PCP Theorem by Gap Amplification

Bernhard Vesenmayer *
JASS 2006

Abstract

The PCP Theorem provides a new classification of NP. Since the original proof by [AS98], several new proofs occured. While the first proof used highly sophisticated methods, the new approaches try to use simpler ones. In this paper the proof by [Din05] is presented. It uses the equivalence of this problem to the NP-hardness of gap-3SAT, i.e. it is this hardness that is shown. The satisfiability gap of a set of constraints is the smallest fraction of unsatisfied constraints over all assignments for the variables. Gap-3SAT is the problem of deciding whether this gap is 0 or greater than a positive constant. In this paper therefore 3SAT is polynomially reduced to gap-3SAT via gap amplification, i.e. the gap is blown up to a constant fraction.

1 Introduction

There are by now many structural different proofs of the famous PCP-Theorem. Most of them are quite technical and use highly sophisticated methods. Originally it has been stated via interactive proofs. In [FGL+96] a surprising connection has been found between this theorem and a formulation via gap-3SAT. That means they have shown, that the PCP-Theorem is equivalent to stating that gap-3SAT is NP-hard. In other words, it is NP-hard to distinguish between $\mathrm{UNSAT}(\mathscr{C}) = 0$ and $\mathrm{UNSAT}(\mathscr{C}) = \alpha > 0$ for a constraint system $\mathscr{C}$. This is the basis of the proof by [Din05] that is presented in this paper. Naturally, if the constraint system is not satisfiable, then $\mathrm{UNSAT}(\mathscr{C}) \geq \frac{1}{n}$, where n is the number of constraints in $\mathscr{C}$. Via a polynomial amplification algorithm this gap is blown up to a constant α which renders this problem equivalent to gap-3SAT.

This paper is part of a talk at JASS'06. It is therefore not self-containing, but relies on the preparative paper [Bul06] in which underlying structures and basic results are presented. Only the existence of assignment tester is used without proof. The proofs of the other results are mostly taken out of [Din05] and can be found here or in [Bul06].

2 Preprocessing

In order to apply the main amplification lemma, we need a 'nicely' structured graph. In this section is shown that any constraint graph can be polynomially transformed into such

*TU München, Boltzmannstraße 3, 85747 Garching bei München, Germany, (e-mail: vesenmay@ma.tum.de).

a 'nice' graph. By nicely structured is meant that the graph is regular, of constant-degree and expanding. The transformation includes two steps which are formalized in the next two propositions.

Proposition 2.1 (Constant degree) *Any constraint graph $G = \langle (V, E), \Sigma, \mathscr{C} \rangle$ can be transformed into a $(d_0 + 1)$-regular graph $G' = \langle (V', E'), \Sigma, \mathscr{C}' \rangle$ such that $|V'| = 2|E|$ and*

$$c \cdot \mathrm{UNSAT}(G) \leq \mathrm{UNSAT}(G') \leq \mathrm{UNSAT}(G)$$

for some global constants $d_0, c > 0$.

Proof. Fix d_0 and define $d := d_0 + 1$. For each n let X_n be a d_0-regular expander on n vertices with $h(X_n) \geq h_0$ as is guaranteed by [Bul06, Lemma 17]. Let d_v be the degree of $v \in V$. Replace each v by X_{d_v} and put equality constraints on these edges. Denote this with $[v]$. Let $[V] := \cup_{v \in V}[v]$ and E_1 be the union of the edges. For every $(v, w) \in E$ put an edge between one of the vertices of $[v]$ and $[w]$, such that each such vertice only sees one such external edge. The constraints on these edges are given by the constraints of the original graph. Denote this set of edges by E_2. That means we have constructed a d-regular graph G':

$$G' := ([V], \mathbf{E} = E_1 \cup E_2)$$

Let us now show that this graph meets the requirements. We start with the upper bound. Let $\sigma : V \to \Sigma$ be an assignment of G. Define $\sigma' : [V] \to \Sigma$ by:

$$\forall v \in V, x \in [v] : \sigma'(x) = \sigma(v)$$

Since the constraints of the additional edges of E_1 are met in this case, we see that the fraction of unsatisfied edges decreases or stays constant. That means:

$$\mathrm{UNSAT}(G') \leq \mathrm{UNSAT}(G)$$

The other bound is more complex. Let $\sigma' : [V] \to \Sigma$ be a best assignment. Define $\sigma : V \to \Sigma$ by:

$$\forall v \in V : \sigma(v) := \mathrm{maxarg}_{a \in \Sigma}(P_{x \in [v]}(\sigma'(x) = a))$$

