# On Radhakrishnan's Proof of PCP Gap Amplification

Petar Maymounkov[*]

**Abstract**

This exposition gives a simplified version of Radhakrishnan's construction of a PCP gap amplification transform [4], which is itself a simplified version of Dinur's construction [3]. The preprocessing and alphabet reduction transforms are left unchanged, and are thus not addressed here. Briefly, we simplify the analysis by pretending that all vertices have opinions about the color of *all* other vertices. This helps avoid conditioning in the analysis of the non-truncated random walk, and makes our arguments conceptually crisper. The "truncation argument" is reduced to a couple of lines as opposed to over two pages in [4]. We also include a proof, due to [6], to a somewhat counter-intuitive random walk lemma used in [4]. Our exposition contains various fixes of small "bugs" as well as restated, clarified and detailed versions of the more important arguments in [4].

# 1 Gap amplification

## 1.1 Notation

We use $\mathbf{I}[\cdot]$ to denote the indicator function.

## 1.2 Setting

In this essay we give a simplified exposition of Radhakrishnan's PCP gap amplification transform [4], which is itself a simplified, slightly modified and conceptually cleaner version of Dinur's original amplification transform [3]. Recall that out starting point is an instance of the Generalized Graph Coloring (GGC) problem. In particular, we are given an undirected graph $G = (V, E)$, together with an alphabet $\Sigma$ of "allowed colors." Each edge $e = (u, v)$ of $G$ is associated with a *constraint* $c_e : \Sigma \times \Sigma \to \{0, 1\}$ which determines whether any particular coloring of $u$ and $v$ is *valid* or not. To be more precise, we define that the first argument to $c_e$ to be the color of the vertex that comes first in lexicographic order. We will not specify the representation of $c_e$ here, because it is not essential for the purposes of this proof. We

[*] Massachusetts Institute of Technology, petar@csail.mit.edu

briefly mention that [3] represents the constraints via boolean formulas, while [4] represents them via systems of quadratic polynomials.

We further restrict our attention to GGC instances $G$ with the following properties:

1. Each vertex of $G$ has $d/2$ self-loops

2. $G$ is an $(\eta, d)$-expander

We abuse notation a little by using $G$ to refer both to the underlying graph as well as to the problem instance. Let $\omega : V \to \Sigma$ be a coloring. Define:

$$\text{UNSAT}_\omega\,(G) = \mathbf{Pr}_{e=(u,v)\in E}\big[c_e(\omega(u),\omega(v)) = 0\big] \quad \text{and} \quad \text{UNSAT}\,(G) = \min_\omega\ \text{UNSAT}_\omega\,(G)$$

When $G$ is satisfiable $\text{UNSAT}\,(G) = 0$, otherwise $\text{UNSAT}\,(G) \geq 1/|E|$. The goal of this section is to construct a GGC instance $G^t$, for any $t \geq 1$, with the following properties:

1. $G^t$ is only linearly bigger than $G$. More precisely, $G^t$ is defined on the same vertex set $V$, the number of edges in $G^t$ is at most $d^t$ times the number of edges in $G$, and each original constraint $c_e$ is repeated (typically in conjunction with others) at most $d^t$ times in $G^t$.

2. If $\text{UNSAT}\,(G) = 0$, then $\text{UNSAT}\,(G^t) = 0$, otherwise $\text{UNSAT}\,(G^t) \geq \epsilon$ for some universal constant $\epsilon$ related to $t$.

## 1.3 Construction

Before we proceed to the definition of the new GGC instance $G^t$, we would like to highlight the "isomorphism" between a GGC instance and the corresponding PCP. For each GGC instance $G$, the natural corresponding PCP is one where:

1. The proof $\Pi : V(G) \to \Sigma$ corresponds to a coloring of $G$.

2. The verifier $V$ picks one random edge $e = (u,v) \in E(G)$, queries the proof for the colors at vertices $u$ and $v$, say $\omega(u)$ and $\omega(v)$ respectively, and verifies that the colors satisfy the constraints at $e$, i.e. $c_e(\omega(u),\omega(v)) = 1$.

