

A New Multilayered PCP and the Hardness of Hypergraph Vertex Cover*

Irit Dinur [†] Venkatesan Guruswami [‡] Subhash Khot [§] Oded Regev [¶]

April, 2004

Abstract

Given a k -uniform hypergraph, the Ek -Vertex-Cover problem is to find the smallest subset of vertices that intersects every hyperedge. We present a new multilayered PCP construction that extends the Raz verifier. This enables us to prove that Ek -Vertex-Cover is NP-hard to approximate within a factor of $(k - 1 - \varepsilon)$ for arbitrary constants $\varepsilon > 0$ and $k \geq 3$. The result is nearly tight as this problem can be easily approximated within factor k . Our construction makes use of the biased Long-Code and is analyzed using combinatorial properties of s -wise t -intersecting families of subsets.

We also give a different proof that shows an inapproximability factor of $\lfloor \frac{k}{2} \rfloor - \varepsilon$. In addition to being simpler, this proof also works for super-constant values of k up to $(\log N)^{1/c}$ where $c > 1$ is a fixed constant and N is the number of hyperedges.

Keywords: Hypergraph Vertex Cover, Hardness of Approximation, PCP, Multilayered Outer Verifier, Long Code.

1 Introduction

A k -uniform hypergraph $H = (V, E)$ consists of a set of vertices V and a collection E of k -element subsets of V called hyperedges. A *vertex cover* of H is a subset $S \subseteq V$ such that every hyperedge in E intersects S , i.e., $e \cap S \neq \emptyset$ for each $e \in E$. An *independent set* in G is a subset whose complement is a vertex cover, or in other words a subset of vertices that contains no hyperedge entirely within it. The Ek -Vertex-Cover problem is the problem of finding a minimum size vertex cover in a k -uniform hypergraph. This problem is alternatively called the minimum hitting set problem with sets of size k (and is equivalent to the set cover problem where each element of the universe occurs in exactly k sets).

*A preliminary version of this paper appears in the *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, June 9-11, 2003.

[†]The Miller Institute, University of California at Berkeley, iritd@cs.berkeley.edu.

[‡]Department of Computer Science, University of Washington, Seattle, WA 98195, venkat@cs.washington.edu. Part of this work was done while the author was at UC Berkeley as a Miller Research Fellow.

[§]Institute for Advanced Study, Princeton, NJ 08544 and Georgia Institute of Technology, Atlanta, GA 30332; khot@math.ias.edu. Work done while at Princeton University, supported by Prof. Sanjeev Arora's David and Lucile Packard Fellowship, NSF Grant CCR-0098180, and an NSF ITR grant.

[¶]Department of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel, odedr@cs.tau.ac.il. Most of this work was done while the author was at the Institute for Advanced Study, Princeton, NJ and supported by NSF grant CCR-9987845.

The Ek -Vertex-Cover problem is a fundamental NP-hard optimization problem. For $k = 2$, it is just the famous vertex cover problem on graphs. Owing to its NP-hardness, one is interested in how well it can be approximated in polynomial time. A very simple algorithm that is often taught in an undergraduate algorithms class is the following: greedily pick a maximal set of pairwise disjoint hyperedges and then include all vertices in the chosen hyperedges in the vertex cover. It is easy to show that this gives a factor k approximation algorithm for Ek -Vertex-Cover. State of the art techniques yield only a tiny improvement, achieving a $k - o(1)$ approximation ratio [14]. This raises the question whether achieving an approximation factor of $k - \varepsilon$ for any constant $\varepsilon > 0$ could be NP-hard. In this paper, we prove a nearly tight hardness result for Ek -Vertex-Cover,

Theorem 1.1 (Main Theorem) *For every integer $k \geq 3$ and every $\varepsilon > 0$, it is NP-hard to approximate Ek -Vertex-Cover within a factor of $(k - 1 - \varepsilon)$.*

Previous Hardness Results

The vertex cover problem on hypergraphs where the size of the hyperedges is unbounded is nothing but the Set-Cover problem. For this problem there is a $\ln n$ approximation algorithm [23, 20] and a matching hardness factor of $(1 - o(1)) \ln n$ due to Feige [10]. Feige showed that an approximation algorithm achieving a factor of $(1 - o(1)) \ln n$ would imply $\text{NP} \subseteq \text{DTIME}(n^{O(\log \log n)})$, where n is the number of hyperedges. The best NP-hardness result (where the reduction is polynomial time) is due to Raz and Safra [25], who showed that it is NP-hard to approximate Set-Cover to within $c \cdot \log n$ for some small $c > 0$. The first explicit hardness result shown for Ek -Vertex-Cover was due to Trevisan [27] who considered the approximability of bounded degree instances of several combinatorial problems, and specifically showed an inapproximability factor of $\Omega(k^{1/19})$ for Ek -Vertex-Cover. Holmerin [18] showed that E4-Vertex-Cover is NP-hard to approximate within $(2 - \varepsilon)$. Independently, Goldreich [12] showed a direct “FGLSS”-type [11] reduction (involving no use of the long-code, a crucial component in most recent PCP constructions) attaining a hardness factor of $(2 - \varepsilon)$ for Ek -Vertex-Cover for some constant k . Later, Holmerin [19] showed that Ek -Vertex-Cover is NP-hard to approximate within a factor of $\Omega(k^{1-\varepsilon})$, and also that it is NP-hard to approximate E3-Vertex-Cover within factor $(3/2 - \varepsilon)$.

More recently Dinur, Guruswami and Khot gave a fairly simple proof of an $\alpha \cdot k$ hardness result for Ek -Vertex-Cover, for some $\alpha > \frac{1}{3}$. The proof takes a combinatorial view of Holmerin’s construction and instead of Fourier analysis uses some properties concerning intersecting families of finite sets. The authors also give a more complicated reduction that shows a factor $(k - 3 - \varepsilon)$ hardness for Ek -Vertex-Cover. The crucial impetus for that work came from the recent result of Dinur and Safra [9] on the hardness of approximating vertex cover (on graphs), and as in [9] the notion of biased long codes and some extremal combinatorics relating to intersecting families of sets play an important role. In addition to ideas from [9], the factor $(k - 3 - \varepsilon)$ hardness result also exploits the notion of covering complexity introduced by Guruswami, Håstad and Sudan [13]. Neither the $\alpha \cdot k$ result nor the $k - 3 - \varepsilon$ result has been published (an ECCC manuscript exists, [7]) since they have been subsumed by the work presented herein.