That means $\sigma(v)$ is the most popular value of $[v]$. Let $F \subset E$ be the set of edges that reject σ, $\mathbf{F} \subset \mathbf{E}$ be the set of edges that reject σ'. Let S be the set of those vertices in $[V]$ that were outvoted in the definition of σ, i.e. $S := \cup_{v \in V}\{x \in [v]; \sigma'(x) \neq \sigma(v)\}$. Every external edge in E_2 corresponding to an edge in F either rejects σ' or has at least one endpoint in S. This gives us:

$$|\mathbf{F}| + |S| \geq |F| = \alpha \cdot |E|,$$

where $\alpha := \frac{|F|}{|E|}$. First recall that $|\mathbf{E}| = d|E|$. We now have two cases:

1. $|\mathbf{F}| \geq \frac{\alpha}{2}|E| = \frac{\alpha}{2d}|\mathbf{E}|$: Therefore we have the desired lower bound:

$$\mathrm{UNSAT}(G') = \mathrm{UNSAT}_{\sigma'}(G') \geq \frac{1}{2d}\mathrm{UNSAT}_{\sigma}(G) \geq \frac{1}{2d}\mathrm{UNSAT}(G).$$

2. $|\mathbf{F}| < \frac{\alpha}{2}|E|$: Then $|S| \geq \frac{\alpha}{2}|E|$. Fix $v \in V$ and define:

$$\begin{aligned} S^v &:= [v] \cap S \\ S^v_a &:= \{x \in S^v; \sigma'(x) = a\} \end{aligned}$$

Then we have $|S^v_a| \leq \frac{|[v]|}{2}$. Therefore the edge expansion property of X_{d_v} yields:

$$|E(S^v_a, [v] \setminus S^v_a)| \geq h_0 \cdot |S^v_a|.$$

All of those edges reject σ', because of their equality constraints. That means:

$$|\mathbf{F}| \geq \sum_{v \in V} \sum_{a \in \Sigma} h_0 |S^v_a| = \sum_{v \in V} h_0 |S^v| = h_0 |S| \geq h_0 \frac{\alpha}{2} |E| = \frac{h_0 \alpha}{2d} |\mathbf{E}|$$

And therefore the lower bound follows as above: $\mathrm{UNSAT}(G') \geq \frac{h_0}{2d}\mathrm{UNSAT}(G)$

Recalling that d is independent of G, the proposition follows with $c := \min(\frac{1}{2d}, \frac{h_0}{2d})$. □

Proposition 2.2 (Expanderizing) *Let $d_0, h_0 > 0$ be some global constants. Any d-regular constraint graph G can be transformed into G' such that*

1. *G' is $(d+d_0+1)$-regular, has self-loops, and $\lambda(G') \leq d+d_0+1-\frac{h_0^2}{d+d_0+1} < \deg(G')$,*
2. *$size(G') = O(size(G))$, and*
3. *$\frac{d}{d+d_0+1} \cdot \mathrm{UNSAT}(G) \leq \mathrm{UNSAT}(G') \leq \mathrm{UNSAT}(G)$.*

Proof. This transformation will be done by adding edges with trivial constraints. Let $E_{loop} := \{(v, v); v \in V\}$ and $X = (V, E')$ be a d_0-regular expander on $|V|$ vertices with $h(X) \geq h_0$ (again, the existence is guaranteed by [Bul06, Lemma 17]). Define:

$$G' = (V, E \cup E' \cup E_{loop})$$

This clearly is a $(d + d_0 + 1)$-regular graph of the same size as G. We can now apply the general result for expanders (see [Bul06, Lemma 19]) and get:

$$\lambda(G') \leq d + d_0 + 1 - \frac{h(G')^2}{d(G')} \leq d + d_0 + 1 - \frac{h_0^2}{d + d_0 + 1} < d + d_0 + 1$$

As the new edges are always satisfied and the number of edges is increased by a factor $\frac{d+d_0+1}{d}$, the fraction of unsatisfied constraints drops at most by this factor. □

Now we can finally state the lemma that allows us to transform a given constraint graph into a nicely-structured one.

