How easily can proofs be checked?

The Riemann Hypothesis is true (12^{th} Revision)

By

Ayror Sappen

 # Pages to follow: 15783

Slide borrowed from Madhu Sudan

PCP Thm as hardness of approximation

$\exists$ Alg computing H from G s.t.

- If G is 3-colorable $\Rightarrow$ H is 3-colorable
- If G is NOT 3-colorable $\Rightarrow$

  H is Really NOT 3-colorable : $\text{unsat}(H) > \varepsilon$

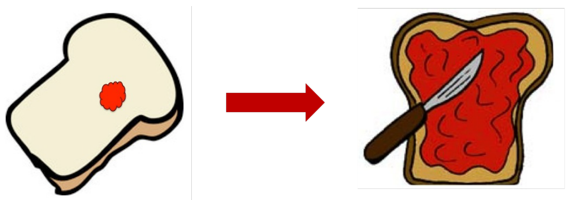

PCP Thm as a robust local proof checking :

Every $L \in NP$ has a verifier s.t.
Given $x \overset{?}{\in} L$ and a witness π, the verifier reads x, tosses $O(\log|x|)$ random coins, reads $\leq q = O(1)$ bits from π, and accepts/rejects

- If $x \in L \Rightarrow \exists \pi$ s.t. $\operatorname{Prob}_r\left[\text{Ver}^{\pi}(x,r) \text{ accepts}\right] = 1$
- If $x \notin L \Rightarrow \forall \pi$ s.t. $\operatorname{Prob}_r\left[\text{Ver}^{\pi}(x,r) \text{ accepts}\right] < 1-\varepsilon$

Context :

- PCPs originate in crypto 1980's –
- FGLSS, ALMSS '90–'91 : connection to inapprox
- Theory of inapprox :
  - phase 1 : basic inapprox
  - phase 2 : tight inapprox
  - phase 3 : UGC & SDP
- Modern Research

crypto & practical motivation

verified computation

zero knoledge PCPs

IOPPs (add interaction)

succint proofs

High dim expansion

local testability

$\Leftrightarrow$

cosystolic expansion

related to:

LTCs

quantum LDPCs

&

agreement tests

Basic PCP thm by gap amplification

Thm: $\exists \varepsilon_0 > 0$, poly time alg that takes G to H

- $|H| = \text{poly}(G)$
- G 3-col $\longrightarrow$ H 3-col
- G not 3-col $\longrightarrow$ $\text{unsat}(H) > \varepsilon$

$$G = G_0 \curvearrowright G_1 \curvearrowright \cdots \curvearrowright H$$

$G_i \curvearrowright G_{i+1}$: $\quad |G_{i+1}| = O(|G_i|)$

G_i 3col $\longrightarrow$ G_{i+1} 3-col

G_i not 3-col $\longrightarrow$ $\text{unsat}(G_{i+1}) > 2 \cdot \text{unsat}(G_i)$

$\log |G|$ iterations $\longrightarrow$ $|H| = \text{poly}(|G|)$

G 3col $\longrightarrow$ H 3-col

G not 3-col $\longrightarrow$ $\text{unsat}(H) \geq \Omega(1)$

$\square$

Main step $\quad G \xrightarrow{\text{balls}} BG \to G'$

1. gap-amplification / powering by balls

   $\text{unsat}(BG) > t \cdot \text{unsat}(G)$

2. to-3COL :

