Course "Proofs and Computers", JASS'06

PCP-Theorem by Gap Amplification

Bernhard Vesenmayer

Faculty of Mathematics
TU Munich

March 29, 2006

Preprocessing

In this step we want to transform a given constraint graph into a 'nice' one. That means for the resulting graph G' should hold:

1. G' is of constant degree d, i.e. independent of the input graph G,
2. G' is d-regular with self-loops, and $\lambda(G') \leq \lambda < d$,
3. $\beta_1 \cdot \mathrm{UNSAT}(G) \leq \mathrm{UNSAT}(G') \leq \mathrm{UNSAT}(G)$.

Constant degree

Proposition 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*

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

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

Construction of G'

1. Let X_n be a d_0-regular expander on n vertices with $h(X_0) > h_0$.

Construction of G'

1. Let X_n be a d_0-regular expander on n vertices with $h(X_0) > h_0$.
2. Let d_v be the degree of v. Replace v by X_{d_v}, equipped with equality constraints, denoted by $[v]$. Let E_1 be the set of such edges.

Construction of G'

1. Let X_n be a d_0-regular expander on n vertices with $h(X_0) > h_0$.
2. Let d_v be the degree of v. Replace v by X_{d_v}, equipped with equality constraints, denoted by $[v]$. Let E_1 be the set of such edges.
3. For every $(u, v) \in E$ put an edge between one vertice of $[u]$ and $[v]$, such that each vertice only sees one such 'external' edge. Let E_2 be the set of those edges.

Construction of G'

1. Let X_n be a d_0-regular expander on n vertices with $h(X_0) > h_0$.
2. Let d_v be the degree of v. Replace v by X_{d_v}, equipped with equality constraints, denoted by $[v]$. Let E_1 be the set of such edges.
3. For every $(u, v) \in E$ put an edge between one vertice of $[u]$ and $[v]$, such that each vertice only sees one such 'external' edge. Let E_2 be the set of those edges.
4. Define $G' := ([V], \mathbf{E} = E_1 \cup E_2)$, where $[V] := \cup_{v \in V}[v]$.

Properties: Upper bound

1. Let $\sigma : V \to \Sigma$ be a best assignment of G.
2. Define $\sigma' : [V] \to \Sigma$ by: $\forall v \in V, x \in [v] : \sigma'(x) = \sigma(v)$

$$\Rightarrow \mathrm{UNSAT}(G) = \mathrm{UNSAT}_{\sigma}(G) \geq \mathrm{UNSAT}_{\sigma'}(G') \geq \mathrm{UNSAT}(G')$$

Properties: Lower bound

1. Let $\sigma' : [V] \to \Sigma$ be a best assignment.
2. Define $\sigma : V \to \Sigma$ by: $\sigma(v) := \mathrm{maxarg}_{a \in \Sigma}(P_{x \in [v]}(\sigma'(x) = a))$
3. Let $F \subset E, \mathbf{F} \subset \mathbf{E}$ be the set of edges that reject σ, σ'.
4. Let $S := \cup_{v \in V}\{x \in [v]; \sigma'(x) \neq \sigma(v)\}$.

$$\Rightarrow |\mathbf{F}| + |S| \geq |F| = \alpha \cdot |E|, \text{ where } \alpha := \frac{|F|}{|E|}.$$

Properties: Lower bound - Two cases

We have to distinguish between two cases:

1. $|\mathbf{F}| \geq \frac{\alpha}{2}|E|$
2. $|\mathbf{F}| < \frac{\alpha}{2}|E|$

The first case is simple:

$$|\mathbf{F}| \geq \frac{\alpha}{2}|E| = \frac{\alpha}{2d}|\mathbf{E}|$$

$$\begin{aligned} \Rightarrow \mathrm{UNSAT}(G') &= \mathrm{UNSAT}_{\sigma'}(G') \\ &\geq \frac{1}{2d}\mathrm{UNSAT}_{\sigma}(G) \\ &\geq \frac{1}{2d}\mathrm{UNSAT}(G). \end{aligned}$$

Properties: Lower bound - Two cases

Second case: $|\mathbf{F}| < \frac{\alpha}{2}|E| \Rightarrow |S| \geq \frac{\alpha}{2}|E|$. Define the following:

