCP to reduce from:
  - Constraint satisfaction problem called *Label Cover*
  - Example early Label Cover-based success:

    $\Omega(\log n)$ hardness for set cover [Lund-Yannakakis]

Label Cover

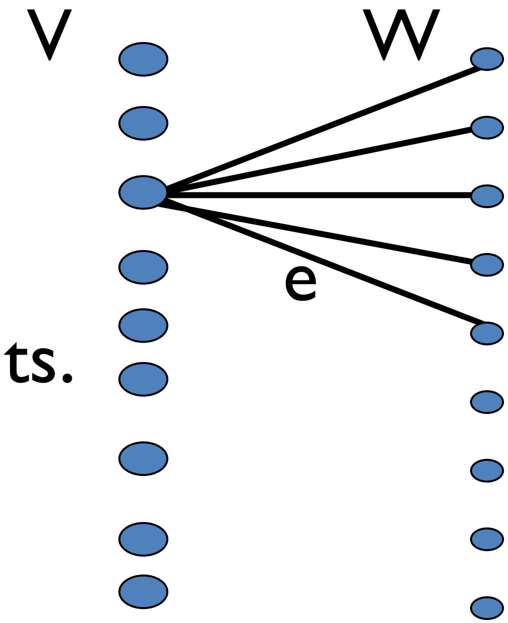

- Binary (arity two) CSP over large domain [R] with "projection" constraints.
- Instance consists of:
  - Bipartite graph G=(V,W,E)
  - For each $e = (v,w) \in E$, a function $\pi_e : [R] \to [R]$.
- Assignment (labeling) $A: V \cup W \to [R]$ *satisfies* an edge $e=(v,w)$ if $\pi_e(A(v)) = A(w)$
  - Value of instance = *maximum fraction of edges satisfied by a labeling*
- Raz's parallel repetition theorem gives a strong (value 1 vs $1/R^{\gamma}$) gap hardness for Label Cover

Strong(er)/optimal PCPs

- Improvements in PCP params, aimed in part at better hardness results (using specific predicates for PCP check)
  - Label Cover used as "outer" PCP
  - Composed with "inner" PCP (trading off soundness for much smaller # queries)
    - Paradigm: **Encode** labels & test codewords
    - "Inner" task: Check if $\pi(a)=b$ reading very few bits of Enc(a) and Enc(b)

- Which code to use?

Brilliant invention of [Bellare-Goldreich-Sudan]: LONG CODE

Long Code (aka Dictator functions)

- For $a \in \{1,2,\ldots,R\}$, $\mathrm{LONG}(a) : \{0,1\}^R \rightarrow \{0,1\}$

  $\mathrm{LONG}(a)(x) = x_a$ for every $x \in \{0,1\}^R$

  - Very redundant (encodes log R bits into 2^R bits)
    - has the value of *every* function $[R] \rightarrow \{0,1\}$ at a
    - but doesn't hurt to have around if R is constant.
  - Surprisingly useful!

The first optimal PCPs

- [Håstad'96]: zero "amortized free bit" PCP ⇒ factor $n^{1-\varepsilon}$ inapproximability for Clique
- [Håstad'97]: $\text{Gap3LIN}_{1-\varepsilon,\frac{1}{2}+\varepsilon}$ is NP-hard (3-query PCP with completeness $1-\varepsilon$ and soundness $\frac{1}{2}+\varepsilon$). Optimal!

  Also, similar result mod p, and NP-hardness of
  - $\text{Gap3SAT}_{1,7/8+\varepsilon}$ & $\text{Gap-4-Set-Splitting}_{1,7/8+\varepsilon}$ Optimal!
- Several hardness results (currently best known, under NP ≠ P) via gadget reductions from Håstad's results:
  - MaxCut: 16/17 Max2SAT: 21/22 MaxDiCut: 11/12
  - NAE-3SAT: 15/16 3-Set-Splitting: 19/20 3-coloring: 32/33 (these are with perfect completeness)

Approximation resistance

- The tight results for CSPs showed "approx. resistance"
  - hard to beat naive random assignment.
- An exception: *Max 3MAJ*
  - 2/3+ε hardness via reduction from 3LIN
  - Factor 2/3 algorithm [Zwick]
- In parallel with hardness revolution, sophisticated SDP rounding methods developed. Eg.
  - 7/8 algo for Max3SAT [Karloff-Zwick]
  - Factor ½ algo for Max3CSP and factor 5/8 for satisfiable 3CSP.