Our result and techniques

In this paper we improve upon all the above hardness results by proving a factor $(k - 1 - \varepsilon)$ inapproximability result for Ek -Vertex-Cover. Already for $k = 3$, this is an improvement from $3/2 - \varepsilon$ to $2 - \varepsilon$. Improving our hardness factor from $(k - 1 - \varepsilon)$ to $(k - \varepsilon)$ appears highly non-trivial (although it was recently proven under a certain conjecture [22]). Note that such a bound would

imply a factor $2 - \varepsilon$ hardness for vertex cover on graphs, a problem that is notoriously difficult. While our proof shares some of the extremal combinatorics flavor of [9] and [7], it draws its strength mainly from a new multilayered outer verifier system for NP languages. This multilayered system is constructed using the Raz verifier [24] as a building block.

The Raz verifier, which serves as the starting point or “outer verifier” in most (if not all) recent hardness results, can be described as follows. There are two sets of (non-Boolean) variables Y and Z , and for certain pairs of $y \in Y$ and $z \in Z$, a constraint $\pi_{y \rightarrow z}$. The constraints are projections, i.e., for each assignment to y there exists exactly one assignment to z such that the constraint $\pi_{y \rightarrow z}$ is satisfied. The goal is to find an assignment A to the variables so that a maximum number of constraints $\pi_{y \rightarrow z}$ are satisfied, i.e., have the property $\pi_{y \rightarrow z}(A(y)) = A(z)$. By the PCP Theorem [2, 1] together with the Parallel Repetition Theorem [24], we know that for every $\varepsilon > 0$ it is NP-hard to distinguish between the case where all the constraints can be satisfied and the case where no more than a fraction ε of the constraints can be satisfied.

In [7], the $\alpha \cdot k$ hardness result is obtained by replacing every Y variable by a block of vertices (representing its Long-Code). Hyperedges connect y_1 -vertices to y_2 -vertices only if there is some $z \in Z$ such that $\pi_{y_1 \rightarrow z}, \pi_{y_2 \rightarrow z}$ are constraints in the system. This construction has an inherent symmetry between blocks which deteriorates the projection property of the constraints, limiting the hardness factor one can prove to at most $k/2$. Since this is a relatively simple reduction, we include a proof showing a hardness of approximation factor of $(\lfloor k/2 \rfloor - \varepsilon)$. The improvement over the factor $k/3$ in [7] is obtained using a better result on s -wise t -intersecting set families. We also remark that this result itself suffices for the recent reduction from Ek -Vertex-Cover to asymmetric K -center [5]. Moreover, this result has the advantage of working for much larger super-constant values of k (up to $(\log N)^{1/c}$ for some absolute constant c where N is the number of hyperedges).

Another way of reducing the Raz verifier to Ek -Vertex-Cover is by maintaining the asymmetry between Y and Z , introducing a block of vertices for each variable in Y and in Z (representing their Long-Code). Each constraint $\pi_{y \rightarrow z}$ can be emulated by a set of hyperedges, where each hyperedge consists of both y -vertices and z -vertices. The hyperedges can be chosen so that if the initial PCP instance was satisfiable, then taking a particular fraction $1/k$ of the vertices in each block will be a vertex cover. However, this reduction has a basic “bipartiteness” flaw: the underlying constraint graph, being bipartite with parts Y and Z , has a vertex cover of size at most one half of the number of vertices. Taking all the vertices of, say, the Z variables will be a vertex cover for the hypergraph regardless of whether or not the initial PCP instance was satisfiable. This, once again, limits the gap to no more than $k/2$.

We remark that this “bipartiteness” flaw naturally arises in other settings as well. One example is approximate hypergraph coloring, where indeed our multilayered PCP construction has been successfully used for showing hardness, see [8, 21].

The Multilayered PCP. We overcome the $k/2$ limit by presenting a new, multilayered PCP. In this construction we maintain the projection property of the constraints that is a strong feature of the Raz verifier, while overcoming the “bipartiteness” flaw. In the usual Raz verifier we have two “layers”, the first containing the Y variables and the second containing the Z variables. In the multilayered PCP, we have ℓ layers containing variables X_1, X_2, \dots, X_ℓ respectively. Between every pair of layers i_1 and i_2 , we have a set of projection constraints that represent an instance of the Raz verifier. In the multilayered PCP, it is NP-hard to distinguish between (i) the case where there exists an assignment that satisfies *all* the constraints, between *every* pair of layers, and (ii) the case where for *every* i_1, i_2 it is impossible to satisfy more than a fraction ε of the constraints between X_{i_1} and X_{i_2} .

In addition, we prove that the underlying constraint graph no longer has the “bipartiteness” obstacle, i.e. it no longer has a small vertex cover and hence a large independent set. Indeed we show that the multilayered PCP has a certain “weak-density” property: for any set containing an ε fraction of the variables there are many constraints between variables of this set. This guarantees that “fake” independent sets in the hypergraph (i.e., independent sets that occur because there are no constraints between the variables of the set) contain at most ε fraction of the vertices.

We mention that the PCP presented by Feige in [10] has a few structural similarities with ours. Most notably, both have more than two types of variables. However, while in our construction the types are layered with decreasing domain sizes, in Feige’s construction the different types are all symmetric. Furthermore, and more importantly, the constraints tested by the verifier in Feige’s construction are not projections while this is a key feature of our multilayered PCP, crucially exploited in our analysis.

We view the construction of the multilayered PCP as a central contribution of our paper, and believe that it could be a powerful starting point for other hardness of approximation results as well. In fact, as mentioned above, our multilayered construction has already been used in obtaining strong hardness results for coloring 3-uniform hypergraphs [8, 21] (namely the hardness of coloring a 2-colorable 3-uniform hypergraph using an arbitrary constant number of colors), a problem for which no non-trivial inapproximability results are known using other techniques. We anticipate that this new outer verifier will also find other applications besides the ones in this paper and in [8, 21].

The Biased Long-Code. Our hypergraph construction relies on the Long-Code that was introduced in [3], and more specifically, on the biased Long-Code defined in [9]. Thus, each PCP variable is represented by a block of vertices, one for each “bit” of the biased Long-Code. More specifically, in x ’s block we have one vertex for each subset of R , where R is the set of assignments for the variable x . However, rather than taking all vertices in a block with equal weight, we attach weights to the vertices according to the p -biased Long-Code. The weight of a subset F is set to $p^{|F|}(1-p)^{|R \setminus F|}$, highlighting subsets of cardinality $p \cdot |R|$. Thus we actually construct a weighted hypergraph; we will later describe how it can be translated to a non-weighted one.

