CMPUT 675: Computational Complexity Theory **Winter 2019**

Lecture 22 (Mar 28): The PCP Theorem: Part II

Lecturer: Zachary Friggstad *Scribe: Zachary Friggstad*

22.1 Preliminaries

We continue our proof of the PCP Theorem.

Recall, between Lectures 21 and 22, these notes begin with a slight recap of what we discussed at the end of lecture 20 and contain a bit of the discussion at the start of lecture 23. They also contain more detail than what was discussed in the lectures and they lift the assumption that the expanders are loopless. I wanted a comprehensive writeup you could use as reference if you wanted to verify the details (and it was good for me to be reminded of some details too).

22.2 Gap Amplification, Continued

Refer to the notes in Lecture 21 for notation. We already described the gap amplification construction and are now analyzing why $\text{sat}(\phi')$ is at most $1 - 4 \cdot \epsilon$ for appropriate settings of constants.

22.2.1 Analysis

Recall we are saying $\text{sat}(\phi) = 1 - \epsilon$.

Let $\overline{\sigma} : V \to [W]^{d^{t+1}}$ be an optimal assignment for ϕ'. For each $u \in V$, we determine an assignment $\sigma(u) \in [W]$ for u as follows:

- Perform **Lazy Random Walk 2** starting from u. If it ends at a vertex a that has an opinion for u, output a's opinion for u from $\overline{\sigma}(a)$. Otherwise, output `None`.
- Let $\sigma(u)$ be the most likely opinion for u (apart from `None`) that is output by this procedure. Break ties arbitrarily.

Claim 1 *Performing* **Lazy Random Walk 2** *from u ends at a vertex a having an opinion for u that matches $\sigma(u)$ with probability $\geq \frac{1}{2W}$*

Proof. The probability that **Lazy Random Walk 2** takes more than t steps is $(1 - 1/t)^{t+1} \leq 1/e \leq 1/2$. So with probability at least $1/2$, it ends at a vertex having an opinion for v.

Conditioning on the event it ends at a vertex having an opinion for u, the probability this opinion is $\sigma(u)$ is at least $\frac{1}{W}$ because we took the most popular opinion for u.

Thus, with probability at least $\frac{1}{2W}$, this walk ends at a vertex having opinion $\sigma(u)$ for u. ■

Now, σ leaves at least an ϵ-fraction of clauses of ϕ unsatisfied because $\text{sat}(\phi) = 1 - \epsilon$. Let $F \subseteq E$ be any particular set of $\epsilon \cdot |E|$ constraints that are not satisfied by σ.

Let us analyze the probability that the sampled constraint is not satisfied by $\overline{\sigma}$. Say a and b are the start and end vertices of the random walk used to sample the constraint. For sure, the constraint will not be satisfied if the walk took a step $u \to v$ with $uv \in F$ where:

- a has an opinion for u (i.e. the shortest $a - u$ path in G has length $\leq t$),
- b has an opinion for v, and
- a's opinion for u in $\overline{\sigma}(a)$ equals $\sigma(u)$, and b's opinion for v in $\overline{\sigma}(b)$ equals $\sigma(v)$.

Say such step is *faulty*. Let N_F be the random variable denote the number of steps of the constraint that are faulty. If $N_F > 0$, then the sampled constraint is not satisfied.

Main Goal: We will show $\mathbf{Pr}[N_F > 0] = \Omega(t) \cdot \epsilon$ where the Ω term suppresses constants depending on W and (d', λ') (the expansion of the constraint graph). We then set t appropriately so this is at least $6 \cdot \epsilon$ (unless ϵ is at least some constant already, as we will see).

22.2.2 The Second Moment Method

Our goal is to show $\mathbf{Pr}[N_F > 0]$ is large enough. The following gives a tool to do that. Intuitively, if we want a random variable to be positive, then having a high expected value would suffice unless the random variable is often zero and is only nonzero with a small probability (but on very large values). Being nonzero with very small probability but taking very high values in this case indicates the variance is very high.

So if the random variable has a high expected value but controlled second moment, then we have a good bound on the probability of it being nonzero.

Lemma 1 *Let X be a nonnegative random variable zero with $0 < \mathbf{E}[X^2] < \infty$. Then $\mathbf{Pr}[X > 0] \geq \frac{\mathbf{E}[X]^2}{\mathbf{E}[X^2]}$.*

Proof.
Quick Reminder: The set of all real-valued random variables Y with bounded second moments $\mathbf{E}[Y^2]$ in a probability space is a vector space over $\mathbb{R}$. Further, $\langle Y, Z \rangle = \mathbf{E}[Y \cdot Z]$ is an inner product in this space. So the Cauchy-Schwartz inequality applies: the proof we saw in Lecture 19 extends immediately to this case.