We now describe the construction of the GGC instance $G^t$:

1. $G^t$ is defined over the same vertex set $V$ as $G$ itself

2. The alphabet of $G^t$ is $\Sigma^{d^{t+1}}$. The intention is that if $G$ is satisfiable, the only correct coloring of $G^t$ is the one where each vertex $v \in G^t$ "knows" the satisfying coloring of all vertices in $G$ that are no more than $t$ steps away from $v$. This we call the *intended coloring* of $G^t$. In this setting, for an arbitrary coloring $\omega^t : V \to \Sigma^{d^{t+1}}$ of $G^t$ we let $\omega^t_a(u)$ be the "opinion" of $a \in G^t$ about the color of $u \in G$ for a purported satisfying assignment of $G$. Whenever $d_G(a,u) > t$, we define $\omega^t_a(u) = \star$ alluding that $v$ has no opinion about $u$'s color.

3. To define the edges of $G^t$ consider the following random walk process on $G$ that outputs two vertices $a$ and $b$:

   (a) Pick a $u_0 \in V$ uniformly at random and set $a := u_0$. Repeat the following step until a stopping condition is reached:

   (b) Having chosen $u_0, u_1, \ldots, u_{i-1}$ let $e_i$ be a random edge leaving $u_{i-1}$ and arriving at $u_i$. Add $u_i$ to the walk and stop this process with probability $1/t$, in which case set $b := u_i$, otherwise repeat.

   Now consider the outcomes $(a, b)$ of the above process. We would ideally like each such outcome to correspond to an edge $(a, b)$ in $G^t$. However, we need to take extra care in ensuring that the walk corresponding to each edge in $G^t$ occurs with the same probability across all edges and furthermore we would like to "discard" all walks of more than $t$ steps (because otherwise we would have to have infinitely many edges).

   This is achieved using the following technical trick. We think of the choice of a random walk as a uniform sample from $W = [n] \times [d]^t \times [t]^t$, which encodes the choice of initial vertex, out-edges at every step as well as stopping conditions (should they occur within the first $t$ steps). A walk ends within $t$ steps if we encounter "1" in the $[t]^t$ dimension of $W$, otherwise we say that the walk is *truncated* (or discarded). The edges of $G^t$ are now in 1-to-1 correspondence with the elements of $W$ in the following way. If $w \in W$ represents a walk of at most $t$ steps, than the corresponding edge in $G^t$ simply connects the starting and ending vertices of the walk. Otherwise, the walk corresponds to a self-loop at the starting vertex with an always-accepting constraint. (Note that the same random walk may be represented by more than one element of $W$.)

4. An edge $(a, b) \in G^t$ is associated with the following constraints. Let $(e_1, \ldots, e_T)$ be the edges in the walk from $a$ to $b$ in $G$. We say that $\omega_a^t$ and $\omega_b^t$ *pass the test* at $e_i = (u_{i-1}, u_i)$ if at least one of the following conditions holds:

   (a) $\omega_a^t(u_{i-1}) = \star$ or $\omega_b^t(u_i) = \star$

   (b) $\omega_a^t(u_{i-1}) = \omega_b^t(u_{i-1})$ and $\omega_a^t(u_i) = \omega_b^t(u_i)$ and $c_{(u_{i-1},u_i)}(\omega_a^t(u_{i-1}), \omega_a^t(u_i)) = 1$

   The constraint at $e = (a, b) \in G^t$, denoted $c^t_{(a,b)}(w_a^*, w_b^*)$, simply ensures that $w_a^*$ and $w_b^*$ pass the test at all $e_1, \ldots, e_T$.

This completes the construction of $G^t$. Claims regarding the size of $G^t$ are straightforward to verify.

## 1.4 Analysis

The main technical result is now the following "amplification" lemma:

**Lemma 1.1** (Lemma 6.1 in [3], Lemma 5.7 in [4]). *Given a GGC instance $G$ such that $G$ is an $(\eta, d)$-expander and each vertex of $G$ has at least $d/2$ self-loops, the following hold true for all $t \geq 1$:*

1. *If $G$ is satisfiable, then so is $G^t$ under the intended coloring.*

2. *If $G$ is unsatisfiable, then:*

$$\text{UNSAT}\ (G^t) \geq \frac{3-\sqrt{5}}{8} \cdot \frac{1}{1+d^2/\eta^2} \cdot t \cdot \min(\text{UNSAT}\ (G), 1/2t)$$

The first part is straightforward.