The vertex cover in the hypergraph is shown to have relative size of either $1-p$ in the good case or almost 1 in the bad case. Choosing large $p = 1 - \frac{1}{k-1-\varepsilon}$, yields the desired gap of $\frac{1}{1-p} \approx k-1-\varepsilon$ between the good and bad cases. The reduction uses the following combinatorial property: *a family of subsets of a set R , where each subset has size $p|R|$, either contains very few subsets, or it contains some $k-1$ subsets whose common intersection is very small.* We will later show that this property holds for $p < 1 - \frac{1}{k-1}$ and hence we obtain a gap of $k-1-\varepsilon$. As can be seen, this property does not hold for $p > 1 - \frac{1}{k-1}$ and hence one cannot improve the $k-1-\varepsilon$ result by simply increasing p .

Weighted versus unweighted

As mentioned above, our construction yields a weighted hypergraph: each vertex is associated with a weight and the goal is to minimize the weight of the vertex cover. By appropriately duplicating vertices, we can obtain an unweighted hypergraph (for more detail, see [6, 9]). We note that in all of our constructions, duplicating vertices does not change the asymptotic size of the hypergraph and hence all of our results carry over to the unweighted case.

Location of the gap

All our hardness results have the gap between sizes of the vertex cover at the “strongest” location. Specifically, to prove a factor $(k - 1 - \varepsilon)$ hardness we show that it is hard to distinguish between k -uniform hypergraphs that have a vertex cover of weight $\frac{1}{k-1} + \varepsilon$ from those whose minimum vertex cover has weight at least $(1 - \varepsilon)$. This result is stronger than a gap of about $(k - 1)$ achieved, for example, between vertex covers of weight $\frac{1}{(k-1)^2}$ and $\frac{1}{k-1}$. In fact, by adding dummy vertices, our result implies that for any $c < 1$ it is NP-hard to distinguish between hypergraphs whose minimum vertex cover has weight at least c from those which have a vertex cover of weight at most $(\frac{c}{k-1} + \varepsilon)$. Put another way, our result shows that for k -uniform hypergraphs, for $k \geq 3$, there exist an $\alpha = \alpha(k) > 0$ such that for arbitrarily small $\varepsilon > 0$, it is NP-hard to find an independent set consisting of a fraction ε of the vertices even if the hypergraph is promised to contain an independent set comprising a fraction α of the vertices. We remark that such a result is not known for graphs and seems out of reach of current techniques. (The 1.36 hardness result for vertex cover on graphs due to Dinur and Safra [9], for example, shows that it is NP-hard to distinguish between cases when the graph has an independent set of size $0.38 \cdot n$ and when no independent set has more than $0.16 \cdot n$ vertices.) This gap location feature was crucial to the recent tight $\Omega(\log^* n)$ inapproximability result for asymmetric K -center [5].

Organization

We begin in Section 2.1 by developing the machinery from extremal combinatorics concerning intersecting families of sets that will play a crucial role in our proof. This is followed by a statement of the starting point of our reduction, which is the standard gap instance obtained by combining the PCP theorem with Raz’s parallel repetition theorem. In Section 3, we present a relatively simple reduction that shows an inapproximability factor of $\lfloor k/2 \rfloor - \varepsilon$ for arbitrary $\varepsilon > 0$. As explained earlier, this is as far as we can go using just the “bipartite” Raz outer verifier. In Section 4 we present the multilayered PCP construction, which gives hope to break the $k/2$ barrier. In Section 5, we present our reduction to a gap version of Ek -Vertex-Cover which allows us to prove a factor $(k - 1 - \varepsilon)$ inapproximability result for this problem. Finally, we discuss the case when k is not a constant but instead is a growing function of the number of hyperedges in Section 6.

2 Preliminaries

2.1 Intersecting Families

In this section we define the notion of an s -wise t -intersecting family and prove an important property of such families. For a comprehensive survey, see [15]. Denote $[n] = \{1, \dots, n\}$ and $2^{[n]} = \{F \mid F \subseteq [n]\}$. We start with a definition:

Definition 2.1 *A family $\mathcal{F} \subseteq 2^{[n]}$ is called s -wise t -intersecting if for every s sets $F_1, \dots, F_s \in \mathcal{F}$, we have $|F_1 \cap \dots \cap F_s| \geq t$.*

For example, the family of all sets $F \in 2^{[n]}$ such that $[t] \subseteq F$ is an s -wise t -intersecting family. A more interesting example is given by the family of all sets $F \in 2^{[n]}$ such that $|F \cap [t+s]| \geq t+s-1$. Each set in this family has at most one “hole” in $[t+s]$. Therefore, any intersection of s sets of this family contains at least $t+s-s = t$ elements from the range $[t+s]$. More generally, for any $j \geq 0$, we can define the family of all sets $F \in 2^{[n]}$ such that $|F \cap [t+js]| \geq t+(s-1)j$; it is easy to see that this is an s -wise t -intersecting family.

We need another important definition:

Definition 2.2 For a bias parameter $0 < p < 1$, and a ground set R , the weight of a set $F \subseteq R$ is

$$\mu_p^R(F) \stackrel{\text{def}}{=} p^{|F|} \cdot (1-p)^{|R \setminus F|}$$

When R is clear from the context we write μ_p for μ_p^R . The weight of a family $\mathcal{F} \subseteq 2^R$ is $\mu_p(\mathcal{F}) = \sum_{F \in \mathcal{F}} \mu_p(F)$.

The weight of a subset is precisely the probability of obtaining this subset when one picks every element in R independently with probability p .

The following is the main lemma of this section. It shows that for any $p < \frac{s-1}{s}$, a family of non-negligible μ_p -weight (i.e., $\mu_p(\mathcal{F}) \geq \varepsilon$) cannot be s -wise t -intersecting for sufficiently large t .

Lemma 2.3 For arbitrary $\varepsilon, \delta > 0$, and integer $s \geq 2$, let $p = 1 - \frac{1}{s} - \delta$. Then, there exists $t = t(\varepsilon, s, \delta)$ such that for any s -wise t -intersecting family $\mathcal{F} \subseteq 2^{[n]}$, $\mu_p(\mathcal{F}) < \varepsilon$. Moreover, it is enough to choose $t(\varepsilon, s, \delta) = \Omega(\frac{1}{\delta^2}(\log \frac{1}{\varepsilon} + \log(1 + \frac{1}{s\delta^2})))$.

Proof: In order to prove this lemma, we need to introduce the notion of a left-shifted family. Performing an (i, j) -shift on a family consists of replacing the element j with the element i in all sets $F \in \mathcal{F}$ such that $j \in F$, $i \notin F$ and $(F \setminus \{j\}) \cup \{i\} \notin \mathcal{F}$. A left-shifted family is a family which is invariant with respect to (i, j) -shifts for any $1 \leq i < j \leq n$. For any family \mathcal{F} , by iterating the (i, j) -shift for all $1 \leq i < j \leq n$ we eventually get a left-shifted family which we denote by $S(\mathcal{F})$. The following simple lemma summarizes the properties of the left-shift operation.