   $\text{unsat}(G') \approx \text{unsat}(BG)$

$$G \longrightarrow BG \qquad \text{"Balls of } G\text{"}$$

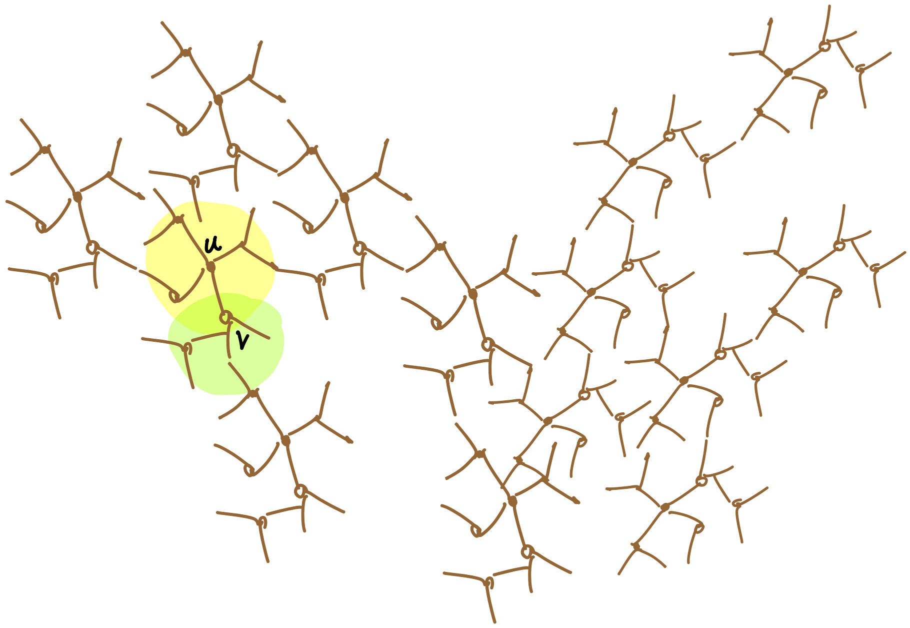

Label - Cover :

Def: a constraint graph is a tuple $G = (V, E, \mathcal{C}, \Sigma)$

(V, E) is a graph and $\forall\, uv \in E \quad C_{uv} \subseteq \Sigma \times \Sigma$

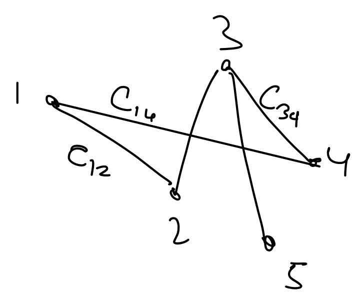

$$\text{sat}(G) = \max_{\sigma : V \to \Sigma} \text{Prob}\left[(\sigma(u), \sigma(v)) \in C_{uv} \right]$$

$$\text{unsat}(G) := 1 - \text{sat}(G)$$

The label cover problem:

find a labelling $\sigma : V \to \Sigma$ maximizing

$$G \longrightarrow BG$$

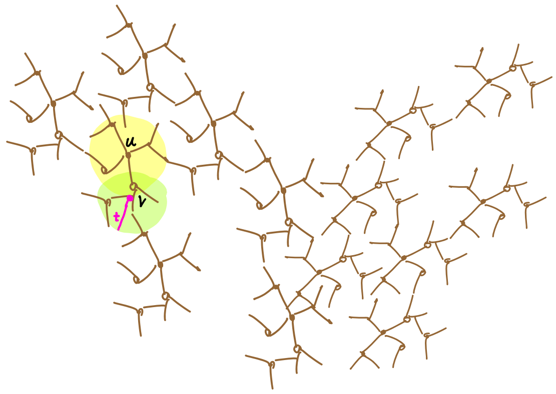

3-coloring $\rightsquigarrow$ label-cover ($|\Sigma| = \exp(t)$)

1 vs $1-\varepsilon$ $\rightsquigarrow$ 1 vs $1 - t \cdot \varepsilon$

agreement testing

to binary : G constraint graph over large Σ

$\downarrow$

G' constraint graph over $\{0,1\}$

$\approx$ same unsat value ($sat(G)=1 \rightarrow sat(G')=1$
$unsat(G') = \Omega(unsat(G))$)