Let Z be the random variable that is 0 in any event X is zero and is 1 in any event where $X > 0$. Then:

$$\mathbf{E}[X]^2 = \mathbf{E}[X \cdot Z] = \langle X, Y \rangle^2 \leq ||X||_2^2 \cdot ||Y||_2^2 = \mathbf{E}[X^2] \cdot \mathbf{Pr}[X > 0].$$

■

22.2.3 Bounding $\mathbf{E}[N_F]$

Lemma 2 *If G has no loops, $\mathbf{E}[N_F] \geq \frac{t}{4W^2} \cdot \epsilon$.*

Proof. Fix a particular edge $uv \in F$ and direction $u \to v$ of that edge. Consider the random walk $a = v_0, v_1, \ldots, v_k = b$ when sampling a constraint. Let $B_{u,v}$ be the event that a's opinion for u is $\sigma(u)$ and b's

opinion for v is $\sigma(v)$: the event $B_{u,v}$ does not happen if either a or b does not have an opinion for u or v, respectively.

Let $Y^k_{u,v}$ be the event that the step $u \to v$ is taken exactly k times (the other direction $v \to u$ can be taken any number of times). By Lemma 6 from Lecture 21, the endpoints a, b from the **Lazy Random Walk 1** distribution conditioned on the event $Y^k_{u,v}$ are distributed independently and are distributed as though they were obtained by **Lazy Random Walk 2** starting from u, v respectively. So by Claim 1, $\mathbf{Pr}[B_{u,v}|Y^k_{u,v}] \geq \frac{1}{4W^2}$.

Let $N_{u,v}$ denote the number of times the step $u \to v$ is taken if $B_{u,v}$ happens, otherwise let it be 0. Then,

$$\begin{aligned}
\mathbf{E}[N_{u,v}] &= \sum_{k\geq 1} k \cdot \mathbf{Pr}[B_{u,v} \wedge Y^k_{u,v}] \\
&= \sum_{k\geq 1} k \cdot \mathbf{Pr}[B_{u,v}|Y^k_{u,v}] \cdot \mathbf{Pr}[Y^k_{u,v}] \\
&\geq \sum_{k\geq 1} \frac{k}{4W^2} \cdot \mathbf{Pr}[Y^k_{u,v}] \\
&= \frac{1}{4W^2} \cdot \mathbf{E}[\#u \to v \text{ steps}].
\end{aligned}$$

To compute the expected number of $u \to v$ steps, observe for each $i \geq 0$ that conditioned on the $(i+1)$'st step occurring, the step $v_i \to v_{i+1}$ is as if it was sampled uniformly at random among all possible edges and all possible directions of this edge. This is because each v_j is uniformly distributed (the random walk $v_0, v_1, \ldots, v_i$ starts from the uniform distribution and G is regular).

The probability that the $(i+1)$'st step occurs is also $(1-1/t)^i$. So the probability that $v_i \to v_{i+1}$ is even taken and is equal to $u \to v$ is $(1-1/t)^i/(2|E|)$.

Therefore, continuing from above we have

$$\mathbf{E}[N_{u,v}] \geq \frac{1}{4W^2} \sum_{i=0}^{\infty} \left(1 - \frac{1}{t}\right)^i \cdot \frac{1}{2|E|} = \frac{1}{4W^2} \cdot \frac{1}{2|E|} \sum_{i=0}^{\infty} (1-1/t)^i = \frac{t}{4W^2} \cdot \frac{1}{2|E|}.$$

Finally, $N_F = \sum_{uv \in F} N_{u,v} + N_{v,u}$ so

$$\mathbf{E}[N_F] = \sum_{uv \in F} \mathbf{E}[N_{u,v}] + \mathbf{E}[N_{v,u}] \geq \frac{t}{4W^2} \cdot \frac{2|F|}{2|E|} = \frac{t}{4W^2} \cdot \epsilon.$$

■

With Loops

Now we consider the possibility of G having loops. We still have that v_i is sampled from the uniform distribution (if we condition the random walk on reaching step i), this holds even in the presence of loops. So step $v_i v_{i+1}$ is a step across some edge in F in some direction with probability at least $1/2|E|$ (rather than exactly this amount), given that step $i+1$ was taken at all.