Lemma 2.4 ([15], p. 1298, Lemma 4.2) For any family $\mathcal{F} \subseteq 2^{[n]}$, there exists a one-to-one and onto mapping τ from \mathcal{F} to $S(\mathcal{F})$ such that $|F| = |\tau(F)|$ for every $F \in \mathcal{F}$. In other words, left-shifting a family maintains its size and the size of the sets in the family. Moreover, if \mathcal{F} is an s -wise t -intersecting family then so is $S(\mathcal{F})$. ■

It can be easily checked that the examples given after Definition 2.1 are all left-shifted. As the next lemma shows, those examples essentially represent all possible sets in any left-shifted s -wise t -intersecting family:

Lemma 2.5 ([15], p. 1311, Lemma 8.3) Let $\mathcal{F} \subseteq 2^{[n]}$ be a left-shifted s -wise t -intersecting family. Then, for every $F \in \mathcal{F}$, there exists a $j \geq 0$ with $|F \cap [t + sj]| \geq t + (s-1)j$.

Proof: For completeness, we sketch the proof of the lemma. For two equally sized sets $A = \{a_1, \dots, a_l\}$, $1 \leq a_1 < \dots < a_l \leq n$ and $B = \{b_1, \dots, b_l\}$, $1 \leq b_1 < \dots < b_l \leq n$ we say that $A \preceq B$ if $a_i \leq b_i$ for all $i = 1, \dots, l$. Then, we claim that for such A, B , if $[n] \setminus A \in \mathcal{F}$ then also $[n] \setminus B \in \mathcal{F}$. We prove this by induction: for $i = 0, \dots, l$ define the set $F_i = [n] \setminus \{a_1, \dots, a_i, b_{i+1}, \dots, b_l\}$. Notice that $F_l = [n] \setminus A$ and therefore $F_l \in \mathcal{F}$. Next, we show that $F_i \in \mathcal{F}$ implies that $F_{i-1} \in \mathcal{F}$. If $a_i = b_i$ then the claim is obvious. Otherwise, $a_i < b_i$ and hence $b_i \in F_i$. Since \mathcal{F} is left-shifted and $a_i \notin F_i$ it follows that $F_i \setminus \{b_i\} \cup \{a_i\} = F_{i-1}$ is in \mathcal{F} . This proves our claim since $F_0 = [n] \setminus B$.

Let us now complete the proof of the lemma. Assume on the contrary that there exists $F \in \mathcal{F}$ such that for all $j \geq 0$, $|F \cap [t + sj]| < t + (s-1)j$. Let $A = \{a_1, \dots, a_l\}$ be such that $F = [n] \setminus A$ and assume that $a_1 < \dots < a_l$. The above condition implies that F contains at least i ‘‘holes’’ in $[t + (i-1)s]$ and therefore $a_i \leq t + (i-1)s$. It also implies that $l \geq \lfloor (n-t)/s \rfloor + 1$.

For $k = 0, \dots, s-1$, define the set $B_k = \{b_{k,1}, \dots, b_{k,l}\}$ by $b_{k,i} = \min\{t+k+(i-1)s, n-(l-i)\}$. Since $a_i \leq t + (i-1)s$ and $a_i \leq n - (l-i)$ (just because $a_i < a_{i+1} < \dots < a_l \leq n$ are all integers) we obtain that $a_i \leq b_{0,i}$. In other words, $A \preceq B_0$. Moreover, since $b_{k,i} \leq b_{k+1,i}$, we

see that $A \preceq B_0 \preceq B_1 \preceq \dots \preceq B_{s-1}$. Using the claim we proved before, we obtain that for all $k = 0, \dots, s-1$, $[n] \setminus B_k \in \mathcal{F}$. But this is a contradiction since $([n] \setminus B_0) \cap \dots \cap ([n] \setminus B_{s-1}) \subseteq [t-1]$ and in particular its size is less than t . \blacksquare

We now complete the proof of Lemma 2.3. We follow the general outline of the proof of Theorem 8.4 in [15], p. 1311. Let \mathcal{F} be an s -wise t -intersecting family where t will be determined later. According to Lemma 2.4, $S(\mathcal{F})$ is also s -wise t -intersecting and $\mu_p(S(\mathcal{F})) = \mu_p(\mathcal{F})$. By Lemma 2.5, for every $F \in S(\mathcal{F})$, there exists a $j \geq 0$ such that $|F \cap [t + sj]| \geq t + (s-1)j$. We can therefore bound $\mu_p(S(\mathcal{F}))$ from above by the probability that such a j exists for a random set chosen according to the distribution μ_p . We now prove an upper bound on this probability, which will give the desired bound on $\mu_p(S(\mathcal{F}))$ and hence also on $\mu_p(\mathcal{F})$.

The Chernoff bound [4] says that for a sequence of m independent random variables X_1, \dots, X_m on $\{0, 1\}$ such that for all i , $\Pr[X_i = 1] = p$ for some p ,

$$\Pr\left[\sum X_i > (p + \tau)m\right] \leq e^{-2m\tau^2}.$$

Hence, for any $j \geq 0$, $\Pr[|F \cap [t + sj]| \geq t + (s-1)j]$ is at most

$$\Pr[|F \cap [t + sj]| - p(t + sj) \geq \delta(t + sj)] \leq e^{-2(t+sj)\delta^2}.$$

Summing over all $j \geq 0$ we get:

$$\mu_p(S(\mathcal{F})) \leq \sum_{j \geq 0} e^{-2(t+sj)\delta^2} = e^{-2t\delta^2} / (1 - e^{-2s\delta^2}) \leq e^{-2t\delta^2} \left(1 + \frac{1}{2s\delta^2}\right)$$

which is smaller than ε for $t = \Omega\left(\frac{1}{\delta^2} \left(\log \frac{1}{\varepsilon} + \log\left(1 + \frac{1}{s\delta^2}\right)\right)\right)$. \blacksquare

2.2 The PCP Theorem and the Parallel Repetition Theorem

As is the case with many inapproximability results (e.g., [3], [16], [17], [26]), we begin our reduction from the Raz verifier described next. Let R be some parameter and let Ψ be a collection of two-variable constraints, where the variables are of two types, denoted Y and Z . Let R_Y denote the range of the Y -variables and R_Z the range of the Z -variables¹, where $|R_Z| \leq |R_Y|$ and both are at most $R^{O(1)}$. Assume each constraint $\pi \in \Psi$ depends on exactly one $y \in Y$ and one $z \in Z$, furthermore, for every value $a_y \in R_Y$ assigned to y there is exactly one value $a_z \in R_Z$ to z such that the constraint π is satisfied. Therefore, we can write each constraint $\pi \in \Psi$ as a function from R_Y to R_Z , and use notation $\pi_{y \rightarrow z} : R_Y \rightarrow R_Z$. Furthermore, we assume that the underlying constraint graph is bi-regular, i.e., every Y -variable appears in the same number of constraints in Ψ , and every Z -variable appears in the same number of constraints in Ψ . Both of these numbers are assumed to be at most $R^{O(1)}$.