Lemma 2.3 (Preprocessing) *There exist constants $0 < \lambda < d$ and $\beta_1 > 0$ such that any constraint graph G can be transformed into a constraint graph G', denoted $G' = prep(G)$, such that*

1. *G' is d-regular with self-loops, and $\lambda(G') \leq \lambda < d$.*
2. *G' has the same alphabet as G, and $size(G') = O(size(G))$.*
3. *$\beta_1 \cdot \mathrm{UNSAT}(G) \leq \mathrm{UNSAT}(G') \leq \mathrm{UNSAT}(G)$.*

Proof. Apply Proposition 2.1 on G, then apply Proposition 2.2 on the result. The lemma is proven with $\beta_1 = c \cdot \frac{d}{d+d_0+1}$. □

3 Powering

In this section we construct out of a given constraint graph $G = \langle (V, E), \Sigma, \mathscr{C} \rangle$ a new graph $G^t = \left\langle (V, \mathbf{E}), \Sigma^{d^t}, \mathscr{C}^t \right\rangle$ for every $t \in \mathbb{N}$ as follows:

1. The vertices of G^t are the same as the vertices of G
2. u and v are connected by k edges in $\mathbf{E}$ iff the number of t-step paths from u to v in G is exactly k.
3. The alphabet of G^t is Σ^{d^t}, where every vertex specifies values for all its neighbours reachable in t steps.
4. The constraint associated with an edge $\mathbf{e} = (u, v) \in \mathbf{E}$ is satisfied iff the assignments for u and v are consistent with an assignment that satisfies all of the constraints of the path, induced by the t neighbourhoods of u and v. (In case this results in two different values for a vertex, then the constraint is automatically not satisfied.)

In the sequel $e \in E$ is called edge and $\mathbf{e} \in \mathbf{E}$ is called path. This construction yields the desired amplification of the satisfiability-gap. It is based on the fact, that the number of edges in G^t increases by d^{t-1}, but each edge in G is possibly included in td^{t-1} paths. Therefore, if a constraint corresponding to an edge e is not satisfied this yields to possibly td^{t-1} rejections of the corresponding assignment for G^t.

First we start with a technical proposition:

Proposition 3.1 *For every $p \in [0, 1]$ and $c > 0$ there exists some $0 < \tau \leq 1$ such that if $|l_1 - l_2| \leq (\sqrt{l_1} \wedge \sqrt{l_2})$, then*

$$\forall k, |k - pl_1| \vee |k - pl_2| \leq c(\sqrt{l_1} \wedge \sqrt{l_2}) : \tau \leq \frac{P(B_{l_1,p} = k)}{P(B_{l_2,p} = k)} \leq \frac{1}{\tau}$$

Proof. The proposition is perfectly symmetric in l_1 and l_2. That means without loss of generality we can assume $l_1 \geq l_2$. Write therefore $l_1 = l_2 + r$ for some $0 \leq r \leq \sqrt{l_1}$. We have

$$\begin{aligned} P(B_{l_1,p} = k) &= \binom{l_2 + r}{k} p^k (1-p)^{l_2+r-k} \\ &= \frac{l_2 + 1}{l_2 + 1 - k} \cdot \frac{l_2 + 2}{l_2 + 2 - k} \cdots \frac{l_2 + r}{l_2 + r - k} \binom{l_2}{k} \cdot p^k (1-p)^{l_2-k} (1-p)^r \\ &= \frac{l_2 + 1}{l_2 + 1 - k} \cdot \frac{l_2 + 2}{l_2 + 2 - k} \cdots \frac{l_2 + r}{l_2 + r - k} P(B_{l_2,p} = k) \end{aligned}$$

For all $a \leq r \leq \sqrt{l_1}$ we have with $l_2 - k \leq l_2 - p + c\sqrt{l_1} \leq (1-p)l_2 + c\sqrt{l_1}$:

$$\begin{aligned} \frac{l_2 + a}{l_2 + a - k} &\geq \frac{l_2}{(1-p)l_2 + (c+1)\sqrt{l_1}} \geq \frac{1}{1-p}\left(1 - \frac{c+1}{(1-p)\sqrt{l_1}}\right) \\ \frac{l_2 + a}{l_2 + a - k} &\leq \frac{l_2 + \sqrt{l_1}}{(1-p)l_2 - c\sqrt{l_1}} \leq \frac{1}{1-p}\left(1 + \frac{4c}{(1-p)\sqrt{l_1}}\right) \end{aligned}$$