For the second part, let $\epsilon =$UNSAT $(G)$. Fix a coloring $\omega^t$ of $G^t$. We would like to argue that $q = \mathbf{Pr}_{(a,b)\in G^t}[c_e^t(\omega_a^t, \omega_b^t) = 0]$ is sufficiently large (as desired by the lemma). Deriving a lower bound on $q$ is equivalent (by construction) to deriving a lower bound on the event that the corresponding random walk fails a test or is longer than $t$ steps. The latter occurs with probability $(1-1/t)^t \asymp 1/e$. We will however analyze the random walk as if it were not truncated, and at the end we will argue that truncation does not change the outcome by a lot.

Since we will be analyzing the "ideal" not-truncated walk, we will have to pretend that each vertex of $G^t$ has an opinion about the colors of *all* vertices of $G$. And thus this imaginary coloring of $G^t$ will be of the form $\omega^\infty : V \to \Sigma^{|V|}$. We will assume that $\omega^\infty$ is any such coloring which agrees with $\omega^t$ in the natural way (i.e. in terms of opinions about colors of vertices). At the end of our probabilistic analysis we will apply a "truncation argument" which will ensure that we do not consider any outcomes that might have been produced by looking at the parts of $\omega^\infty$ not present in $\omega^t$.

From here onwards we shall use $\kappa$ to denote a random walk over vertices $(w_0, \ldots, w_{\ell(\kappa)})$ and edges $(e_1, \ldots, e_{\ell(\kappa)})$ as defined above, where $\ell(\kappa)$ is the length of the walk. Let $W^{\#}_{\kappa,u\to v}$ be the number of $i$'s for which $w_i = u$ and $w_{i+1} = v$. We use $\kappa_a = w_0$ and $\kappa_b = w_{\ell(\kappa)}$. We distinguish two flavors of random walks: an after-walk and a before-walk. The former denotes a walk as defined above. The latter differs from it in that stopping decisions are made *before* edge traversals. We use $\kappa^*$ to denote a before-walk. Central to our analysis is the next lemma:

**Lemma 1.2** (Random Walk Lemma)**.** *Let $G$ be an undirected $d$-regular multi-graph with at least one self-loop at each vertex. Let $\kappa$ be an after-walk on $G$, and let $Q_k = W^{\#}_{\kappa,u\to v}$. Then for all $a, b, u, v \in V$ and $k \in \mathbb{Z}^+$:*

1. *$\kappa_a$ and $\kappa_b$ are independent conditioned on $Q = k$*

2. *$\mathbf{Pr}_\kappa[\kappa_b = b \mid Q = k] = \mathbf{Pr}_{\kappa^*}[\kappa_b^* = b \mid \kappa_a^* = v]$*

3. *$\mathbf{Pr}_\kappa[\kappa_a = a \mid Q = k] = \mathbf{Pr}_{\kappa^*}[\kappa_a^* = u \mid \kappa_b^* = a]$*

Since the statement of this lemma (in particular part 2) is somewhat counter-intuitive, we give a detailed proof here:

*Proof of Lemma 1.2.* This proof is borrowed from [6]. Part 2. It is straightforward that:

$$p_b = \mathbf{Pr}_\kappa[\kappa_b = b \mid Q \geq k] = \mathbf{Pr}_{\kappa^*}[\kappa_b^* = b \mid \kappa_a^* = v]$$

This is so since once $(u, v)$ is traversed $k$ times there are no additional restrictions on the walk. Thus $p_b$ is independent of $k$.

$$\begin{aligned} p_b &= \mathbf{Pr}_\kappa[\kappa_b = b \mid Q \geq 1] \\ &= \frac{\mathbf{Pr}_\kappa[\kappa_b = b \wedge Q \geq 1]}{\mathbf{Pr}_\kappa[Q \geq 1]} \\ &= \frac{\mathbf{Pr}_\kappa[\kappa_b = b \wedge Q = 1] + \mathbf{Pr}_\kappa[\kappa_b = b \wedge Q \geq 2]}{\mathbf{Pr}_\kappa[Q = 1] + \mathbf{Pr}_\kappa[Q \geq 2]} \end{aligned}$$

But we know that:

$$p_b = \frac{\mathbf{Pr}_\kappa[\kappa_b = b \wedge Q \geq 2]}{\mathbf{Pr}_\kappa[Q \geq 2]}$$

And thus $\mathbf{Pr}_\kappa[\kappa_b = b \mid Q = 1] = p_b$. Applying the above argument inductively produces $\mathbf{Pr}_\kappa[\kappa_b = b \mid Q = k] = p_b$.