$\Sigma = \{0,1\}^s$

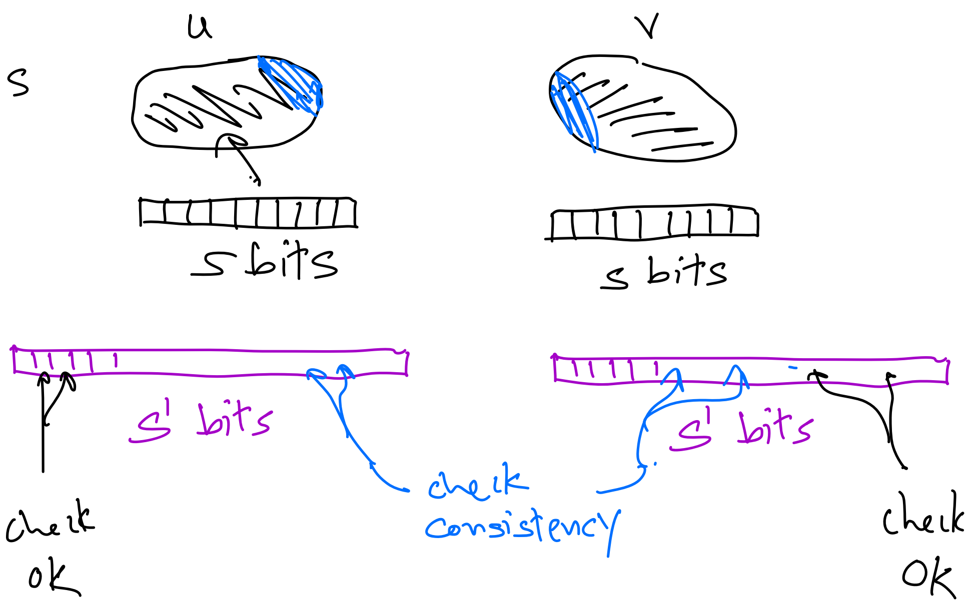

PCP encodings

[new bit $\forall$ quadratic function of old bits

$\#$ newbits $= 2^{s^2}$.]

1 values are ok

2 values are CONS

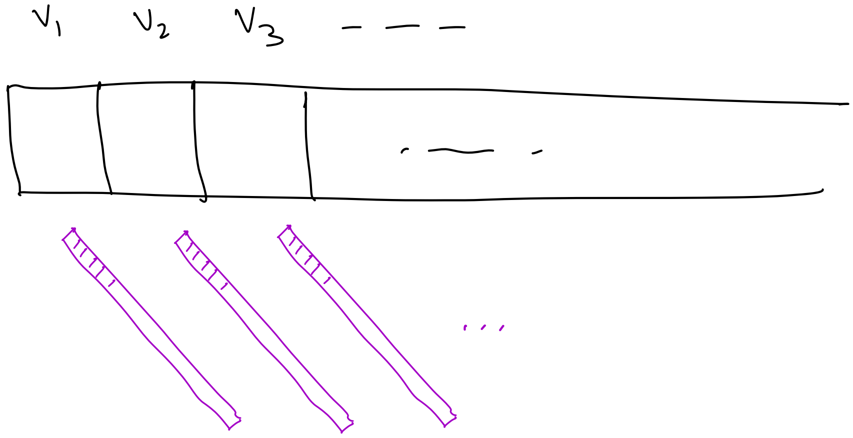

← the new proof is over $\{0,1\}$-vars.

but, the verifier queries $O(1)$ locations.

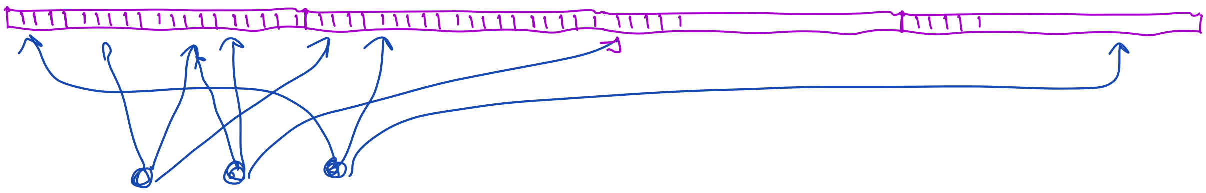

Convert to constraint graph:

$\mathcal{U}$ = vertex for each proof bit

V = vertex $\forall$ possible query pattern (how many ?)

For the rest of the notes, please see:

https://cs.nyu.edu/~khot/PCP-Spring20.html

using gadgets

inapprox for 3-COLORING $\Longrightarrow$

inapprox for 3SAT
max-cut
vertex cover
clique
ind. set

"gap-preserving" reduction

Q: Given a 3SAT formula
what's the best algorithm for
satisfying maximal # clauses?

(slides by Subhash Khot)

Håstad's 3-Bit PCP

Theorem NP has PCP verifier that

- uses $O(\log n)$ random bits
- 3 queries, linear predicate (over $\mathbb{F}_2$)
- $x \in L \Rightarrow \exists \, \Pi \quad \Pr[\text{Accept}] \geq 1-\varepsilon$

  $x \notin L \Rightarrow \forall \, \Pi \quad \Pr[\text{Accept}] \leq \frac{1}{2}+\varepsilon$

$$\equiv$$

Theorem Given 3-Lin instance S

$$\vdots$$

$$x \oplus y \oplus z = 0$$

$$x \oplus w \oplus u = 1$$

$$\vdots$$

It is NP-hard to distinguish bet^{n}

(YES) $\quad OPT(S) \geq 1-\varepsilon$

(NO) $\quad OPT(S) \leq \frac{1}{2}+\varepsilon$.