1. $S^v := [v] \cap S$
2. $S_a^v := \{x \in S^v; \sigma'(x) = a\}$
3. $|S_a^v| \leq \frac{|[v]|}{2} \Rightarrow |E(S_a^v, [v] \setminus S_a^v)| \geq h_0 \cdot |S_a^v|$
4. $|\mathbf{F}| \geq \sum_{v \in V} \sum_{a \in \Sigma} h_0 |S_a^v| = h_0|S| \geq h_0 \frac{\alpha}{2}|E| = \frac{h_0 \alpha}{2d}|\mathbf{E}|$

$$\begin{aligned} \Rightarrow \mathrm{UNSAT}(G') &= \mathrm{UNSAT}_{\sigma'}(G') \\ &\geq \frac{h_0}{2d}\mathrm{UNSAT}_{\sigma}(G) \\ &\geq \frac{h_0}{2d}\mathrm{UNSAT}(G). \end{aligned}$$

Expanderizing

Proposition 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. $\frac{d}{d+d_0+1} \cdot \mathrm{UNSAT}(G) \leq \mathrm{UNSAT}(G') \leq \mathrm{UNSAT}(G)$.

Construction of G'

1. Let $E_{loop} := \{(v, v); v \in V\}$.

Construction of G'

1. Let $E_{loop} := \{(v, v); v \in V\}$.
2. Let $X = (V, E')$ be a d_0-regular expander on $|V|$ vertices with $h(X) \geq h_0$.

Construction of G'

1. Let $E_{loop} := \{(v, v); v \in V\}$.
2. Let $X = (V, E')$ be a d_0-regular expander on $|V|$ vertices with $h(X) \geq h_0$.
3. Define $G' = (V, E \cup E' \cup E_{loop})$ with trivial constraints on the new edges.

Properties

1. G' is $(d + d_0 + 1)$-regular and of the same size as G.
2. $\lambda(G') \leq d + d_0 + 1 - \frac{h(G')^2}{d(G')} \leq d + d_0 + 1 - \frac{h_0^2}{d+d_0+1} < d + d_0 + 1$
3. The fraction of unsatisfied edges drops at most by $\frac{d}{d+d_0+1}$. □

Preprocessing

Lemma 1 (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. Consecutively apply the previous two propositions. □

Powering

In this step we amplify the satisfiability gap. Let G be a given constraint graph. Consider the graph G^t which is defined as follows:

1. The vertices of G^t are the same as the vertices of G

Powering

In this step we amplify the satisfiability gap. Let G be a given constraint graph. Consider the graph G^t which is defined 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 **E** iff the number of t-step paths from u to v in G is exactly k.

Powering

In this step we amplify the satisfiability gap. Let G be a given constraint graph. Consider the graph G^t which is defined 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 **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.

Powering

In this step we amplify the satisfiability gap. Let G be a given constraint graph. Consider the graph G^t which is defined 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 induced by the t neighbourhoods of u and v.

Powering

Lemma 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 $\lambda(G) \leq \lambda$

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

Assignments

1. Let $\tilde{\sigma}$ be a best assignment for G^t, i.e. $\mathrm{UNSAT}_{\tilde{\sigma}}(G^t) = \mathrm{UNSAT}(G^t)$.
2. Let $\tilde{\sigma}(v)_w$ denote the value that is specified for w by v.
3. For every $1 \leq j \leq t$ and $v \in V$ let $X_{v,j}$ be a random variable with distribution

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

4. Now define $\sigma : V \to \Sigma$ as follows:

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

$N(\mathbf{e})$

1. Let $F \subset E$ be a subset of the edges that reject σ, such that $\mathrm{UNSAT}(G) = \frac{|F|}{|E|}$.
2. Let $I_r := \{\frac{t}{2} - r < i \leq \frac{t}{2} + r\}$.
3. For each $\mathbf{e} = (v_0, \ldots, v_t) \in \mathbf{E}$ define the following random variable

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

Then $\mathrm{UNSAT}_{\tilde{\sigma}}(G^t) \geq P(N > 0) \geq \frac{E^2(N)}{E(N^2)}$.

$E(N)$

1. By definition we have:
   $E(N) = \sum_{i \in I_r} E(N_i)) = \sum_{i \in I_r} P(N_i > 0)$.
2. We can gain the probability of $P(N_i > 0)$ as follows:

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

3. For any $l \in I_r$ we can decompose as follows ($B_{l,p}$ is binomially distributed with $p = 1 - \frac{1}{d}$ and $X'_{u,k}$ is defined like $X_{u,k}$ but without loops):

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

$E(N)$

Now we need the following technical proposition, stating the distribution of B_{l_1} and B_{l_2} are quite similar:

Proposition 3

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

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

1. Define $J := \{k \in \mathbb{N}; |k - pl| \vee |k - p\frac{t}{2}| \leq c(\sqrt{\frac{t}{2}} \wedge \sqrt{l})\}$.
2. Choose c such that $P(B_{\frac{t}{2},p} \notin J) \leq \frac{1}{2|\Sigma|}$.