The only other adjustment is that N_F is the sum of at least $|F|$ (rather than exactly $2|F|$) terms of the form $N_{u,v}$ because $u \to v$ and $v \to u$ are identical steps if uv is a loop. That is, recall in the adjacency matrix of a graph with loops that each loop contributed only 1 to the diagonal entry: we think of a loop as only having 1 endpoint.

Lemma 3 $\mathbf{E}[N_F] \geq \frac{t}{8W^2} \cdot \epsilon$, *even if G has loops.*

22.2.4 Bounding $\mathbf{E}[N_F^2]$

Recall that the constraint graph G of ϕ is an (n, d', λ')-expander.

Lemma 4 *If G has no loops, $\mathbf{E}[N_F^2] \leq 2 \cdot (C' + t \cdot \epsilon) \cdot (t \cdot \epsilon)$ where C' is a constant depending only on d' and λ'.*

Proof. For every $i \geq 0$, let χ_i be the $\{0,1\}$-indicator for the event that step $v_i v_{i+1}$ was taken and that it used an edge in F. Clearly $N_F \leq \sum_{i \geq 0} \chi_i$ always holds, so we proceed as follows.

$$\begin{aligned}
\mathbf{E}[N_F^2] &\leq \sum_{i,j} \mathbf{E}[\chi_i \cdot \chi_j] \\
&\leq 2 \cdot \sum_{0 \leq i \leq j} \mathbf{E}[\chi_i \cdot \chi_j] \\
&= 2 \cdot \sum_{i=0}^{\infty} \mathbf{Pr}[\chi_i = 1] \cdot \sum_{\ell \geq 0} \mathbf{Pr}[\chi_{i+\ell} = 1 | \chi_i = 1]
\end{aligned}$$

We bound the inner term $\mathbf{Pr}[\chi_{i+\ell} = 1 | \chi_i = 1]$. Consider **Lazy Random Walk 1** conditioned on the event that step $v_i v_{i+1}$ is taken and used F. Now, v_i is distributed uniformly in this distribution so step $v_i v_{i+1}$ is uniformly distributed among all edges in F and their two directions. That is, the random walk starting from v_{i+1} is like we chose v_{i+1} by uniformly choosing an edge in F and then randomly picking one of its endpoints to be v_{i+1}.

The assignment has you analyze such a random walk. It asks you to show that if you start a random walk $v'_0, v'_1, \ldots$ by first randomly picking F and a random endpoint of it, and then proceeding to perform a normal random walk after that, the probability that step $v'_\ell v'_{\ell+1}$ crosses an edge in F is at most $\frac{|F|}{|E|} + \lambda'^{\ell}$.

Interpreting this in our setting, this means for each $\ell \geq 0$

$$\mathbf{Pr}[\chi_{i+\ell} | \chi_i = 1] \leq \left(1 - \frac{1}{t}\right)^{\ell} \left(\frac{|F|}{|E|} + \lambda'^{\ell}\right) = \left(1 - \frac{1}{t}\right)^{\ell} (\epsilon + \lambda'^{\ell})$$

where we recall the $(1-1/t)^\ell$ term comes from the fact that **Lazy Random Walk 1** stops after each step with probability $1/t$. So

$$\sum_{\ell \geq 0} \mathbf{Pr}[\chi_{i+\ell} = 1 | \chi_i = 1] = \sum_{\ell \geq 0} \left(1 - \frac{1}{t}\right)^{\ell} \left(\frac{|F|}{|E|} + \lambda'^{\ell}\right) = C' + \epsilon \sum_{\ell \geq 0} (1 - 1/t)^{\ell} = C' + t \cdot \epsilon$$

where C' depends only on d', λ'.

Overall, we can continue our bound

$$\mathbf{E}[N_F^2] \leq 2(C' + t \cdot \epsilon) \sum_{i=0}^{\infty} \mathbf{Pr}[\chi_i = 1] = 2(C' + t \cdot \epsilon) \sum_{i=0}^{\infty} (1 - 1/t)^i \cdot \epsilon = 2(C' + t \cdot \epsilon) \cdot (t \cdot \epsilon).$$

■

With Loops
Now, if G has loops we again need a small revision. Conditioning on $\chi_i = 1$ does not mean step $v_i v_{i+1}$ is sampled by taking a random edge in F and picking a random endpoint. It is true that, before conditioning, v_i is a uniform vertex (given that it the walk reached step i). Say F has k_1 loops and k_2 nonloops.