This yields the proposition with $\tau = e^{-\frac{4c+1}{1-p}}$. □

Clearly we have that if $\mathrm{UNSAT}(G) = 0$, then $\mathrm{UNSAT}(G^t) = 0$. In the other case we have the amplification of the satisfiability gap:

Lemma 3.2 (Powering) *Let $\lambda < d$, and $|\Sigma|$ be arbitrary constants. There exists a constant $\beta_2 = \beta_2(\lambda, d, |\Sigma|) > 0$, such that for every $t \in \mathbb{N}$ and for every d-regular constraint graph $G = \langle (V, E), \Sigma, \mathscr{C} \rangle$ with self-loops and $\lambda(G) \leq \lambda$*

$$\mathrm{UNSAT}(G^t) \geq \beta_2 \sqrt{t} \cdot \min\left(UNSAT(G), \frac{1}{t}\right).$$

Proof. Let $\tilde{\sigma}$ be a best assignment for G^t, i.e. $\mathrm{UNSAT}_{\tilde{\sigma}}(G^t) = \mathrm{UNSAT}(G^t)$. By definition, $\tilde{\sigma}(v)$ specifies values for every w in its t-neighbourhood. Let $\tilde{\sigma}(v)_w$ denote this value. For every $1 \leq j \leq t$ and $v \in V$ let $X_{v,j}$ be a random variable with distribution

$$\forall a \in \Sigma : P(X_{v,j} = a) := \frac{\text{\# j-step paths starting from } v \text{ and ending at } w \text{ with } \tilde{\sigma}(w)_v = a}{\text{\# j-step paths starting from } v}$$

That means $X_{v,j}$ represents the value assigned to v by a random point within j-distance. With this define the corresponding assignment σ for G as follows:

$$\sigma(v) := \mathrm{maxarg}_{a \in \Sigma}(P(X_{v,\frac{t}{2}} = a))$$

Let $F \subset E$ be a subset of the edges that reject σ, such that $\mathrm{UNSAT}(G) = \frac{|F|}{|E|}$, and $I_r := \{\frac{t}{2} - r < i \leq \frac{t}{2} + r\}$. For each $\mathbf{e} \in \mathbf{E}$ define the following random variable

$$\begin{aligned} N_i(\mathbf{e}) &:= \begin{cases} 1 & \text{if } \mathbf{e} = (v_0, \ldots, v_t) \text{ where } (v_{i-1}, v_i) \in F \wedge \tilde{\sigma}(v_0)_{v_{i-1}} = \sigma(v_{i-1}) \wedge \tilde{\sigma}(v_t)_{v_i} = \sigma(v_i) \\ 0 & \text{otherwise} \end{cases} \\ N(\mathbf{e}) &:= \sum_{i \in I_r} N_i \end{aligned}$$

If $N(\mathbf{e}) > 0$, then this path has got a rejecting edge "in the middle", i.e. it clearly rejects $\tilde{\sigma}$. That means $P(N > 0) \leq \mathrm{UNSAT}_{\tilde{\sigma}}(G^t)$. In the sequel we will estimate this probability. To this end we make use of [Bul06, Lemma 23], i.e. $P(N > 0) \geq \frac{E^2(N)}{E(N^2)}$. Therefore we have to cope with $E(N)$ and $E(N^2)$.

1. $\mathbf{E}(\mathbf{N})$: From the definition we can gain the probability of $P(N_i > 0)$ as follows:

$$\begin{aligned} P(N_i > 0) &= P((u, v) \in F) \cdot P(X_{u,i-1} = \sigma(u)) P(X_{v,t-i} = \sigma(v)) \\ &= \frac{|F|}{|E|} \cdot P(X_{u,i-1} = \sigma(u)) P(X_{v,t-i} = \sigma(v)) \end{aligned}$$