$E(N)$

1. By construction of σ we have $P(X_{u,\frac{t}{2}} = \sigma(u)) \geq \frac{1}{|\Sigma|}$.
2. Apply the proposition with $r := -\frac{3}{2} + \frac{1}{2}\sqrt{1+2t}$ (then the condition $|l_1 - l_2| \leq \sqrt{l_1} \wedge \sqrt{l_2}$ is met) and with $l_1 = \frac{t}{2}, l_2 = i-1$:

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

$E(N)$

Therefore we finally have:

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

$E(N^2)$

Define $Z_i(\mathbf{e}) = 1$ iff $e_i \in F$, else $Z_i(\mathbf{e}) = 0$, $Z(\mathbf{e}) := \sum_{i \in I_r} Z_i(\mathbf{e})$.

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

$E(N^2)$

Having that $\lambda < d$, we know that $\sum_{i=0}^{\infty}\left(\frac{\lambda}{d}\right)^i = K < \infty$. Gathering all those results, we can estimate $E(N^2)$:

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

Powering: Two Cases

Now we can show the amplification. However we have to distinguish between two cases:

1. $\mathrm{UNSAT}(\mathbf{G}) > \frac{\mathbf{1}}{\sqrt{\mathbf{t}}}$
2. $\mathrm{UNSAT}(\mathbf{G}) \leq \frac{\mathbf{1}}{\sqrt{\mathbf{t}}}$

The first case is simple:

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

Powering: Two Cases

Let $\mathrm{UNSAT}(\mathbf{G}) \leq \frac{\mathbf{1}}{\sqrt{\mathbf{t}}}$, then we have:

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

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

Composition

Definition 4

[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 $\mathrm{UNSAT}_{a \cup b}(G) \geq \varepsilon \cdot \mathrm{dist}(a, \mathrm{SAT}(\Phi))$.

Composition

We this we can state the composition lemma:

Lemma 3 (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

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

Composition: Preparation

The assignment tester needs input constraints defined over Boolean variables.

1. Let $e : \Sigma \to \{0,1\}^l$ be an encoding with relative Hamming distance $\varrho > 0$, $l = O(\log |\Sigma|)$.
2. Replace each $v \in V$ by l Boolean variables $[v]$.
3. 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.

Composition: Composition

1. Run $\mathscr{P}$ on any such $\tilde{c}$ and get constraint graphs $G_c = \langle (V_c, E_c), \Sigma_0, \mathscr{C}_c \rangle$.
2. Without loss of generality we can assume $|E_c| = |E_{c'}|$ for any constraints c, c'.
3. Define the resulting graph G' as follows:

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

where $V' = \cup_{c \in \mathscr{C}} V_c$, $E' = \cup_{c \in \mathscr{C}} E_c$, $\mathscr{C}' = \cup_{c \in \mathscr{C}} \mathscr{C}_c$.

Composition: Properties

1. Any c is transformed into a constraint $\tilde{c} : \{0,1\}^{2l} \to \{T, F\}$.
2. Set M as the maximal size of the output graph of $\mathscr{P}$ for an input $\tilde{c}$.
   $\Rightarrow M$ only depends on $|\Sigma|$ and $\mathscr{P}$, and we have $size(G') \leq M \cdot size(G)$.
3. Let $\sigma' : V' \to \Sigma_0$ be a best assignment for G'.
4. Define $\sigma : V \to \Sigma$ by:$\sigma(v) := \mathrm{minarg}_{a \in \Sigma}(\mathrm{dist}(e(a), \sigma'([v])))$.
5. Let c be a constraint over the variables u, v that rejects σ.
6. That means $\mathrm{dist}(\sigma'|_{[u] \cup [v]}, \mathrm{SAT}(\tilde{c})) \geq \frac{\varrho}{4}$.
7. By definition at least a fraction of $\varepsilon \cdot \frac{\varrho}{4}$ of the constraints in $\mathscr{C}_c$ reject σ'.

$$\Rightarrow \mathrm{UNSAT}(G') \geq \frac{\varepsilon \varrho}{4} \cdot \mathrm{UNSAT}(G).$$

$\square$

Main Theorem

Finally we can state the main amplification theorem:

Theorem 4 (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 of the Main Theorem

1. Define $G' = (prep(G))^t \circ \mathscr{P}$ using the previous lemmata.
2. Each lemma only incurs a linear blowup of the size.
3. Choose $t = \left\lceil \left(\frac{2}{\beta_1\beta_2\beta_3}\right)^2 \right\rceil$ and $\alpha = \frac{\beta_3\beta_2}{\sqrt{t}}$:

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

Proof of the Main Theorem

1. It is NP-hard to decide if $\mathrm{UNSAT}(G) = 0$ or not for a given constraint graph G with $|\Sigma| = 7$.

Proof of the Main Theorem

1. It is NP-hard to decide if $\mathrm{UNSAT}(G) = 0$ or not for a given constraint graph G with $|\Sigma| = 7$.
2. Let G_0 be such an instance and G_i be the the outcome after applying the main theorem on G_{i-1}.