Conditioning only on taking step i, the edge v_iv_{i+1} is obtained by sampling a random vertex and a random incident edge. So the probability F is taken is $\frac{k_1+2k_2}{|E|}$. Think of starting a random walk in G as follows: among all $k_1 + 2k_2$ endpoints of edges in F (counting the same vertex multiple times if it is the endpoint of multiple edges in F), pick one randomly. Then continue a normal random walk from this vertex. The **brownie points** version of the assignment question asks you to prove that step $v_\ell v_{\ell+1}$ crosses an edge in F with probability at most $\frac{|F|}{|E|} \leq 4 \cdot \lambda'^{\ell}$.

The other part of the proof that needs to be modified is $\mathbf{Pr}[\chi_i = 1]$ is not exactly $(1 - 1/t)^i \cdot \epsilon$, but is instead at most $2 \cdot (1 - 1/t)^i \cdot \epsilon$ for similar reasons. But all of this slight loss can be absorbed in the constant.

Lemma 5 $\mathbf{E}[N_F^2] \leq 2 \cdot (C + t \cdot \epsilon) \cdot (t \cdot \epsilon)$ *where C is a constant depending only on d' and λ', even if G has loops.*

22.2.5 Finishing the Bound

Recall $\mathbf{Pr}[N_F]$ is a lower bound on the total weight/probability of constraints of ϕ' that are not satisfied by $\overline{\sigma}$. That is, $\text{sat}(\phi') \leq 1 - \mathbf{Pr}[N_F > 0]$. From Lemmas 3 and 5 and the second moment method Lemma 1, we immediately conclude

Corollary 1 $\mathbf{Pr}[N_F > 0] \geq \left(\frac{t^2}{64W^4}\epsilon^2\right) / \left(2(C + t \cdot \epsilon) \cdot (t \cdot \epsilon)\right) = \frac{\beta \cdot t \cdot \epsilon}{2(C+t\epsilon)}$ *where β depends only on W.*

In general, the best we can say is $\epsilon \leq 1$ which lower-bounds this probability by $\beta' \cdot \epsilon$ for some constant β' depending on W, d, λ'. In other words, if the unsatisfiability of ϕ was low enough (i.e. if ϵ is big enough) then we don't seem to get any controllable improvement.

But, if $\epsilon \leq 1/t$ then in fact we can say $\mathbf{Pr}[N_F > 0] \geq \beta'' \cdot t \cdot \epsilon$ where β'' depends on W, D and λ' but not t. So we are free to set t to be any constant.

Choice of t: Set $t := \lceil 4/\beta'' \rceil$.

So if $\epsilon \leq t$ then $\text{sat}(\phi') \leq 1 - 4\epsilon$. If $\epsilon > t$, we can still say $\text{sat}(\phi') \leq 1 - \beta'/t$. Set $\epsilon' := \beta'/t > 0$, the constant mentioned in the **Output** specification of the start of this Gap Amplification. Again, ϵ' depends only on W, d', λ' yields

Corollary 2 $sat(\phi') \leq 1 - \min\{4\epsilon, \epsilon'\}$.

22.2.6 Truncating the Walk

The last thing that must be done is truncating the random walk so it generates only $O(n)$ constraints. With foresight, set τ to be minimum such that $(1 - 1/t)^\tau \cdot (\tau + t) \cdot 2 \leq \frac{1}{16W^2}$. The value $\tau = \lceil 10 \cdot t \cdot \log_2 W \rceil$ suffices, just remember $(1 - 1/t)^t \leq e^{-1}$ when you verify this.

Constraints - Real Version

Perform **Lazy Random Walk 1** except force it to stop if it takes more than τ steps. Let $a = v_0, v_1, \ldots, v_k = b$ be the walk and note $k \leq \tau$. Add an $a - b$ edge/constraint to ϕ'.

If the random walk stopped within τ steps, the constraint is the same as before: it is satisfied if and only if for *every* step v_iv_{i+1} for which a has an opinion for v_i and b has an opinion for v_{i+1}, the v_iv_{i+1} constraint in ϕ is satisfied by this opinion.

If the random walk did not stop within τ steps and we had to truncate it, this new constraint is always satisfied.