Optimal PCPs: queries vs soundness

- [G.-Lewin-Sudan-Trevisan] 3-query *adaptive* PCP with perfect completeness & soundness $\frac{1}{2}+\varepsilon$
- [Samorodnitsky-Trevisan] $1+\varepsilon$ *amortized query complexity:* k queries, soundness $2^{-k+o(k)}$
  - $2^{-k+o(k)}$ hardness for Max k-CSP.
  - Later with perfect completeness [Håstad-Khot]
  - Useful starting points in some reductions, eg. Low-congestion path routing, Clique.

Low-soundness Multiprover systems

- [Arora-Sudan; Raz-Safra] O(1) prover 1-round proof systems with $\exp(-(\log n)^{\Omega(1)})$ soundness
  - *NP-hardness* of $\Omega(\log n)$ factor for set cover
  - Similar result for 2-prover case open
    - Would have more applications, like hardness of lattice problems

- [Dinur-Fischer-Kindler-Raz-Safra] O(1) prover systems with
  - $\exp(-(\log n)^{0.99})$ soundness.
  - Proof of BGLR *"sliding scale" conjecture* for up to $(\log n)^{0.99}$ bits read

Covering PCPs

- Notion of soundness tailored to coloring problems
  - Covering soundness = minimum number of proofs that can "cover" all constraints (for every check, at least one proof should cause acceptance)
  - [G.-Håstad-Sudan] 4-query PCP with $\omega(1)$ covering soundness
    - Super-constant hardness for coloring 2-colorable 4-uniform hypergraphs.
    - Later also for 3-uniform hypergraphs [Khot] [Dinur-Regev-Smyth]

<u>Frontier:</u> Rule out 5-coloring 3-colorable graphs in polytime

PCPs till ~ 2000 summary

- Label Cover hardness
  - versatile starting point for inapproximability (continues to be prominent)
- Label Cover + Long Code + Fourier analysis paradigm
- Tight hardness results for several CSPs of arity ≥ 3
- Arity 2 CSPs not well understood (results only via gadgets)

PCPs in the last decade

New Proof Composition methods

Low soundness error 2-query PCPs

Short PCPs ($n^{1+o(1)}$ size):
- Best known n $(\log n)^{O(1)}$

New "outer" PCPs
- multilayered, smooth, mixing, Dinur-Safra, etc.
- Conjectural forms: Unique Games, 2-to-1, …

New proofs, notions:
- Dinur's gap amplification
- Robust PCPs
- PCP of proximity (PCPP)

Dictatorship tests and new "inner" PCPs
- New analytic machinery

New proofs and notions

Dinur's proof

- Gap amplification: Reduce Gap-3Color$_{1,1-\delta}$ to Gap-3Color$_{1,1-2\delta}$ provided $\delta < 10^{-6}$
  - Apply O(log n) times starting with δ=1/m
  - Shows that Gap-3Color$_{1,\alpha}$ is NP-hard for some constant α<1 (this implies the PCP theorem)
- *PCP via inapproximability instead of other way around*
- Requires elements of old PCP in alphabet reduction
  - Constant sized PCP: variant called "assignment tester" that checks if assignment x is close to satisfying circuit C

Robust PCPs

- PCP soundness: When $\varphi \notin$ SAT, for all proofs π, with probability ρ, check C_I rejects π_I (I=randomly chosen query positions)

Proof π

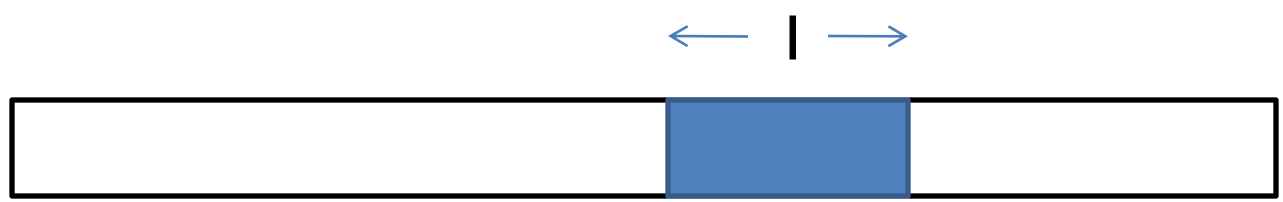

Check $C_I(\pi_I) = 1$

- Robust PCP: stronger soundness guarantee
  - π_I far from satisfying C_I (with good prob.)
  - Formally,

$$\mathbb{E}_I\left[\Delta(\pi_I, \mathrm{SAT}(C_I))\right] > \eta$$

(η = robust soundness)