Part 3. An after-walk is represented by a sequence $\kappa = (a, e_1, s_1, \ldots, e_\ell, s_\ell)$ with $a \in V$, $e_i \in [d]$ and $s_i \in \{0, 1\}$, where each entry is chosen independently. We obtain a new distribution by applying the probabilistic isomorphism $\kappa \mapsto (\kappa_b, e_\ell, 0, \ldots, e_2, 0, e_1, 1)$. Let the walk $\tau$ be drawn according to the new distribution, then:

$$\mathbf{Pr}_\kappa[\kappa_a = a \mid Q = k] = \mathbf{Pr}_\tau[\tau_b = a \mid W^{\#}_{\tau, v \to u} = k] = \mathbf{Pr}_{\tau^*}[\tau^*_b = a \mid \tau^*_a = u]$$

Part 1. It is easily seen that the value of $\kappa_a$ conditioned on $Q = k$ depends only on $u$. Similarly, the value of $\kappa_b$ conditioned on $Q = k$ depends only on $v$. Therefore $\kappa_a$ and $\kappa_b$ are independent. ∎

We start by defining a fictitious coloring $\omega$ of $G$ (induced by $\omega^\infty$) which will be used only for analysis purposes. For every $u \in V$ define:

$$\omega(u) = \arg\max_{\sigma \in \Sigma} \mathbf{Pr}_{\kappa^*}[\omega^\infty_{\kappa^*_b}(u) = \sigma \mid \kappa^*_a = u]$$

Let $F \subseteq E$ be a maximal set of unsatisfied edges of $G$ under $\omega$ such that $|F| \leq 1/t$. Since $G$ is $\epsilon$-far from satisfiable, we have that $\min(\epsilon, 1/2t) \leq |F|/|E| \leq 1/t$.

Let us now consider the random walk $\kappa$ taken by the verifier. We call an edge $e_i$ *faulty* if $e_i \in F$ and $\omega^\infty(a)$ and $\omega^\infty(b)$ fail the test at $e_i$. Let $X$ be a random variable that equals the number of faulty edges on the verifier's walk. We will show the following two claims:

**Claim 1.1** (Average analysis). $\mathbf{E}[X] \geq \alpha t \cdot |F|/|E|$, *where* $\alpha = \frac{\sqrt{5}-1}{2}$.

**Claim 1.2** (Variance analysis). $\mathbf{E}[X^2] \leq 4\beta t \cdot |F|/|E|$, *where* $\beta = 1 + d^2/\eta^2$.

**Lemma 1.3** (Theorem 4.3.1 in [5]). *For every integral non-negative random variable* $X$*:* $\mathbf{Pr}[X > 0] \geq \mathbf{E}[X]^2/\mathbf{E}[X^2]$.

Together they give us our main result:

$$\mathbf{Pr}_{(e_1,\ldots,e_T)}[\text{some } e_i \text{ is faulty}] \geq \mathbf{Pr}[X > 0] \geq \frac{\mathbf{E}[X]^2}{\mathbf{E}[X^2]} \geq \frac{\alpha^2}{4\beta} \cdot t \cdot \frac{|F|}{|E|} \geq \frac{\alpha^2}{4\beta} \cdot t \cdot \min(\epsilon, 1/2t)$$

In the following two sections we prove the two claims. Then in the final section we give an argument that explains why the truncated random walk (simulated by $G^t$) behaves essentially identically to the "ideal" non-truncated one analyzed so far.

### 1.4.1 Average analysis

In this section, when we talk about an edge $e \in E$ traversed by the verifier's walk we silently do not distinguish the direction of traversal. This makes the exposition cleaner. To complete the argument one must apply it twice: once for each direction. Or, equivalently, assume that $G$ is directed where each undirected edge is replaced by two opposing directed ones.

Let $\kappa$ be the random walk of the verifier, and condition on the event that a fixed edge $e = (u, v) \in F$ is in $\kappa$. $e$ is faulty if at least one of the following three events holds:

1. $A_1 = \mathbf{I}\big[\omega_a^\infty(u) \neq \omega_b^\infty(u)\big]$, or
2. $A_2 = \mathbf{I}\big[\omega_a^\infty(v) \neq \omega_b^\infty(v)\big]$, or
3. $A_3 = \mathbf{I}\big[\omega_a^\infty(u) = \omega(u) \wedge \omega_b^\infty(v) = \omega(v)\big]$. (In this case $c_e(\omega(u), \omega(v)) = 0$ since $e \in F$ by assumption.)