   For any $l \in I_r$ we can decompose as follows. Distinguish between the loops in a path and the rest. That means a path is determined by the number of loops and their position and the rest of the path without loops. This considerations yield:

$$P(X_{u,l} = \sigma(u)) = \sum_{k=0}^{l} P(B_{l,p} = k) P(X'_{u,k} = \sigma(u)),$$

where $B_{l,p}$ is binomially distributed with $p = 1 - \frac{1}{d}$ and $X'_{u,k}$ is defined like $X_{u,k}$ but without loops. In the sequel we will use 3.1 to estimate that value. First set $r := -\frac{1}{2} + \frac{1}{2}\sqrt{1+2t}$ (then $|\frac{t}{2} - l| \leq \sqrt{\frac{t}{2}} \wedge \sqrt{l}$) and $M := \max\{t \in \mathbb{N}; \frac{t+r}{\frac{t}{2}-r}\}$. Now choose $c > 0$ such that the following is met:

$$\frac{p(1-p)}{c^2} M < \frac{1}{2|\Sigma|}.$$

Set $J := \{k \in \mathbb{N}; |k - pl| \vee |k - p\frac{t}{2}| \leq c(\sqrt{\frac{t}{2}} \wedge \sqrt{l})\}$. Then we have with Chebyshev

$$\begin{aligned} P(B_{\frac{t}{2},p} \notin J) &\leq P\Big(|B_{l,p} - pl| > c(\sqrt{\frac{t}{2}} \wedge \sqrt{l})\Big) + P\Big(|B_{\frac{t}{2},p} - p\frac{t}{2}| > c(\sqrt{\frac{t}{2}} \wedge \sqrt{l})\Big) \\ &\leq \frac{lp(1-p)}{c^2(\frac{t}{2} \wedge l)} + \frac{\frac{t}{2}p(1-p)}{c^2(\frac{t}{2} \wedge l)} \\ &\leq \frac{p(1-p)}{c^2} \frac{t+r}{\frac{t}{2}-r} \\ &\leq \frac{1}{2|\Sigma|} \end{aligned}$$

Remember, by construction of σ we have $P(X_{u,\frac{t}{2}} = \sigma(u)) \geq \frac{1}{|\Sigma|}$. We can now apply 3.1 with $l_1 = \frac{t}{2}, l_2 = i - 1$:

$$\begin{aligned} P(X_{u,i-1} = \sigma(u)) &\geq \sum_{k \in J} P(B_{i-1,p} = k) P(X'_{u,k} = \sigma(u)) \\ &\geq \tau \cdot \sum_{k \in J} P(B_{\frac{t}{2},p} = k) P(X'_{u,k} = \sigma(u)) \\ &\geq \tau(P(X_{u,\frac{t}{2}} = \sigma(u)) - \frac{1}{2|\Sigma|}) \\ &\geq \frac{\tau}{2|\Sigma|} \end{aligned}$$

The same holds for $P(X_{v,t-i} = \sigma(v))$. Therefore we finally have:

$$\begin{aligned} E(N) &= \sum_{i \in I_r} E(N_i) = \sum_{i \in I_r} P(N_i > 0) = \sum_{i \in I_r} \frac{|F|}{|E|} P(X_{u,i-1} = \sigma(u)) P(X_{v,t-i} = \sigma(v)) \\ &\geq |I_r| \frac{\tau^2}{4|\Sigma|^2} \frac{|F|}{|E|} \geq \Omega(\sqrt{t}) \cdot \frac{|F|}{|E|} \end{aligned}$$

2. $\mathbf{E}(\mathbf{N^2})$: To this end we define the following random variables:

$$\begin{aligned} Z_i(\mathbf{e}) &:= \begin{cases} 1 & \text{if } e_i \in F \\ 0 & \text{otherwise} \end{cases} \\ Z(\mathbf{e}) &:= \sum_{i \in I_r} Z_i(\mathbf{e}) \end{aligned}$$

Then we clearly have $N(\mathbf{e}) \leq Z(\mathbf{e})$ and we can start to calculate.

$$\begin{aligned} E(Z^2) &= \sum_{i\in I_r} E(Z_i^2) + 2\sum_{i<j;i,j\in I_r} E(Z_iZ_j) = |I_r|\frac{|F|}{|E|} + 2\sum_{i<j;i,j\in I_r} E(Z_iZ_j) \\ E(Z_iZ_j) &= P(Z_iZ_j > 0) = P(Z_i > 0)P(Z_j > 0|Z_i > 0) = \frac{|F|}{|E|}P(Z_j > 0|Z_i > 0) \\ P_{|\mathbf{e}|=t}(\ & Z_j\ (\mathbf{e}) > 0|Z_i(\mathbf{e}) > 0) = P_{|\mathbf{e}'|=t-i+1}(Z_{j-i+1}(\mathbf{e}') > 0|Z_1(\mathbf{e}') > 0) \leq \frac{|F|}{|E|} + \left(\frac{\lambda}{d}\right)^{j-i} \end{aligned}$$