Robust PCPs & PCPPs

- Check if π_l satisfies C_l or is *far* from satisfying C_l ***recursively,*** using another "inner" PCP
- Inner primitive: *PCP of proximity* (PCPP)
  - Input: circuit C
  - Proof: (purported) satisfying assignment x and proof of proximity σ that x satisfies C
  - Verifier (Assignment Tester): read few bits in *both* x and σ;
    - For satisfying x, $\exists\ \sigma$ with acc. prob. 1
    - If x is δ-far from satisfying C, then $\forall\ \sigma$, $\Omega(1)$ rej. prob.
- Useful when proximity parameter δ of inner PCPP is < robust distance η of outer PCP

Composition streamlined

- Robust PCPs compose "syntactically" with "inner" PCPs of proximity (when PCPP proximity parameter < robustness)
- [Ben-Sasson, Goldreich, Harsha, Sudan, Vadhan] Can check that original polynomial and Hadamard based PCPs can be made robust PCPPs.
  - Simplifies details of composition
  - Used to give near-linear size PCPs
- PCPP also used in Dinur's alphabet reduction step
  - Explicit coding used to create distance between inconsistent assignments

Talk Plan

- New proofs and notions
- Robust PCPs and PCPs for proximity
- Short PCPs
- Low-soundness error Label Cover
- Unique Games, Dictatorship tests, etc.
- NP-hardness via structured outer PCPs
- Some challenges

Quasi-linear PCPs

- [BGHSV] PCPs of length $n \cdot 2^{(\log n)^{\epsilon}}$ ($O(1/\varepsilon)$ queries)
- [Ben-Sasson, Sudan] *Univariate* polynomial based PCP
  - *Proof of proximity* for Reed-Solomon codes which makes it locally testable
  - $n (\log n)^{O(1)}$ sized PCP with $(\log n)^{O(1)}$ queries
  - $O(\log \log n)$ steps of Dinur's gap amplification gives

    $n (\log n)^{O(1)}$ sized PCP with $O(1)$ queries

- Implication for approximation: APX-hardness via reduction from 3SAT hold for $2^{\tilde{\Omega}(n)}$ time algos, under the ETH

- New proofs and notions
- Robust PCPs and PCPs for proximity
- Short PCPs
- Low-soundness error Label Cover
- Unique Games, Dictatorship tests, etc.
- NP-hardness via structured outer PCPs
- Some challenges

Label cover with o(1) soundness

- Håstad's 7/8+ε hardness for 3SAT requires soundness error of Label Cover to be $\ll \varepsilon$
- Getting this via parallel repetition makes Label Cover instance size $n^{\Omega(\log (1/\epsilon))}$ (with large hidden constant factor)
- Can we get such soundness with LC size say $O_\varepsilon(n^3)$?

Answer: [Moshkovitz-Raz] YES. In fact, with *near-linear* size !
(Was very surprising to me)

For soundness ε, Label Cover with $n^{1+o(1)}\,\mathrm{poly}(1/\varepsilon)$ vertices.

- Worse dependence on domain size: $R = \exp(\mathrm{poly}(1/\varepsilon))$ instead of $\mathrm{poly}(1/\varepsilon)$ in Raz.

Inapproximability consequence

- $7/8+\varepsilon$ approx. for 3SAT requires $\exp(\Omega_\epsilon(n^{1-o(1)}))$ time (under "Exponential Time Hypothesis")
  - constant factor: *doubly exponentially* small in ε
  - Similar claims for other hardness results based on Label cover + Long code testing
- Sharp *complexity dichotomy* at the approximation threshold: polytime vs. exponential time.
  - Eg. unlikely that there's a factor $7/8+\varepsilon$ approx. algo for Max3SAT with runtime $\exp(n^{10\sqrt{\varepsilon}})$

[Moshkovitz-Raz] approach

- Start with Label Cover of low soundness error ε but large alphabet Σ
  - $\varepsilon = (\log n)^{-\beta}$, $|\Sigma| = poly(n)$ (based on low-degree testing in list-decoding regime)
  - *reduce alphabet size* via composition

- New composition method for low-error regime that does not increase # provers
  - Based on "Locally Decode/Reject Codes"

<u>Next</u>: Few words on an alternate approach (giving polynomial instead of near-linear size)

Derandomized parallel repetition

- *u-parallel repetition* of Label Cover reduces soundness error from (say) 0.99 to $\varepsilon=2^{-O(u)}$, but blows up size to $\approx n^u$
  - Can't get $o(1)$ soundness with polytime reduction (const. u)
- Derandomization: Can we select a *poly-sized subset* of all possible u questions?
  - Limitation [Feige-Kilian]: for u=O(log n) and poly(n) size subsets, $\varepsilon \geq 1/\text{poly}(\log n)$
- [Dinur-Meir] Match this lower bound, combining
  - Derandomized *direct product testing* based on *subspaces* [Impagliazzo-Kabanets-Wigderson]
  - Structured "linear" PCPs
    - Identify proof coordinates (vertices) with F^m
    - Edges corresponding to 2-query checks form a *subspace* of F^{2m}