The following theorem follows by combining the PCP Theorem with Raz's Parallel Repetition Theorem. The PCP given by this theorem will be called the Raz verifier henceforth.

Theorem 2.6 (PCP Theorem [1, 2] + Raz's Parallel Repetition Theorem [24]) *Let Ψ be as above. There exists a universal constant $\gamma > 0$ such that for every (large enough) constant R it is NP-hard to distinguish between the following two cases:*

¹Readers familiar with the Raz verifier may prefer to think concretely of $R_Y = [7^u]$ and $R_Z = [2^u]$ for some number u of repetitions.

- **Yes:** There is an assignment $A : Y \rightarrow R_Y, A : Z \rightarrow R_Z$ such that all $\pi \in \Psi$ are satisfied by A , i.e., $\forall \pi_{y \rightarrow z} \in \Psi, \pi_{y \rightarrow z}(A(y)) = A(z)$.
- **No:** No assignment can satisfy more than a fraction $\frac{1}{R^Y}$ of the constraints in Ψ . ■

Remark: The reduction, from 3SAT say, proving the above hardness can be assumed to run in time $n^{O(\log R)}$ where n is the size of the 3SAT instance. Also recall that the constraint graph of Ψ is bi-regular with Y and Z degrees d_l and d_r where both $d_l, d_r \leq R^{O(1)}$.

3 A Factor $k/2$ inapproximability result

In this section, we prove the factor $(\lfloor k/2 \rfloor - \varepsilon)$ hardness result for Ek -Vertex-Cover, for $k \geq 4$. We will show this by proving a factor $(k/2 - \varepsilon)$ hardness for even $k \geq 4$ (the result for odd k follows by a trivial reduction that adds a new vertex to every hyperedge of a hard instance of the $(k - 1)$ -uniform hypergraph).

Let $IS(G)$ denote the weight of vertices contained in the largest independent set of the hypergraph G .

Theorem 3.1 *Let $k \geq 4$ be an arbitrary even integer. Then for every $\varepsilon > 0$, it is NP-hard to approximate Ek -Vertex-Cover within a factor of $(k/2 - \varepsilon)$. More specifically, for every $\varepsilon, \delta > 0$, it is NP-hard to distinguish, given a k -uniform hypergraph G , between the following two cases:*

- $IS(G) \geq 1 - \frac{2}{k} - \delta$
- $IS(G) \leq \varepsilon$.

Proof: Start with a PCP instance, as given in Theorem 2.6, namely a set of local constraints Ψ over variables $Y \cup Z$, whose respective ranges are R_Y, R_Z . For parameters, we pick t to be larger than $t(\frac{\varepsilon}{2}, \frac{k}{2}, \delta)$ in Lemma 2.3; say, $t = O(\frac{1}{\delta^3} \log \frac{1}{\varepsilon})$ with some large enough constant. Moreover, take $R > (\frac{2t^2}{\varepsilon})^{1/\gamma}$ where $\gamma > 0$ is the universal constant from Theorem 2.6. From Ψ , we now construct a (weighted) k -uniform hypergraph G whose minimum vertex cover has weight $\approx 2/k$ or ≈ 1 depending on whether Ψ is satisfiable or not.

For a set R we denote by 2^R the power set of R . The vertex set of G is

$$V \stackrel{def}{=} Y \times 2^{R_Y}$$

namely for each $y \in Y$ we construct a block of vertices denoted $V[y] = \{y\} \times 2^{R_Y}$ corresponding to all possible subsets of R_Y . The weight of each vertex $\langle y, F \rangle \in V$ is

$$\Lambda(\langle y, F \rangle) \stackrel{def}{=} \frac{1}{|Y|} \cdot \mu_{1 - \frac{2}{k} - \delta}(F)$$

The hyperedges are defined as follows. For every pair of local-constraints $\pi_{y_1 \rightarrow z}, \pi_{y_2 \rightarrow z} \in \Psi$ sharing a common variable $z \in Z$, we add the hyperedge

$$\left\{ \langle y_1, F_1^1 \rangle, \langle y_1, F_2^1 \rangle, \dots, \langle y_1, F_{k/2}^1 \rangle \right\} \cup \left\{ \langle y_2, F_1^2 \rangle, \langle y_2, F_2^2 \rangle, \dots, \langle y_2, F_{k/2}^2 \rangle \right\}$$

if and only if there is no $r_1 \in \bigcap_{1 \leq j \leq k/2} F_j^1$ and $r_2 \in \bigcap_{1 \leq j \leq k/2} F_j^2$ such that $\pi_{y_1 \rightarrow z}(r_1) = \pi_{y_2 \rightarrow z}(r_2)$. Formally,

$$E \stackrel{def}{=} \bigcup_{\pi_{y_1 \rightarrow z}, \pi_{y_2 \rightarrow z} \in \Phi} \left\{ \left\{ \langle y_1, F_1^1 \rangle, \dots, \langle y_1, F_{k/2}^1 \rangle, \langle y_2, F_1^2 \rangle, \dots, \langle y_2, F_{k/2}^2 \rangle \right\} \mid \pi_{y_1 \rightarrow z}(F_1^1 \cap \dots \cap F_{k/2}^1) \cap \pi_{y_2 \rightarrow z}(F_1^2 \cap \dots \cap F_{k/2}^2) = \emptyset \right\}$$

where the union is taken over all pairs of local-constraints with a common variable z .

Lemma 3.2 (Completeness) *If Ψ is satisfiable then $IS(G) \geq 1 - \frac{2}{k} - \delta$.*

Proof: Assume a satisfying assignment $A : Y \cup Z \rightarrow R_Y \cup R_Z$ for Ψ . The following set is an independent set of G :

$$\mathcal{I} = \{\langle y, F \rangle \in V[y] \mid y \in Y, A(y) \in F\}. \quad (1)$$

For every hyperedge $e = \{\langle y_1, F_1^1 \rangle, \dots, \langle y_1, F_{k/2}^1 \rangle, \langle y_2, F_1^2 \rangle, \dots, \langle y_2, F_{k/2}^2 \rangle\}$ either $A(y_1) \notin F_1^1 \cap \dots \cap F_{k/2}^1$ or $A(y_2) \notin F_1^2 \cap \dots \cap F_{k/2}^2$, otherwise since $A(y_1), A(y_2)$ agree on every common Z -variable, we have $\pi_{y_1 \rightarrow z}(F_1^1 \cap \dots \cap F_{k/2}^1) \cap \pi_{y_2 \rightarrow z}(F_1^2 \cap \dots \cap F_{k/2}^2) \neq \emptyset$, and e would not have been a hyperedge.