The last inequality is the reason why the regularity condition on G had to be postulated. There we used the result given in [Bul06, Theorem 21], stating that $P(Z_s(\mathbf{e}) > 0|Z_1(\mathbf{e}) > 0) \leq \frac{|F|}{|E|} + \left(\frac{|\lambda(G)|}{d}\right)^s$, if G is d-regular. Having that $\lambda < d$, we know that $K := \sum_{i=0}^{\infty}\left(\frac{\lambda}{d}\right)^i < \infty$. Gathering all those results, we can estimate $E(N^2)$:

$$\begin{aligned} E(N^2) &\leq E(Z^2) \leq |I_r|\frac{|F|}{|E|} + 2\frac{|F|}{|E|}\sum_{i<j;i,j\in I_r}\left(\frac{|F|}{|E|} + \left(\frac{\lambda}{d}\right)^{j-i}\right) \\ &\leq O(\sqrt{t})\frac{|F|}{|E|} + 2\left(|I_r|^2\left(\frac{|F|}{|E|}\right)^2 + |I_r|\frac{|F|}{|E|}K\right) \\ &= O(\sqrt{t})\frac{|F|}{|E|} + O(t)\left(\frac{|F|}{|E|}\right)^2. \end{aligned}$$

Finally we can show the gap amplification. Recall that $\frac{|F|}{|E|} = \mathrm{UNSAT}(G)$ and distinguish between these two cases:

1. $\mathrm{UNSAT}(\mathbf{G}) \leq \frac{\mathbf{1}}{\sqrt{\mathbf{t}}}$: In this case we have

$$E(N^2) \leq O(\sqrt{t})\frac{|F|}{|E|} + O(t)\left(\frac{|F|}{|E|}\right)^2 \leq O(\sqrt{t})\frac{|F|}{|E|}.$$

This finally yields:

$$\mathrm{UNSAT}(G^t) \geq P(N > 0) \geq \frac{\Omega(\sqrt{t})^2\frac{|F|^2}{|E|^2}}{O(\sqrt{t})\frac{|F|}{|E|}} \geq \Omega(\sqrt{t})\frac{|F|}{|E|} \geq \Omega(\sqrt{t})\min(\mathrm{UNSAT}(G), \frac{1}{t})$$

2. $\mathrm{UNSAT}(\mathbf{G}) > \frac{\mathbf{1}}{\sqrt{\mathbf{t}}}$: We know that $E(N^2) \leq O(t)\frac{|F|}{|E|}$. This yields the desired result in this case, too:

$$\begin{aligned} \mathrm{UNSAT}(G^t) &\geq P(N > 0) \geq \frac{\Omega(\sqrt{t})^2\frac{|F|^2}{|E|^2}}{O(t)\frac{|F|}{|E|}} \geq \Omega(1)\frac{|F|}{|E|} \geq \Omega(1)\frac{1}{\sqrt{t}} \\ &= \Omega(\sqrt{t})\frac{1}{t} \geq \Omega(\sqrt{t})\min(\mathrm{UNSAT}(G), \frac{1}{t}) \end{aligned}$$

Therefore the proposition follows. □

4 Composition

The previous step amplified the satisfiability gap, but has blown up the alphabet. It remains to reduce the size of the alphabet. This will be done making use of *composition*. We will take a constraint of the given constraint graph. This will be put into an 'assignment tester' $\mathscr{P}$, which produces a constraint graph on an alphabet Σ_0 with $|\Sigma_0| = O(1)$. The set of all such graphs that were produced this way is put together into the resulting graph.
Let for the sequel $\mathrm{SAT}(\Phi) \subset \{0,1\}^n$ denote the set of assignments that satisfy a given boolean circuit Φ.

Definition 4.1 [Assignment Tester] An Assignment Tester with alphabet Σ_0 and rejection probability $\varepsilon > 0$ is a polynomial-time transformation $\mathscr{P}$ whose input is a circuit Φ over Boolean variables X, and whose output is a constraint graph $G = \langle (V, E), \Sigma_0, \mathscr{C} \rangle$ such that $X \subset V$, and such that the following holds. Let $V' = V \setminus X$, and let $a : X \to \{0,1\}$ be an assignment.