Composing without an extra prover

- [Dinur-Harsha] alternate composition method to reduce alphabet size keeping #provers at 2
  - Based on "decodable PCPs"
  - Exploit equivalence between robust PCPs and Label Cover

- Applying this to [Dinur-Meir] 2-query PCP, gives:
  - Label Cover of *fixed polynomial* (though not near-linear) size with soundness ε, $\forall$constants $\varepsilon > 0$

Label Cover ⇔ Robust PCPs

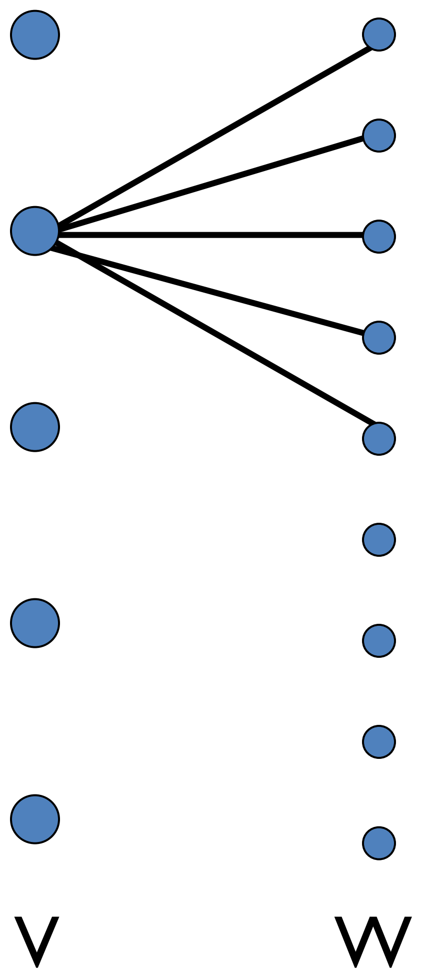

Label Cover ⇒ Robust PCPs

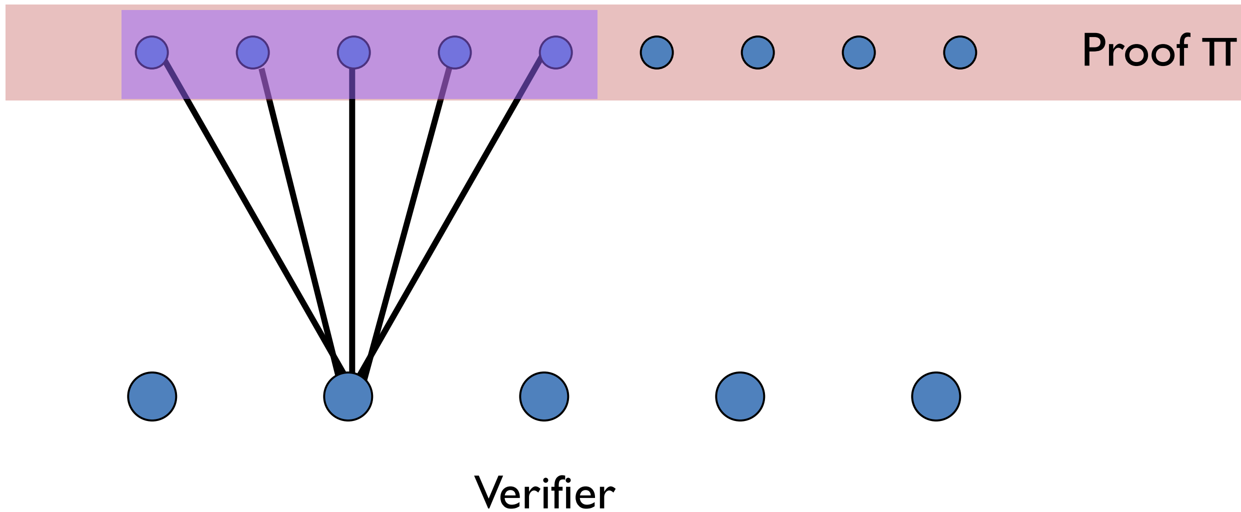

<u>Verifier</u>

1. Selects a random "big" vertex **v** ∈ **V**
2. Reads entire neighborhood of **v**
3. Accepts iff there is a value for **v** that would cause all edge constraints to accept.