Now note that the weight of the family $\{F \mid A(y) \in F\}$ w.r.t. the bias parameter $(1 - 2/k - \delta)$ is $(1 - 2/k - \delta)$. Hence the weight of the independent set \mathcal{I} in (1) is $1 - 2/k - \delta$. \blacksquare

Lemma 3.3 (Soundness) *If $IS(G) \geq \varepsilon$, then there exists an assignment $A : Y \cup Z \rightarrow R_Y \cup R_Z$ that satisfies more than a fraction $1/R^\gamma$ of the constraints in Ψ .*

Proof: Let $S \subset V$ be such an independent set of weight at least ε . We consider the set $Y' \subseteq Y$ of all variables y for which the weight (under Λ) of $S \cap V[y]$ in $V[y]$ is at least $\varepsilon/2$. A simple averaging argument shows that $|Y'| \geq |Y|\varepsilon/2$.

For each $y \in Y'$, define

$$\mathcal{F}_y = \{F \in 2^{R_Y} \mid \langle y, F \rangle \in S\}.$$

Thus we have $\forall y \in Y', \mu_{1-\frac{2}{k}-\delta}(\mathcal{F}_y) \geq \frac{\varepsilon}{2}$. By Lemma 2.3, applied with the choice $s = k/2$ and $\varepsilon/2$ in place of ε , there must exist $F_1^y, F_2^y, \dots, F_{k/2}^y \in \mathcal{F}_y$ such that $|F_1^y \cap F_2^y \cap \dots \cap F_{k/2}^y| \leq t$ where t is as defined before. Let us denote by $B(y)$ the intersection $F_1^y \cap F_2^y \cap \dots \cap F_{k/2}^y$; we will refer to $B(y)$ as the *core* of assignments for y . Intuitively, for any y for which the independent set S has large intersection with $V[y]$, we can decode a small collection of potential values from R_Y , specifically the core $B(y)$, that may be assigned to y .

We next translate these cores into an assignment satisfying more than a fraction $\frac{1}{R^\gamma}$ of the constraints in Ψ . Observe that for every $z \in Z$ and $y_1, y_2 \in Y', y_1 \neq y_2$ such that $\pi_{y_1 \rightarrow z}, \pi_{y_2 \rightarrow z} \in \Psi$, we must have

$$\pi_{y_1 \rightarrow z}(B(y_1)) \cap \pi_{y_2 \rightarrow z}(B(y_2)) \neq \emptyset. \quad (2)$$

Indeed, otherwise the set $\{\langle y_1, F_1^{y_1} \rangle, \dots, \langle y_1, F_{k/2}^{y_1} \rangle, \langle y_2, F_1^{y_2} \rangle, \dots, \langle y_2, F_{k/2}^{y_2} \rangle\}$ would be a hyperedge whose vertices all lie in S , contradicting the assumption that S is an independent set. Let us now slightly modify the definition of $B(y)$: for any y such that $B(y)$ is empty, we add to $B(y)$ some arbitrary element of R_Y . Notice that Equation 2 still holds; moreover, it holds also in the case that y_1 and y_2 are equal.

Let $Z' \subseteq Z$ denote the set of all Z variables of the instance Ψ that participate in some constraint with some $y \in Y'$. Formally, $Z' \stackrel{\text{def}}{=} \{z \mid \pi_{y \rightarrow z} \in \Psi \text{ for some } y \in Y'\}$. Associate each such $z \in Z'$ with an arbitrary $y \in Y'$ for which $\pi_{y \rightarrow z} \in \Psi$, and let $B(z) \stackrel{\text{def}}{=} \pi_{y \rightarrow z}(B(y)) \subseteq R_Z$. Now define a random assignment A by independently selecting for each $y \in Y', z \in Z'$ a random value from $B(y), B(z)$ respectively. Assign the rest of the variables $(Y \setminus Y') \cup (Z \setminus Z')$ with any arbitrary value. To complete the proof, we prove:

$$E_A \left[\left| \{ \pi_{y \rightarrow z} \in \Psi \mid \pi_{y \rightarrow z} \text{ is satisfied by } A \} \right| \right] \geq \frac{\varepsilon}{2t^2} \cdot |\Psi|. \quad (3)$$

Here the expectation is taken over the choice of the random assignment A . We will show that for every $y \in Y'$, each constraint $\pi_{y \rightarrow z} \in \Psi$ is satisfied by A with probability $\geq \frac{1}{t^2}$; thus the expected number of local-constraints satisfied by A is $\frac{|Y'|}{|Y|} \cdot \frac{1}{t^2} |\Psi| \geq \frac{\varepsilon}{2t^2} \cdot |\Psi|$ (because each $y \in Y$ appears in the same number of local constraints).

So, let $\pi_{y \rightarrow z} \in \Psi$ for arbitrary $y \in Y'$. Assume z is associated with some $y' \in Y'$ possibly equal to y . By Equation (2), we have

$$\pi_{y \rightarrow z}(B(y)) \cap B(z) = \pi_{y \rightarrow z}(B(y)) \cap \pi_{y' \rightarrow z}(B(y')) \neq \emptyset.$$

Therefore, there is at least one value $a_y \in B(y)$ such that $\pi_{y \rightarrow z}(a_y) \in B(z)$. Since for every $y \in Y'$, $|B(y)| \leq t$, there is at least a $\frac{1}{t^2}$ probability of having $\pi_{y \rightarrow z}(A(y)) = A(z)$, thereby showing (3).