CL-Reduction
Note, for each u there are at most $d'^{\tau} + \sum_{i=0}^{\tau} d'^{i}$ constraints generated with u as the start vertex. The first term is for the truncated constraints and the rest are for the normal constraints. Also, if ϕ is satisfiable then ϕ' is as well for the same reason: if $\sigma : V \to [W]$ satisfies ϕ then choose $\overline{\sigma} : V \to [W]^{d'^{t+1}}$ just like before. That is, $\overline{\sigma}(u)$ should set u's opinion for v to $\sigma(v)$ for every v that is within distance at most t from u. So this is now a CL-reduction.

We just have to analyze $\text{sat}(\phi')$ for this new instance. Let N'_F be the random variable that is 0 if the random walk went too long, otherwise it agrees with N_F. Note $N'_F \leq N_F$ so we can bound $\mathbf{E}[N'^2_F]$ by the same bound[1] that we used for $\mathbf{E}[N_F^2]$. We want to show for this choice of τ that $\mathbf{E}[N'_F] \geq \mathbf{E}[N_F]/2$.

We measure the expected difference of N_F and N'_F. Let χ_i be the $\{0,1\}$ random variable that indicates if $v_i v_{i+1}$ occurred and used an edge in F. We already discussed how $\mathbf{Pr}[\chi_i = 1] \leq (1 - 1/t)^i \cdot 2\epsilon$ (it would be exactly $(1-1/t)^i \cdot \epsilon$ if G had no loops). Observe,

$$\mathbf{E}[\sum_i \chi_i | \geq \tau + 1 \text{ steps taken}] \leq 2\epsilon \cdot \tau + \sum_{\ell \geq 0} \mathbf{Pr}[\chi_\ell] \leq 2\epsilon(\tau + t).$$

Thus,

$$\begin{aligned}
\mathbf{E}[N_F - N'_F] &\leq \mathbf{Pr}[\geq \tau + 1 \text{ steps taken}] \cdot \mathbf{E}[\#\text{ steps in } F | \geq \tau + 1 \text{ steps taken}] \\
&\leq (1 - 1/t)^\tau \cdot (\tau + t) \cdot 2\epsilon \\
&\leq \frac{1}{16W^2} \cdot \epsilon \\
&\leq \frac{\mathbf{E}[N_F]}{2}
\end{aligned}$$

The last bound holds by Lemma 5. Therefore, $\mathbf{E}[N'_F] \geq \frac{\mathbf{E}[N_F]}{2} \geq \frac{1}{16W^2} \cdot \epsilon$.

The rest of the proof proceeds as before, with perhaps very minor adjustments made to the choice of t and the final constant $\epsilon' > 0$.

22.2.7 Making ϕ' Unweighted

The reduction created a weighted CSP instance ϕ' where $\text{sat}(\phi') \leq 1 - \min\{6\epsilon, \epsilon'\}$ meant the maximum *weight* of clauses that can be satisfied is at most the stated term.

Observe the probability any clauses is sampled is an integer multiple of $\left(\frac{1}{d' \cdot t}\right)^\tau$ where t, τ are constants that depend only on d', λ' and W. That is, the random walk would take up to τ steps, each step would first pick a random neighbour of the vertex (with probability $1/d$ each) and then flip a coin with bias $1/t$ to see if it should stop. So if p_e denotes the probability a particular clause/edge e is sampled from the random walk creation, then $(d' \cdot t)^\tau \cdot p_e$ is an integer.

Replace each such clause e with $(d' \cdot t)^\tau \cdot p_e$ copies of itself to get an unweighted CSP instance that is only bigger than ϕ' by a $(d' \cdot t)^\tau = O(1)$ factor. This is our final unweighted CSP from the gap amplification step.

[1] Technically, to do this we should finish the normal random walk without truncating, but output the trivially-satisfied constraint connecting v_0 to v_τ if the walk went too long. Then N_F remains defined as before.

22.3 Alphabet Reduction

Input: A 2CSP_D instance ϕ'. Say $\text{sat}(\phi') = 1 - \epsilon'$.
Output: A $q_0\text{CSP}_2$ instance ϕ'' where $\text{sat}(\phi'') \leq 1 - \epsilon'/3$.

We assume $D = 2^k$ for some integer k. We can do this by "padding": adding extra items to the alphabet such that no clause of ϕ' is satisfied by any assignment that uses one of these new terms. This does not change $\texttt{sat}(\phi')$. In fact, we now view the alphabet for ϕ' as $\{0,1\}^k$.