Label Cover $\Rightarrow$ Robust PCPs

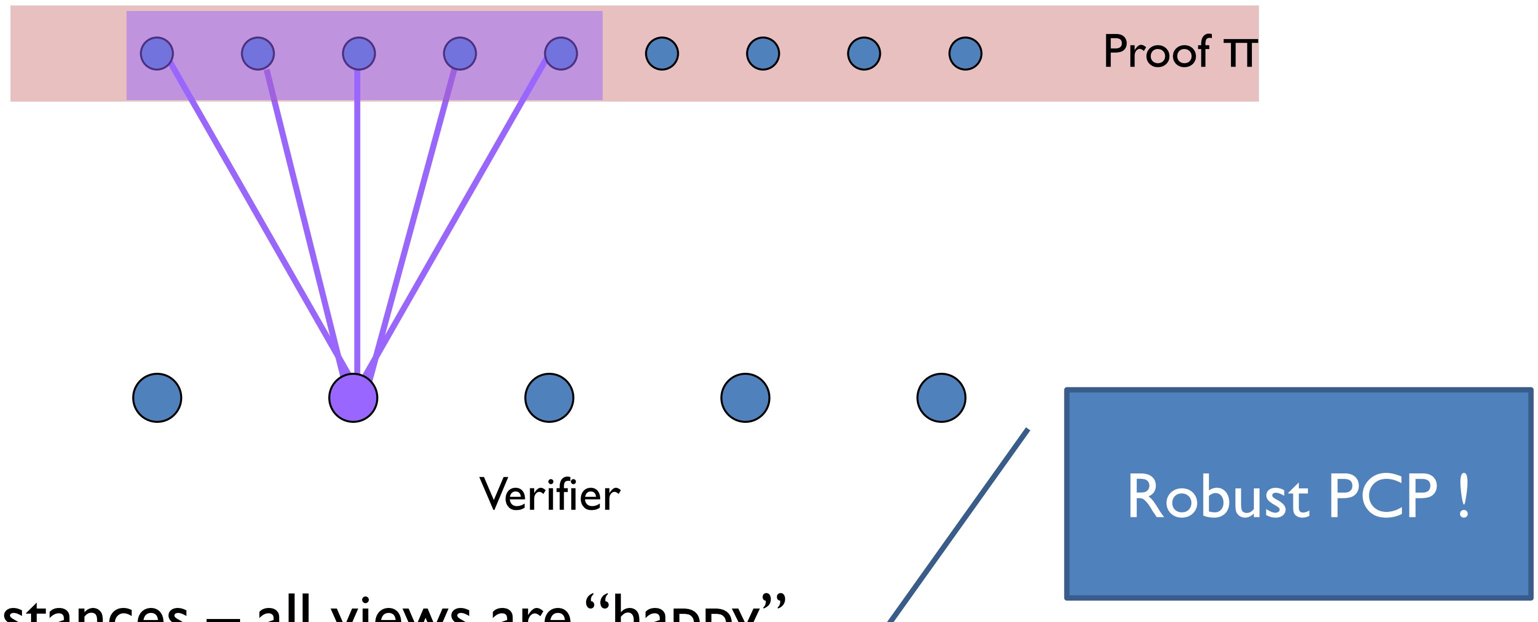

YES instances – all views are "happy"

NO instances – average view is very "unhappy", i.e. view from a random window is at most δ -close to a satisfying view.