Thus, there exists some assignment A that meets the expectation, which means it satisfies $\geq \frac{\varepsilon}{2t^2} > \frac{1}{R^\gamma}$ (by our choice of R) of the local constraints in Ψ . This completes the proof of the soundness Lemma 3.3. ■

Theorem 3.1 now follows from Lemmas 3.2 and 3.3 and Theorem 2.6. ■

Let us compute some parameters of the reduction, which will be useful in Section 6 where we will consider the case with super-constant values of k . First, let us compute R . The condition $R \geq (2t^2/\varepsilon)^{1/\gamma}$, together with our choice of t , implies that it suffices if

$$R = \left(\frac{1}{\varepsilon\delta} \right)^{O(1)}.$$

By the remark following Theorem 2.6, the number of variables in the PCP instance is at most $n^{O(\log R)} = n^{O(\log(1/\varepsilon\delta))}$ where n is the size of the original 3SAT instance. Since there is a block of $2^{R^{O(1)}}$ vertices corresponding to all subsets of R_y for each of these variables, the number of vertices in the hypergraph produced by the reduction is at most

$$n^{O(\log(1/\varepsilon\delta))} 2^{R^{O(1)}} \leq n^{O(\log(1/\varepsilon\delta))} 2^{(1/\varepsilon\delta)^{O(1)}}. \quad (4)$$

The degree of the hypergraph, i.e., the maximum number of hyperedges each vertex appears in, is at most $2^{R^{O(1)}k} \cdot R^{O(1)}$, since each variable in the PCP instance appears in at most $R^{O(1)}$ constraints. Hence, the degree is at most

$$2^{k/(\varepsilon\delta)^{O(1)}}.$$

Finally, the running time of the reduction is polynomial in the number of hyperedges. A bound on the latter can be obtained by combining the two previous bounds:

$$n^{O(\log(1/\varepsilon\delta))} 2^{(1/\varepsilon\delta)^{O(1)}} \cdot 2^{k/(\varepsilon\delta)^{O(1)}} \leq n^{O(\log(1/\varepsilon\delta))} 2^{k/(\varepsilon\delta)^{O(1)}}. \quad (5)$$

4 The Multilayered PCP

As discussed in the introduction, a natural approach to build a hypergraph from the PCP Ψ is to have a block of vertices for every variable y or z and define hyperedges of the hypergraph so as to enforce the constraints $\pi_{y \rightarrow z}$. For every constraint $\pi_{y \rightarrow z}$, there will be hyperedges containing vertices from the block of y and the block of z . However, this approach is limited by the fact that the constraint graph underlying the PCP has a small vertex cover. Since each hyperedge contains

vertices from both the Y and Z “sides”, the subset of all vertices on the Y (resp. Z) “side”, already covers all of the hyperedges regardless of whether the initial PCP system was satisfiable or not.

This difficulty motivates our construction of a multilayered PCP where we have many types of variables (rather than only Y and Z) and the resulting hypergraph is *multipartite*. The multilayered PCP is able to maintain the properties of Theorem 2.6 between *every* pair of layers. Moreover, the underlying constraint graph has a special “weak-density” property that roughly speaking, guarantees that it will have only tiny independent sets (thus any vertex cover for it must contain almost all of the vertices).

4.1 Layering the Variables

Let $\ell, R > 0$. Let us begin by defining an ℓ -layered PCP. In an ℓ -layered PCP there are ℓ sets of variables denoted by X_1, \dots, X_ℓ . The range of variables in X_i is denoted R_i , with $|R_i| = R^{O(\ell)}$. For every $1 \leq i < j \leq \ell$ there is a set of constraints Φ_{ij} where each constraint $\pi \in \Phi_{ij}$ depends on exactly one $x \in X_i$ and one $x' \in X_j$. For any two variables we denote by $\pi_{x \rightarrow x'}$ the constraint between them if such a constraint exists. Moreover, the constraints in Φ_{ij} are projections from x to x' , that is, for every assignment to x there is exactly one assignment to x' such that the constraint is satisfied.

In addition, as mentioned in the introduction, we would like to show a certain “weak-density” property of our multilayered PCP:

Definition 4.1 *An ℓ -layered PCP is said to be weakly-dense if for any δ , $2/\ell < \delta < 1$, given $m \geq \lceil \frac{2}{\delta} \rceil$ layers $i_1 < \dots < i_m$, and given any sets $S_j \subseteq X_{i_j}$ for $j \in [m]$ such that $S_j \geq \delta |X_{i_j}|$, there always exist two sets S_j and $S_{j'}$ such that the number of constraints between them is at least a $\frac{\delta^2}{4}$ fraction of the constraints between the layers X_{i_j} and $X_{i_{j'}}$.*

Theorem 4.2 *There exists a universal constant $\gamma > 0$, such that for any parameters ℓ, R , there is a weakly-dense ℓ -layered PCP $\Phi = \cup \Phi_{ij}$ such that it is NP-hard to distinguish between the following two cases:*

- **Yes:** *There exists an assignment that satisfies all the constraints.*
- **No:** *For every $i < j$, not more than $1/R^\gamma$ of the constraints in Φ_{ij} can be satisfied by an assignment.*

Moreover, the theorem holds even if every variable participates in $R^{O(\ell)}$ constraints.

Proof: Let Ψ be a constraint-system as in Theorem 2.6. We construct $\Phi = \cup \Phi_{ij}$ as follows. The variables X_i of layer $i \in [\ell]$ are the elements of the set $Z^i \times Y^{\ell-i}$, i.e., all ℓ -tuples where the first i elements are Z variables and the last $\ell - i$ elements are Y variables. The variables in layer i have assignments from the set $R_i = (R_Z)^i \times (R_Y)^{\ell-i}$ corresponding to an assignment to each variable of Ψ in the ℓ -tuple. It is easy to see that $|R_i| \leq R^{O(\ell)}$ for any $i \in [\ell]$ and that the total number of variables is no more than $|\Psi|^{O(\ell)}$. For any $1 \leq i < j \leq \ell$ we define the constraints in Φ_{ij} as follows. A constraint exists between a variable $x_i \in X_i$ and a variable $x_j \in X_j$ if they contain the same Ψ variables in the first i and the last $\ell - j$ elements of their ℓ -tuples. Moreover, for any $i < k \leq j$ there should be a constraint in Ψ between $x_{i,k}$ and $x_{j,k}$. More formally, denoting $x_i = (x_{i,1}, \dots, x_{i,\ell})$ for

$$x_i \in X_i = Z^i \times Y^{\ell-i},$$

$$\Phi_{ij} = \left\{ \begin{array}{l} \pi_{x_i, x_j} \mid x_i \in X_i, x_j \in X_j, \\ \forall k \in [\ell] \setminus \{i+1, \dots, j\}, x_{i,k} = x_{j,k} \\ \forall k \in \{i+1, \dots, j\}, \pi_{x_{i,k} \rightarrow x_{j,k}} \in \Psi \end{array} \right\}.$$

As promised, the constraints $\pi_{x_i \rightarrow x_j}$ are projections. Given an assignment $a = (a_1, \dots, a_\ell) \in R_i$ to x_i , we define the consistent assignment $b = (b_1, \dots, b_\ell) \in R_j$ to x_j as $b_k = \pi_{x_{i,k} \rightarrow x_{j,k}}(a_k)$ for $k \in \{i+1, \dots, j\}$ and $b_k = a_k$ for all other k .