Since the above events are not independent, we have:

$$\mathbf{Pr}[e \text{ is faulty} \mid e \in \kappa] \geq \max\big(\mathbf{Pr}[A_1], \mathbf{Pr}[A_2], \mathbf{Pr}[A_3]\big)$$

Let $p_a$ be the probability that $\omega_a^\infty(u) = \omega(u)$, and $p_b$ be the probability that $\omega_b^\infty(v) = \omega(v)$, both of which are independent from each other by Lemma 1.2. Thus $\mathbf{Pr}[A_3] = p_a p_b$.

We turn our attention to $A_1$ now. A little thought would convince the reader that:

$$\mathbf{Pr}[A_1] = \mathbf{Pr}[\omega_a^\infty(u) \neq \omega_b^\infty(u)] \geq 1 - \mathbf{Pr}_{\kappa'}[\omega_{\kappa'_b}^\infty(u) = \omega(u) \mid \kappa'_a = u] = 1 - p_a$$

This follows from the fact that if $\rho_1, \rho_2 : \mathbb{Z} \to \mathbb{R}$ are two distributions over the integers, then:

$$\mathbf{Pr}_{x \sim \rho_1, y \sim \rho_2}[x \neq y] \geq \max_{i \in [2]} \{1 - \max_z \rho_i(z)\}$$

Similarly $\mathbf{Pr}[A_2] \geq 1 - p_b$. Thus:

$$\mathbf{Pr}[e \text{ is faulty} \mid e \in \kappa] \geq \max\,(1 - p_a, 1 - p_b, p_a p_b) \geq \frac{\sqrt{5} - 1}{2} = \alpha$$

Now, let $e$ be a directed version of $e = (u, v) \in E$. Let the random variable $X_e$ equal the number of faulty occurrences of $e$ in the verifier's walk in the specified direction, and let $Y_e$

denote the number of all occurrences. (Note that if one occurrence is faulty, then all are.) Since the stationary distribution of this walk is uniform, it is easily checked that $Y_e = t/2|E|$. Also let $Z_e$ indicate whether $e$ has any faulty occurrences in the walk, hence $X_e = Y_e Z_e$. Observe that $Z_e$ depends solely on $\kappa_a$ and $\kappa_b$, which are independent of $k$ as implied by Lemma 1.2. Thus:

$$\mathbf{Pr}[Y_e = k \wedge Z_e = 1 \mid Y_e \geq 1] = \mathbf{Pr}[Y_e = k \mid Y_e \geq 1] \cdot \mathbf{Pr}[Z_e = 1 \mid Y_e \geq 1],$$

which we prove shortly, we have:

$$\begin{aligned}
\mathbf{E}[X] &= \sum_{e \in F} \mathbf{E}[X_e] \\
&= \sum_{e \in F} \mathbf{E}[Y_e Z_e] \\
&= \sum_{e \in F} \sum_{k \geq 1} \mathbf{Pr}[Y_e Z_e = k] \cdot k \\
&= \sum_{e \in F} \mathbf{Pr}[Y_e \geq 1] \sum_{k \geq 1} \mathbf{Pr}[Y_e = k \wedge Z_e = 1 \mid Y_e \geq 1] \cdot k \\
&= \sum_{e \in F} \mathbf{Pr}[Y_e \geq 1] \sum_{k \geq 1} \mathbf{Pr}[Y_e = k \mid Y_e \geq 1] \cdot \mathbf{Pr}[Z_e = 1 \mid Y_e \geq 1] \cdot k \\
&= \sum_{e \in F} \mathbf{Pr}[Z_e = 1 \mid Y_e \geq 1] \sum_{k \geq 1} \mathbf{Pr}[Y_e \geq 1] \cdot \mathbf{Pr}[Y_e = k \mid Y_e \geq 1] \cdot k \\
&= \sum_{e \in F} \mathbf{Pr}[Z_e = 1 \mid Y_e \geq 1] \cdot \mathbf{E}[Y_e] \\
&= \alpha t \cdot \frac{|F|}{|E|}
\end{aligned}$$

### 1.4.2 Variance analysis

To prove the variance claim, we will need to use the expansion properties of the graph $G$ using the following theorem, which is a restatement of the well-known "bit-recycling" theorem:

**Theorem 1.1.** *For $j > i$:*

$$\mathbf{Pr}[e_j \in F | e_i \in F] \leq \left(1 - \frac{1}{t}\right)^{j-1} \left(\frac{|F|}{|E|} + \left(1 - \frac{\eta^2}{d^2}\right)^{j-i-1}\right)$$

This theorem roughly says that in the random walk the events of the form $e_i \in F$ are approximately pairwise independent.