1. If $a \in \mathrm{SAT}(\Phi)$, there exists $b : V' \to \Sigma_0$ such that $\mathrm{UNSAT}_{a \cup b}(G) = 0$.
2. If $a \notin \mathrm{SAT}(\Phi)$ then for all $b : V' \to \Sigma_0$ holds $\mathrm{UNSAT}_{a \cup b}(G) \geq \varepsilon \cdot \mathrm{dist}(a, \mathrm{SAT}(\Phi))$.

Such an algorithm exists, see for example [Din05, 6.2]. Fixing such an assignment tester $\mathscr{P}$ we can now formulate and prove the following lemma:

Lemma 4.2 (Composition) *Let $\mathscr{P}$ an assignment tester with constant rejection probability $\varepsilon > 0$, and alphabet $\Sigma_0, |\Sigma_0| = O(1)$. There exists $\beta_3 > 0$ that depends only on $\mathscr{P}$, such that any constraint graph $G = \langle (V, E), \Sigma, \mathscr{C} \rangle$ can be transformed into a constraint graph $G' = \langle (V', E'), \Sigma_0, \mathscr{C}' \rangle$, denoted by $G \circ \mathscr{P}$, such that $size(G') = M(|\Sigma|) \cdot size(G)$, and*

$$\beta_3 \cdot \mathrm{UNSAT}(G) \leq \mathrm{UNSAT}(G') \leq \mathrm{UNSAT}(G)$$

Proof. Assignment tester are only defined for constraints over Boolean variables. Therefore we first prepare the graph with that respect. Let $e : \Sigma \to \{0,1\}^l$ be an encoding with relative Hamming distance $\varrho > 0, l = O(\log |\Sigma|)$. Replace each $v \in V$ by l Boolean variables $[v]$. Replace each constraint c over the two variables v, w by $\tilde{c}$ over $[v] \cup [w]$, such that $\tilde{c}$ is satisfied iff the assignment for $[v] \cup [w]$ is a legal encoding via e of an assignment for v and w that satisfies c. Now run $\mathscr{P}$ on any such $\tilde{c}$ and get constraint graphs $G_c = \langle (V_c, E_c), \Sigma_0, \mathscr{C}_c \rangle$. Without loss of generality we can assume $|E_c| = |E_{c'}|$ for any constraints c, c' (otherwise add edges with trivial constraints and set $\varepsilon' := \varepsilon \frac{\min_{c \in \mathscr{C}} |E_c|}{\max_{c \in \mathscr{C}} |E_c|}$). Now define the resulting graph G' as follows:

$$G' = \langle (V', E'), \Sigma_0, \mathscr{C}' \rangle ,$$

where $V' = \cup_{c \in \mathscr{C}} V_c, E' = \cup_{c \in \mathscr{C}} E_c, \mathscr{C}' = \cup_{c \in \mathscr{C}} \mathscr{C}_c$.
Let us now check, if G' has got the desired properties. We know that any c is transformed into a constraint $\tilde{c} : \{0,1\}^{2l} \to \{T, F\}$. There are only finitely many such constraints possible. Set M as the maximal size of the output graph of $\mathscr{P}$ for an input $\tilde{c}$. Then M only depends on $|\Sigma|$ and $\mathscr{P}$, and we have $size(G') \leq M \cdot size(G)$. Let $\sigma' : V' \to \Sigma_0$ be a best assignment for G'. Define $\sigma : V \to \Sigma$ by:

$$\sigma(v) := \mathrm{minarg}_{a \in \Sigma}(\mathrm{dist}(e(a), \sigma'([v]))).$$

That means $\sigma(v)$ is the value whose encoding via e is closest to $\sigma'([v])$. Let c be a constraint over the variables u, v that rejects σ. Now the property of the error correcting code comes in: At least a fraction of $\frac{\varrho}{2}$ of the bits of $\sigma'([u])$ or of $\sigma'([v])$ (or of both of them) has to be changed in order to lead to a satisfying value. That means $\mathrm{dist}(\sigma'|_{[u]\cup[v]}, \mathrm{SAT}(\tilde{c})) \geq \frac{\varrho}{4}$. Since $\mathscr{P}$ is an assignment tester with rejection probability of ε, we have that at least a fraction of $\varepsilon \cdot \frac{\varrho}{4}$ of the constraints in $\mathscr{C}_c$ reject σ'. The assumption $|E_c| = |E_{c'}|$ for any constraints c, c' and $|E'| = \sum_{c\in\mathscr{C}} |E_c|$ therefore guarantees:$\mathrm{UNSAT}(G') \geq \frac{\varepsilon\varrho}{4} \cdot \mathrm{UNSAT}(G)$. □

5 Main theorem

Now we have everything at hand to finally state our main theorem.