Thankfully this is very easy compared to the last section. That is, we have already done most of the work in a previous lecture. The idea is that we will turn each constraint into a circuit and use the PCP verifier from an earlier lecture to check that a setting of the values to the wires of the circuit. The PCP we use is the one that used exponential-size proofs. However, here "exponential-size" will be still be a constant because the alphabet size of ϕ', namely D, is a constant itself. This PCP will still query $O(1)$ bits of the proof where the $O(1)$ does not depend on D: it is a fresh constant.

22.3.1 Constraints as Circuits

Consider a constraint/edge e in ϕ' and say it depends on variables u, v. View the constraint as a circuit C_e taking in $2k$ inputs, one for the encoding of a value for u and one for the encoding of a value for v. Say C_e has ℓ_e wires including the output wire apart from the $2k$ input wires, we can take ℓ to be bounded by $O(2^{2k})$ (the $O()$ hiding an absolute constant independent of k).

Note, it could be that some constraints e of ϕ' are "loops" meaning they depend only on one variable. In such a case, still consider the circuit C_e to have $2k$ inputs and additionally require (via gates in the circuit) that both groups of k inputs encode the same value. value.

Observe: We can now write a system of strictly quadratic constraints whose variables are all $2k + \ell_e$ wires of C_e such that all constraints are satisfied by an assignment if and only if the values on all wires are consistent (i.e. each non-input wire is correct as a function of the wires entering the corresponding gate) and the output wire is 1. Compactly write this system as $A_e \cdot (x \otimes x) = b_e$ where x is a vector of $2k + \ell_e$ variables over $\{0,1\}$ integers mod 2. We already did this earlier in the course. Technically we reduced from $\texttt{SAT}$, but the exact same idea works when we start from circuits with the observation that we can get away with using exactly as many variables as we have wires.

22.3.2 Variables of ϕ''

Recall for $x \in \{0,1\}^s$ for some s that $\texttt{WH}_x : \{0,1\}^s \to \{0,1\}$ is the function mapping $\texttt{WH}_x(y) = x \circ y$ where we take $x \circ y = \sum_{i=1}^{s} x_i \cdot y_i \pmod 2$.

The variables of ϕ'' come in three categories:

- For each variable v of ϕ', we have 2^k variables giving the truth table of some function $f_v : \{0,1\}^k \to \{0,1\}$. The intent is that f_v should be $\texttt{WH}_a$ for some $a \in \{0,1\}^k$ (i.e. the Walsh-Hadamard encoding of a value v should take in ϕ').
- For each constraint e of ϕ' depending on, say, variables u, v we have $2^{2k+\ell_e}$ variables giving the truth table of some function $g_e : \{0,1\}^{2k+\ell_e} \to \{0,1\}$. The intent is that g_e should be the Walsh-Hadamard encoding of a satisfying assignment $y \in \{0,1\}^{2k+\ell_e}$ to the wires of the circuit C_e.

By satisfying, we mean the value assigned to each wire by y is correct (as a function of the wires entering the corresponding gate) and the output wire for C_e is 1.

- Again for each constraint e of ϕ' depending on variables u, v, we have a further $2^{(2k+\ell_e)^2}$ variables giving the truth table of some function $h_e : \{0,1\}^{(2k+\ell_2)^2} \to \{0,1\}$. The intent is that h_e is the Walsh-Hadamard encoding of $y \otimes y$ where g_e encodes $\mathtt{WH}_y$.

22.3.3 Constraints of ϕ''

For each constraint $e = uv$ of ϕ, we will verify that f_u, f_v, g_e, h_e are "correct" in the following way. Every $O(1)$ term below does not depend on k, it is a fresh constant.

- Run the linearity test from Lecture 15.3.3 on each of f_u, f_v, g_e, and h_e $\Theta(1)$ times so that if any of them is not 0.999-close to being linear, then it is rejected with probability at least 1/2.

  Recall from Lecture 15 that if a function $\Theta(1)$ rounds of the linearity test then it is 0.999-close to the function $\mathtt{HW}_a$: here $\Theta(1)$ is an appropriately large constant. In fact, we have not yet seen a proof of this but it will come next lecture.

  From now on, we suppose f_u, f_v, g_e, h_e are 0.999-close to the $\mathtt{WH}$-encodings of a, b, y, y', respectively, where $a, b \in \{0,1\}^k, y \in \{0,1\}^{2k+\ell_2}$ and $y' \in \{0,1\}^{(2k+\ell_e)^2}$.

- Perform the **concatentation test** (see below) with f_u, f_v, g_e to check that the first k bits of y equal a and the second k bits of y equal b. This is to ensure the input to the circuit C_e is in fact given by the values of u and v.