Let $B_i = \mathbf{I}[e_i \in F]$. Then $\mathbf{Pr}[B_i = 1] = |F|/|E| \cdot \big(1 - 1/t\big)^{i-1}$, and:

$$\begin{aligned}
\mathbf{E}[X^2] &= \mathbf{E}\left[\left(\sum_{i=1}^{\infty}\right)^2\right] \\
&\leq 2 \sum_{1 \leq i \leq j \leq \infty} E[B_i B_j] \\
&= 2 \sum_{i=1}^{\infty} \mathbf{Pr}[B_i = 1] \sum_{i \leq j} \mathbf{Pr}[B_j = 1 \mid B_i = 1] \\
&\leq 2 \sum_{i=1}^{\infty} \mathbf{Pr}[B_i = 1] \left(1 + \sum_{l \geq 1} (1 - 1/t)^l \left(|F|/|E| + \big(1 - \eta^2/d^2\big)^{l-1}\right)\right) \\
&\leq 2 \sum_{i=1}^{\infty} \mathbf{Pr}[B_i = 1] \left(1 + \sum_{l \geq 1} (1 - 1/t)^l \cdot |F|/|E| + \big(1 - \eta^2/d^2\big)^{l-1}\right) \\
&\leq 2t \cdot \frac{|F|}{|E|} \left(1 + t\frac{|F|}{|E|} + \frac{d^2}{\eta^2}\right)
\end{aligned}$$

The variance analysis completes after recalling that $|F|/|E| \leq 1/t$.

### 1.4.3 Truncation argument

We have thus far shown that the ideal verifier, who takes a random walk of unlimited length, rejects every bad coloring with constant probability:

$$\mathbf{Pr}_\kappa[X \geq 1] \geq Q \cdot t \cdot \min(\text{UNSAT}\ (G), 1/2t)$$

for some constant $Q$. Let $\ell(\kappa)$ denote the length of the random walk $\kappa$. Observe that for any fixed $t \geq 1$ we have:

$$\mathbf{Pr}_\kappa[X \geq 1] = \underbrace{\mathbf{Pr}[\ell(\kappa) \leq t]}_{A} \cdot \underbrace{\mathbf{Pr}[X \geq 1 \mid \ell(\kappa) \leq t]}_{B} + \underbrace{\mathbf{Pr}[\ell(\kappa) > t]}_{C} \cdot \underbrace{\mathbf{Pr}[X \geq 1 \mid \ell(\kappa) > t]}_{D}$$

We can choose $t = t(Q)$ so that $C$ becomes arbitrarily small. Thus, $AB$ would have to account for most of the mass in $\mathbf{Pr}_\kappa[X \geq 1]$. This can be re-interpreted as saying that: If the verifier automatically accepts whenever the random walk exceeds $t$ steps, it still has a high probability of rejecting unsat instances.

So as long as each vertex keeps opinions about all vertices' colors in radius $t$, the real verifier will work just as well as the ideal one. Note that by discarding all walks longer than $t$ we are ensuring that the fictitious coloring $\omega^\infty$ is only ever queried in its restriction to $\omega^t$.

## References

[1] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy, *Proof verification and intractability of approximation problems*, in J. ACM'98

[2] S. Arora and S. Safra, *Probabilistic checking of proofs: A new characterization of NP*, in J. ACM'98

[3] Irit Dinur, *The PCP Theorem by Gap Amplification*, in STOC'06

[4] Jaikumar Radhakrishnan and Madhu Sudan, *On Dinur's Proof of the PCP Theorem*, in Bulletin of the AMS

[5] Noga Alon and Joel Spencer, *The Probabilistic Method*, A Wiley Interscience Publication

[6] Venkatesan Guruswami and Ryan O'Donnell, *The PCP Theorem and Hardness of Approximation*, A University of Washington course CSE533, `http://www.cs.washington.edu/education/courses/533/05au/`