Label Cover $\Leftarrow$ Robust PCPs

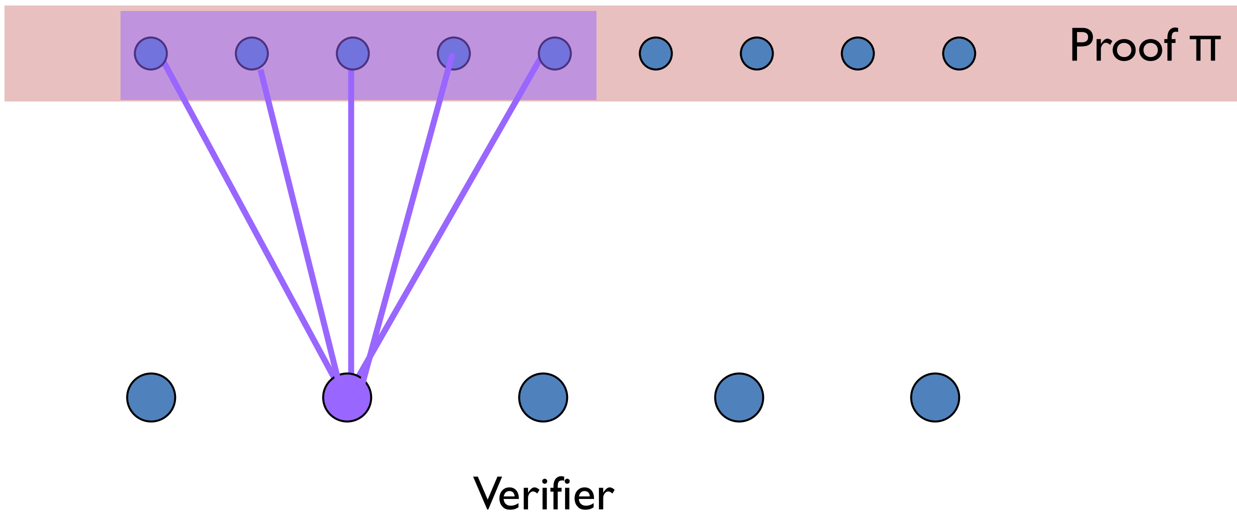

- This transformation is "invertible" (rotate back!)
- $|\Sigma|$ corresponds to the number of accepting configurations, which is $\leq$ exp(# PCP queries)
  - Reducing PCP queries $\Rightarrow$ reducing LC alphabet size

Low-error 2-prover systems summary

- Some very exciting recent constructions
- Frontier: 2-query PCP of *polynomial size and polynomial alphabet* with soundness error $1/(\log n)^{10}$

- New proofs and notions
- Robust PCPs and PCPs for proximity
- Short PCPs
- Low-soundness error Label Cover
- Unique Games, Dictatorship tests, etc.
- NP-hardness via structured outer PCPs
- Some challenges

Long code based "inner" PCPs

- Task: Checking that labels $a = A(v)$ and $b = A(w)$ satisfy some Label Cover projection constraint $\pi(a) = b$
- Check constraints on few (say q) bits of (purported) long codes $f, g : \{0,1\}^R \to \{0,1\}$ of labels a and b
  - If f and g are long codes of "consistent" a, b (i.e., $\pi(a) = b$), accept with prob. 1 (or $1-\varepsilon$)
  - Acceptance with prob. $> s + \varepsilon$ implies one can randomly *"decode"* f, g into labels a', b' s.t. $\pi(a') = b'$ with prob. $\theta = \theta(\varepsilon)$
  - This would imply a soundness $s + 2\varepsilon$ q-query PCP

Let's see Håstad's 3-query PCP, and [...] constructing 2-query PCPs.

Recall:
$f = \text{LONG}(a) \Rightarrow$
$\mathbf{f(x) = x_a}$

Håstad's 3-query PCP

- Pick $x, y \in \{0,1\}^R$ u.a.r.
- Pick $\mu \in \{0,1\}^R$ from the ϵ-biased distribution.
- Define $z \in \{0,1\}^R$ by $z_j = y_j \oplus x_{\pi(j)} \oplus \mu_j$.
- With prob. 1/2 check $g(x) \oplus f(y) \oplus f(z) = 0$, and with prob. 1/2 check $g(x) \oplus f(y) \oplus f(\overline{z}) = 1$ ($\overline{z}$ = coordinate-wise complement of z).

Completeness: If $f(y) = y_a$ and $g(x) = x_b$ and $\pi(a) = b$, then $g(x) \oplus f(y) \oplus f(z) = x_b \oplus y_a \oplus y_a \oplus x_{\pi(a)} \oplus \mu_a = \mu_a$

Suffices for decoding labels

Soundness: Acceptance prob. > ½ +ε ⇒ g, f have non-trivial agreement with "consistent" low-weight linear functions

3 vs. 2 queries

- Get $(1-\varepsilon)$ vs. $(\frac{1}{2}+\varepsilon)$ hardness for Max3LIN (mod 2)
  - Approximation resistant
- Each query point x,y, z is uniformly distributed in $\{0,1\}^R$
  - y,z are correlated, but f has to give value to each separately (and each is uniform)
- What about 2-variable linear equations mod 2?
  - [Goemans-Williamson] algo finds assignment of value $1-O(\sqrt{\varepsilon})$ in $(1-\varepsilon)$-satisfiable instance
  - Matching hardness through a 2-query PCP ?

A 2-query PCP?

Here's a natural test, saving 1 query in Håstad's test:

- Pick $x \in \{0,1\}^R$ u.a.r, and ϵ-biased noise vector $\mu \in \{0,1\}^R$
- Set $y = x \circ \pi \oplus \mu$, i.e., for $j \in [R]$, set $y_j = x_{\pi(j)} \oplus \mu_j$.
- With prob. $1/2$, check $g(x) \oplus f(y) = 0$,
  with prob. $1/2$, check $g(x) \oplus f(\overline{y}) = 1$.