Proof of the Main Theorem

1. It is NP-hard to decide if $\mathrm{UNSAT}(G) = 0$ or not for a given constraint graph G with $|\Sigma| = 7$.
2. Let G_0 be such an instance and G_i be the the outcome after applying the main theorem on G_{i-1}.
3. Set $k := \lceil \log(\alpha |E_0|) \rceil = O(\log n) :\Rightarrow \forall i \leq k : size(G_i) \leq C^i \cdot size(G_0) = \mathrm{poly}(n)$.

Proof of the Main Theorem

1. It is NP-hard to decide if $\mathrm{UNSAT}(G) = 0$ or not for a given constraint graph G with $|\Sigma| = 7$.
2. Let G_0 be such an instance and G_i be the the outcome after applying the main theorem on G_{i-1}.
3. Set $k := \lceil \log(\alpha |E_0|) \rceil = O(\log n) :\Rightarrow \forall i \leq k : size(G_i) \leq C^i \cdot size(G_0) = \mathrm{poly}(n)$.
4. If $\mathrm{UNSAT}(G_0) = 0$ then $\mathrm{UNSAT}(G_i) = 0$.

Proof of the Main Theorem

1. It is NP-hard to decide if $\mathrm{UNSAT}(G) = 0$ or not for a given constraint graph G with $|\Sigma| = 7$.
2. Let G_0 be such an instance and G_i be the the outcome after applying the main theorem on G_{i-1}.
3. Set $k := \lceil \log(\alpha |E_0|) \rceil = O(\log n) :\Rightarrow \forall i \leq k : size(G_i) \leq C^i \cdot size(G_0) = \mathrm{poly}(n)$.
4. If $\mathrm{UNSAT}(G_0) = 0$ then $\mathrm{UNSAT}(G_i) = 0$.
5. If not, then we have by induction:$\mathrm{UNSAT}(G_i) \geq \min(2^i \mathrm{UNSAT}(G_0), \alpha)$.

Proof of the Main Theorem

1. It is NP-hard to decide if $\mathrm{UNSAT}(G) = 0$ or not for a given constraint graph G with $|\Sigma| = 7$.
2. Let G_0 be such an instance and G_i be the the outcome after applying the main theorem on G_{i-1}.
3. Set $k := \lceil \log(\alpha|E_0|) \rceil = O(\log n) :\Rightarrow \forall i \leq k : size(G_i) \leq C^i \cdot size(G_0) = \mathrm{poly}(n)$.
4. If $\mathrm{UNSAT}(G_0) = 0$ then $\mathrm{UNSAT}(G_i) = 0$.
5. If not, then we have by induction:$\mathrm{UNSAT}(G_i) \geq \min(2^i \mathrm{UNSAT}(G_0), \alpha)$.
6. Therefore $2^k \mathrm{UNSAT}(G_0) \geq \alpha$ and therefore $\mathrm{UNSAT}(G_k) \geq \alpha$.

Proof of the Main Theorem

1. It is NP-hard to decide if $\mathrm{UNSAT}(G) = 0$ or not for a given constraint graph G with $|\Sigma| = 7$.
2. Let G_0 be such an instance and G_i be the the outcome after applying the main theorem on G_{i-1}.
3. Set $k := \lceil \log(\alpha|E_0|) \rceil = O(\log n) :\Rightarrow \forall i \leq k : size(G_i) \leq C^i \cdot size(G_0) = \mathrm{poly}(n)$.
4. If $\mathrm{UNSAT}(G_0) = 0$ then $\mathrm{UNSAT}(G_i) = 0$.
5. If not, then we have by induction:$\mathrm{UNSAT}(G_i) \geq \min(2^i \mathrm{UNSAT}(G_0), \alpha)$.
6. Therefore $2^k \mathrm{UNSAT}(G_0) \geq \alpha$ and therefore $\mathrm{UNSAT}(G_k) \geq \alpha$.
7. Now a local gadget reduction takes G_k to 3SAT form, while maintaining the satisfiability gap up to some constant.