- Using local decoding on g_e, h_e, perform the test from Lecture 15.3.2 $\Theta(1)$ times to check that y' encodes $y \otimes y$. This was described in Lecture 15 and we prove that if $y' \neq y \otimes y$, then $\Theta(1)$ repetitions of this test rejects with probability $\geq 1/2$.

- Using local decoding on h_e, perform the test from Lecture 15.3.1 $\Theta(1)$ times to check that $y \otimes y$ indeed satisfies $A_e \cdot (y \otimes y) = b_e$. This was also described in Lecture 15 and we proved that if $y' = y \otimes y$ yet $A_e \cdot (y \otimes y) \neq b_e$, then this test fails with probability $\geq 1/2$.

Summary: Intuitively, the first test checks that the functions are $\mathtt{WH}$ encodings of particular assignments. The second test ensures the input wires to the circuit take the same values as the variables u, v themselves (i.e. the input to the circuit is given by the values of variables u, v). The third and fourth test then ultimately check that the values assigned to the wires of the circuit indeed satisfy the circuit.

22.3.4 Concatenation Test

Input: Functions f_u, f_v, g_e on k, k and $2k+\ell_e$ bits, respectively, such that each is 0.999 close to a linear function.

Perform the following twice.

- Sample $x, x' \sim \{0,1\}^k$ uniformly and independently.
- Construct vector z of length $2k + \ell_e$ by setting the first k bits to x, the second k bits to x', and the remaining ℓ_e bits to 0.
- Sample $\zeta, \zeta' \sim \{0,1\}^k$ and $\xi \sim \{0,1\}^{2k+\ell_e}$ uniformly and independently.

- Reject if
$$f_u(x+\zeta) + f_u(\zeta) + f_v(x'+\zeta') + f_v(\zeta') \neq g_e(z+\xi) + g_e(\xi).$$

Say that f_u, f_v and g_e are 0.999-close to a, b, y, respectively. If the first $2k$ bits of y do not equal the concatenation of a and b, then by the random subsum principle (Lemma 2 from Lecture 15), $(a \circ x) + (b \circ x') \neq (y \circ z)$ with probability $1/2$ over the choice of x, x', z. Further, the probability over ζ, ζ', ξ that any of the 6 queries to f_u, f_v or g_e in the last step above does not give the correct value from the WH-encoding of the corresponding function is at most $6 \cdot 0.001 = 0.006$ (union bound over a single query failing).

So if the first $2k$ bits of y are not the concatenation of a and b then the probability a single iteration does not reject is at most 0.5006. Iterating this twice brings the probability to $< 1/2$. Formally,

Lemma 6 *If f_u, f_v and g_e are 0.999-close to linear functions encoding a, b, y, respectively, but the first $2k$ bits of y is not equal to the concatenation of a, b, then the above test rejects with probability $\geq 1/2$.*

22.3.5 Finishing the Constraints of ϕ''

For each constraint $e = uv$ of ϕ', consider running the four tests from 22.3.3 in parallel. To generate all of these tests, we flipped $r_e = \text{poly}(k, \ell_e) = O(1)$ coins. Each test also depended on $O(1)$ variables of ϕ'', i.e. each test queried the functions f_u, f_v, g_e, h_e $O(1)$ times. Let $q_0 = O(1)$ be an upper bound on the number of distinct queries performed in a single test (some tests query the same bit more than once).

For each of the 2^{r_e} possible ways to flip the coins, generate a constraint for ϕ'' that depends on $\leq q_0$ variables of ϕ'' as follows. The constraint is satisfied by an assignment to these bits if and only if the corresponding tests passed.

It is clear that if ϕ' is satisfiable, then so to is ϕ''. That is, simply let all f_v be the WH encodings of a satisfying assignment for ϕ', let all g_e be the WH encodings of ways to satisfy the corresponding circuit C_e, and let h_e be all the WH encodings of the $\otimes$-product of these corresponding assignments. All tests pass with probability 1, so all constraints are satisfied.

Thus, ϕ'' is a $q_0\text{CSP}_2$ instance whose size is only a constant-factor larger than the size of ϕ'. This is because each constraint $e = uv$ of ϕ' has 2^{r_e} corresponding constraints in ϕ''. So this is a CL-reduction.