Trouble: query y to f is not uniform.
$y_j = y_k$ with prob. close to 1 when $\pi(j) = \pi(k)$

Query y reveals lots of information about projection π
Could form "cheating" f by "piecing together"
many inconsistent long codes, for portions of $\{0,1\}^R$
corresponding to different projections π_e

How to circumvent this?

Unique Games

- Khot's insight: if π is a **bijection**, then $y = x \circ \pi$ is uniformly distributed (since x is); gives no clue about π

Unique Game (UG)

* Label Cover where ***all projection constraints are bijections***

Khot's Unique Games Conjecture (UGC):

$\text{GapUG}_{1-\varepsilon,\varepsilon}$ is NP-hard for R > R(ε)

- UGC $\Rightarrow$ analysis of 2-query test reduces to f=g case
  - show that if f passes w.h.p, then f is "like" a long code (modern term: ***dictator***)
  - just codeword testing, no "consistency" checking

2-query dictator testing

The core question becomes analyzing "noise stability"

$NS_\varepsilon(f) = \text{Prob}_{x,\mu}[\, f(x) = f(x \oplus \mu)\,]$ (assume f is balanced)

- If f = dictator, then $NS_\varepsilon(f) = 1 - \varepsilon$
- If $NS_\varepsilon(f)$ is close to $1 - \varepsilon$, what can we say?

[Bourgain] If $NS_\varepsilon(f) > 1 - \varepsilon^{0.51}$ then f is close to a *junta* (depends on few coordinates)

[Mossel-O'Donnell-Oleszkiewicz] (Majority is Stablest Thm)
If $NS_\varepsilon(f) > 1 - \Theta(\sqrt{\varepsilon})$ then f has an *influential* coordinate.

Both of these can be used in reduction from Unique Games

UGC consequences…

- (2- ε) hardness for Vertex cover [Khot-Regev]
- 0.878.. hardness for Max Cut [Khot-Kindler-Mossel-O'Donnell] (using Majority is Stablest)
- Near-tight hardness for all Boolean 2CSPs [Austrin]
- Optimal hardness for *every* CSP [Raghavendra] (using invariance principle of [Mossel])
- Approximation resistance of *every ordering* CSP [G.-Håstad-Manokaran-Raghavendra-Charikar]

Hardness matching ***LP*** integrality gaps:

- Multiway Cut, Metric Labeling [Manokaran-Naor-Raghavendra-Schwartz]
- Strict CSPs, covering problems [Kumar-Manokaran-Tulsiani-Vishnoi]

- New proofs and notions
- Robust PCPs and PCPs for proximity
- Short PCPs
- Low-soundness error Label Cover
- Unique Games, Dictatorship tests, etc.
- NP-hardness via structured outer PCPs
- Some challenges

Proving UGC predictions without UGC?

Bypassing UGC?

- UGC has predicted many strong results
  - All plausible & consistent with our knowledge
  - "combinatorial" ones (like embedding lower bound) confirmed unconditionally
- Yet, ***no*** UG-completeness result so far
  - Possible that consequences of UGC are true but conjecture itself is false
- Natural question: Can we verify some of UGC's predictions *without* resorting to UGC?
  - Ideally, show NP-hardness
  - Or hardness under weaker assumptions (like optimality of Goemans-Williamson, 2-to-1 conjecture)?

Smooth Label Cover

- [Khot] Label Cover where for each v $\in$ V, the projections π_e for edges e incident on v form a "hash family"

$$\forall a \neq a' \quad \mathrm{Prob}_{e \ni v}[\pi_e(a) = \pi_e(a')] \to 0 \qquad \text{(for large } R\text{)}$$

- Gap-Smooth-Label-Cover$_{1,\varepsilon}$ is NP-hard
  - Reduction from Label Cover (note perfect completeness)

- Such "locally unique" projections have been useful in some NP-hardness reductions
  - 3-coloring 3-uniform hypergraphs [Khot]
  - Learning intersection of two halfspaces [Khot-Saket]
  - Agnostic learning monomials by halfspaces

[Feldman-G.-Raghavendra-Wu]

A Gaussian approximation threshold

L_p Subspace approximation problem ($2 < p < \infty$)

Input: A set of m points $a_1, a_2, \ldots, a_m \in \mathbb{R}^n$.
Goal: Find a $(n-1)$-dimensional subspace (hyperplane) H minimizing

$$\sum_{i=1}^{m} \mathrm{dist}(H, a_i)^p$$

where $\mathrm{dist}(H, a)$ is the min. Euclidean distance between a and any point in H.