Theorem 5.1 (Main) *For any $\Sigma, |\Sigma| = O(1)$, there exist constants $C > 0$ and $0 < \alpha < 1$, such that given a constraint graph $G = \langle (V, E), \Sigma, \mathscr{C}\rangle$ one can construct, in polynomial time, a constraint graph $G' = \langle (V', E'), \Sigma_0, \mathscr{C}'\rangle$ such that*

1. *$size(G') \leq C \cdot size(G)$ and $|\Sigma_0| = O(1)$.*
2. *If $\mathrm{UNSAT}(G) = 0$ then $\mathrm{UNSAT}(G') = 0$*
3. *$\mathrm{UNSAT}(G') \geq \min(2 \cdot \mathrm{UNSAT}(G), \alpha)$.*

Proof. We construct G' using the lemmas of the previous sections:

$$G' = (prep(G))^t \circ \mathscr{P}$$

First we notice that each lemma only incurs a linear blowup of the size. More precisely, the number of edges increases by a constant factor during Preprocessing and Powering. In the Composition step the size grows by a factor that depends only on $|\Sigma^{d^t}|$ and on $\mathscr{P}$ which do not depend on G. Let $\beta_1, \beta_2, \beta_3$ be the constants in the Lemmas 2.3,3.2,4.2. Now choose $t = \left\lceil \left(\frac{2}{\beta_1\beta_2\beta_3}\right)^2 \right\rceil$ and $\alpha = \frac{\beta_3\beta_2}{\sqrt{t}}$. Then we have altogether:

$$\begin{aligned}
\mathrm{UNSAT}(G') &\overset{4.2}{\geq} \beta_3 \cdot \mathrm{UNSAT}((prep(G))^t) \\
&\overset{3.2}{\geq} \beta_3 \cdot \beta_2\sqrt{t} \cdot \min(\mathrm{UNSAT}(prep(G)), \frac{1}{t}) \\
&\overset{2.3}{\geq} \beta_3 \cdot \beta_2\sqrt{t} \cdot \min(\beta_1\mathrm{UNSAT}(G), \frac{1}{t}) \\
&\geq \min(2 \cdot \mathrm{UNSAT}(G), \alpha).
\end{aligned}$$

This proves the theorem. □

At last we can now prove the PCP Theorem. [Bul06, Lemma 9] has already shown that the PCP Theorem is equivalent to showing that Gap-3SAT is NP-hard.

Corollary 5.2 *Gap-3SAT is NP-hard.*

Proof. According to [Bul06, Theorem 14], it is NP-hard to decide if $\mathrm{UNSAT}(G) = 0$ or not for a given constraint graph G with $|\Sigma| = 7$. Let G_0 be such an instance and G_i be the outcome of applying the main theorem on G_{i-1}. Set $k := \lceil \log(\alpha |E_0|) \rceil = O(\log n)$. This way we have for all $i \leq k : size(G_i) \leq C^i \cdot size(G_0) = \mathrm{poly}(n)$. If $\mathrm{UNSAT}(G_0) = 0$ then $\mathrm{UNSAT}(G_i) = 0$. If not, then we have by induction:

$$\mathrm{UNSAT}(G_i) \geq \min(2^i \mathrm{UNSAT}(G_0), \alpha).$$

If $\mathrm{UNSAT}(G_0) > 0$, then $\mathrm{UNSAT}(G_0) \geq \frac{1}{|E_0|}$, so surely we have $2^k \mathrm{UNSAT}(G_0) > \alpha$ and therefore $\mathrm{UNSAT}(G_k) \geq \alpha$. Now a local gadget reduction takes G_k to 3SAT form, while maintaining the satisfiability gap up to some constant. This proves the corollary. □

References

[AS98] S. Arora and S. Safra. *Probabilistic checking of proofs: A new characterisation of NP*. J.ACM, 45(1):70-122, 1998.

[Bul06] Lukas Bulwahn. *The PCP Theorem: An Introduction*. Munich, 2006.

[Din05] Irit Dinur. *The PCP Theorem via Gap Amplification*. Technical Report TR05-46, Electronic Colloquium on Computational Complexity, 2005.

[FGL+96] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy. *Approximating clique is almost NP-complete.* Journal of the ACM, 43(2):268-292, 1996.