22.3.6 Soundness Analysis

Lemma 7 *Say $sat = 1 - \epsilon$, then $sat(\phi'') \leq 1 - \epsilon/2$.*

Proof. Consider any functions $\{f_v\}_{v \in V}$ and $\{g_e, h_e\}_{e \in E}$. We get an assignment $\sigma : V \to [D]$ for ϕ' as follows. For each $v \in V$, if f_v is 0.999-close to WH_a for some $a \in \{0,1\}^k$ then set $\sigma(v) = a$. Otherwise, set $\sigma(v)$ to anything.

Let $\delta \geq \epsilon$ be the fraction of clauses of ϕ' that are not satisfied by σ. For each $e = uv$ that is not satisfied by σ, one of the following must hold:

- One of f_u, f_v, g_e, h_e is not 0.999-close to a linear function.
  Otherwise, let us say that f_u, f_v, g_e, h_e are 0.999-close to the WH-encodings of $\sigma(u), \sigma(v), y, y'$.
- The first $2k$ bits of y are not equal to the concatenation of $\sigma(u)$ and $\sigma(v)$.
- $y' \neq y \otimes y$.

- Finally, if we get here then because e is not satisfied by σ. As none of the previous cases applied, then e is, in particular, not satisfied by $\sigma(u)$ and $\sigma(v)$ so some wire of the circuit is incorrectly given by y (whose first $2k$ bits are $\sigma(u)$ and $\sigma(v)$). As $y' = y \otimes y$, it must be that $A_e \cdot y' \neq b_e$.

In any case, we have that f_u, f_v, g_e, h_e will be rejected by the tests in 22.3.3 with probability $\geq 1/2$. That is, at least half of the constraints generated from this edge e will not be satisfied by f_u, f_v, g_e, h_e.

So, in total, at least a $\delta/2 \geq \epsilon/2$-fraction of constraints of ϕ'' are not satisfied by this assignment as at least $1/2$ of the constraints coming from each of the ϵ-fraction of edges of ϕ' are not satisfied. As this holds for any functions $\{f_v\}_{v \in V}$ and $\{g_e, h_e\}_{e \in E}$, $\mathtt{sat}(\phi'') \leq 1 - \epsilon/2$. ■

22.4 Post-Mortem Discussion

The preparation and alphabet reduction steps did not assume the ϵ in $\mathtt{sat}(\phi) = 1 - \epsilon$ was bounded by a constant, it was only the gap amplification step. If we start with a $q_0\mathrm{CSP}_2$ instance ϕ with $\mathtt{sat}(\phi) = 1 - \epsilon$, we have the resulting instance ψ after applying all reductions to the end of alphabet composition satisfies

$$\mathtt{sat}(\psi) \leq 1 - \frac{1}{2} \cdot \min\{4\epsilon, \epsilon'\} = 1 - \min\{\epsilon'/2, 2\epsilon\}.$$

where ϵ' was defined in Section 22.2.5.

To conclude the proof of the Lemma **??**, simply set $\epsilon_0 = \epsilon'/2$.

A lot of constants were thrown about during these proofs. But these were eventually just tucked into the linear blowup in the size of the reduction. That is, say ϕ and ϕ' are the "before and after" $q_0\mathrm{CSP}_2$ instances referred to in Lemma **??** and say ϕ has m constraints. Then ϕ' has $\leq \gamma \cdot m$ constraints where γ ultimately only depended on q_0 and the parameters (d, λ) of the expander family $\mathcal{G}$ we used. The "stopping point" ϵ_0 also only depends on q_0 and (d, λ).

The value q_0 is an absolute constant one could work out with some diligence, it came from the number of queries in the alphabet reduction step which was totally independent of (d, λ) and the "previous" q_0.

References

AB09 S.Arora and B.Barak, Computational Complexity: A Modern Approach, *Cambridge University Press,New York, NY, USA*, 2009, pp. 126–151.

ALMSS98 S.Arora, C.Lund, R.Motwani, M.Sudan, and M. Szegedy, Proof verification and the hardness of approximation problems, Journal of the ACM 45 (3): 501–555, 1998.

AS98 S.Arora and S.Safra, Probabilistic checking of proofs: A new characterization of NP, Journal of the ACM, 45 (1): 70–122, 1998.

D07 I.Dinur, The PCP theorem by gap amplification, Journal of the ACM, 54 (3): 12, 2007.

GO05 V.Guruswami and R.O'Donnell, CSE 533: The PCP Theorem and Hardness of Approximation, `https://courses.cs.washington.edu/courses/cse533/05au/`, 2005.

R06 J.Radhakrishnan, Gap amplification in PCPs using lazy random walks, In Proceedings of ICALP, 96–107, 2006.