[Deshpande-Tulsiani-Vishnoi]
Factor β_p algorithm where $\beta_p = \mathbb{E}_{g \sim N(0,1)}\left[|g|^p\right]$ is the p'th moment of the standard Gaussian.
And matching $\beta_p - \varepsilon$ Unique-Games-hardness

[G.-Raghavendra-Saket-Wu] NP-hardness of factor $\beta_p - \varepsilon$ approximation, using smooth Label Cover

Other Label Cover variants

Other structured Label Cover instances discovered in the last decade:

- Multilayered PCPs
- 2-to-1 projections (conjectural)
- Dinur-Safra

Multilayered PCP

- Multipartite label cover:
  - L layers of vertices $V_1, V_2, \ldots, V_L$
  - Projection maps between every pair of layers:
    - for edge e between $v_i \in V_i$ and $v_j \in V_j$ $(i < j)$, $\pi_e : [R] \to [R]$ from label of v_i to label of v_j
- Ensure every γ fraction of vertices have many constraints amongst them (for $L > L(\gamma)$)
- Introduced in [Dinur-G.-Khot-Regev] to show factor $(k-1-\varepsilon)$ hardness for vertex cover on k-uniform hypergraphs
- Later use in hypergraph coloring, non-mixed 3SAT, etc.

2-to-1 conjecture

- Label Cover where projection maps are 2-to-1
- Conjecture [Khot]: Gap-2-to-1-$LC_{1,\varepsilon}$ is NP-hard
  - Parallel repetition gives poly($1/\varepsilon$)-to-1 projections
- Consequences:
  - $\sqrt{2}-\varepsilon$ hardness for vertex cover [Dinur-Safra], [Khot]
  - Hardness of O(1)-coloring 4-colorable graphs [Dinur-Mossel-Regev]
  - Hardness of Gap-No-$\text{Two}_{1,5/8+\varepsilon}$ [O'Donnell-Wu]
  - Factor $1 - \frac{1}{k} + O\left(\frac{\ln k}{k^2}\right)$ hardness for Max k-coloring [G.-Sinop]

[Dinur-Safra]

- This remarkable paper showed factor 1.3606 NP-hardness for vertex cover
- Underlying this was a "2-to-1 like" Label Cover
  - label for v consistent with one or two labels for w
  - Soundness: $\forall$ labeling, every ε fraction of vertices has a 100-clique of inconsistent pairs.
- Other applications?

Wrap-up

- PCPs remarkably successful in showing inapproximability (even beyond initial expectations?)
  - Breadth of problems.

    I find it amazing what all Label Cover can be reduced to!
  - Many tight results
- Some notorious problems have withstood resolution
  - Densest subgraph, minimum linear arrangement, bipartite clique, sparsest cut, graph bisection, etc.
  - *Known algorithms have superconstant approx. ratio,*

    but even APX-hardness not known

Cut challenges

Eg. Uniform Sparsest Cut, Minimum Bisection:

Best approximation $(\log n)^{\Omega(1)}$. Hardness evidence:

1. Refuting random 3SAT is hard $\Rightarrow$ Factor 1.1 hardness [Feige]
2. Polytime $(1+\varepsilon)$ approximation $\Rightarrow$ NP has $2^{n^{\epsilon'}}$ time algorithms [Khot, "quasi-random PCPs"]
3. Superconstant hardness under "SSE hypothesis" (stronger than Unique Games conjecture) [Raghavendra-Steurer-Tulsiani]

"Easiness" evidence [G.-Sinop]

- Factor $(1+\varepsilon)/\lambda_r$ approximation in $2^{O_\epsilon(r)} n^{O(1)}$ time where λ_r is the r'th smallest eigenvalue of normalized Laplacian.
  - Factor $3/\lambda_r$ for minimum uncut

Challenges

- Can PCP machinery (even assuming UGC) give strong hardness results for Steiner Tree, TSP, Asymmetric TSP ?
- Lasserre integrality gaps beyond known hardness bound for Vertex Cover, Max Cut (or Unique Games)?
  - Just 4 rounds could improve [GW] and refute UGC !
- Unique-Games-completeness?
- Bypass UGC for some other consequences?
- Other hardness assumptions: eg. Densest subgraph?
  - Or finding indep. Sets of size εn when one of size n/100 exists
- Unchartered terrain for inapproximability:
  - eg. , nearest codeword in algebraic codes, bin packing, …