How many constraints does a variable x participate in? Let $x \in X_i$ for some $i \in [\ell]$. Then $x = (x_1, \dots, x_\ell)$ has a constraint with $x' = (x'_1, \dots, x'_\ell) \in X_j$ if $i < j$ and on each coordinate $k \in \{i+1, \dots, j\}$, Ψ contains a constraint between x'_k and x_k . For each k there are at most $R^{O(1)}$ constraints in Ψ that touch x_k , so altogether there are $R^{O(j-i)}$ such constraints that touch x . This is similar for the case where $j < i$, and summing these together there are at most $\ell \cdot R^{O(\ell)} = R^{O(\ell)}$ constraints that touch any given x .

The completeness of Φ follows easily from the completeness of Ψ . That is, assume we are given an assignment $A : Y \cup Z \rightarrow R_Y \cup R_Z$ that satisfies all the constraints of Ψ . Then, the assignment $B : \bigcup X_i \rightarrow \bigcup R_i$ defined by $B(x_1 \dots x_\ell) = (A(x_1) \dots A(x_\ell))$ is a satisfying assignment.

For the soundness part, assume that there exist two layers $i < j$ and an assignment B that satisfies more than a $1/R^\gamma$ fraction of the constraints in Φ_{ij} . We partition X_i into classes such that two variables in X_i are in the same class iff they are identical except possibly on coordinate j . The variables in X_j are also partitioned according to coordinate j . Since more than $1/R^\gamma$ of the constraints in Φ_{ij} are satisfied, it must be the case that there exist a class $x_{i,1}, \dots, x_{i,j-1}, x_{i,j+1}, \dots, x_{i,\ell}$ in the partition of X_i and a class $x_{j,1}, \dots, x_{j,j-1}, x_{j,j+1}, \dots, x_{j,\ell}$ in the partition of X_j between which there exist constraints and the fraction of satisfied constraints is more than $1/R^\gamma$. We define an assignment to Ψ as

$$A(y) = (B(x_{i,1}, \dots, x_{i,j-1}, y, x_{i,j+1}, \dots, x_{i,\ell}))_j$$

for $y \in Y$ and as

$$A(z) = (B(x_{j,1}, \dots, x_{j,j-1}, z, x_{j,j+1}, \dots, x_{j,\ell}))_j$$

for $z \in Z$. Notice that there is a one-to-one and onto correspondence between the constraints in Ψ and the constraints between the two chosen classes in Φ . Moreover, if the constraint in Φ is satisfied, then the constraint in Ψ is also satisfied. Therefore, A is an assignment to Ψ that satisfies more than $1/R^\gamma$ of the constraints.

To prove that this multilayered PCP is *weakly-dense*, we recall the bi-regularity property mentioned above, i.e., each variable $y \in Y$ appears in the same number of constraints and also each $z \in Z$ appears in the same number of constraints. Therefore, the distribution obtained by uniformly choosing a variable $y \in Y$ and then uniformly choosing one of the variables in $z \in Z$ with which it has a constraint is a uniform distribution on Z .

Take any $m = \lceil \frac{2}{\delta} \rceil$ layers $i_1 < \dots < i_m$ and sets $S_j \subseteq X_{i_j}$ for $j \in [m]$ such that $|S_j| \geq \delta |X_{i_j}|$. Consider a random walk beginning from a uniformly chosen variable $x_1 \in X_1$ and proceeding to a variable $x_2 \in X_2$ chosen uniformly among the variables with which x_1 has a constraint. The random walk continues in a similar way to a variable $x_3 \in X_3$ chosen uniformly among the variables with which x_2 has a constraint and so on up to a variable in X_ℓ . Denote by E_j the indicator variable of the event that the random walk hits an S_j variable when in layer X_{i_j} . From the uniformity of Ψ it

follows that for every j , $\Pr[E_j] \geq \delta$. Moreover, using the inclusion-exclusion principle, we get:

$$\begin{aligned}
1 &\geq \Pr[\bigvee E_j] \geq \sum_j \Pr[E_j] - \sum_{j < k} \Pr[E_j \wedge E_k] \\
&\geq \lceil \frac{2}{\delta} \rceil \cdot \delta - \binom{m}{2} \max_{j < k} \Pr[E_j \wedge E_k] \\
&\geq 2 - \binom{m}{2} \max_{j < k} \Pr[E_j \wedge E_k]
\end{aligned}$$

which implies

$$\max_{j < k} \Pr[E_j \wedge E_k] \geq 1 / \binom{m}{2} \geq \frac{\delta^2}{4}$$

Fix j and k such that $\Pr[E_j \wedge E_k] \geq \frac{\delta^2}{4}$ and consider a shorter random walk beginning from a random variable in X_{i_j} and proceeding to the next layer and so on until hitting layer i_k . Since E_j is uniform on X_{i_j} we still have that $\Pr[E_j \wedge E_k] \geq \frac{\delta^2}{4}$ where the probability is taken over the random walks between X_{i_j} and X_{i_k} . Also, notice that there is a one-to-one and onto mapping from the set of all random walks between X_{i_j} and X_{i_k} to the set Φ_{i_j, i_k} . Therefore, at least a fraction $\frac{\delta^2}{4}$ of the constraints between X_{i_j} and X_{i_k} are between S_j and S_k , which completes the proof of the weak-density property. \blacksquare

5 A factor $(k - 1 - \varepsilon)$ inapproximability result

Theorem 1.1 (Main Theorem) *For any $k \geq 3$, it is NP-hard to approximate Ek -Vertex-Cover within any factor less than $k - 1$. More specifically, for every $\varepsilon, \delta > 0$, it is NP-hard to distinguish, given a k -uniform hypergraph G , between the following two cases*

- $IS(G) \geq 1 - \frac{1}{k-1} - \delta$
- $IS(G) \leq \varepsilon$.

Proof: Fix $k \geq 3$ and arbitrarily small $\varepsilon, \delta > 0$. Define $p = 1 - \frac{1}{k-1} - \delta$. Let Φ be a PCP instance with layers X_1, \dots, X_ℓ , as described in Theorem 4.2, with parameters $\ell = 33\varepsilon^{-2}$ and R large enough to be chosen later. We present a construction of a k -uniform hypergraph $G = (V, E)$. We use the Long Code introduced by Bellare et al. [3]. A Long Code over domain R has one bit for every subset $v \subseteq R$. An encoding of element $x \in R$ assigns bit-value 1 to the sets v s.t. $x \in v$ and assigns 0 to the sets which do not contain x . In the following, the bits in the Long Code will be vertices of the hypergraph. The vertices that correspond to a bit-value 0 are (supposedly) the vertices of a Vertex Cover.

VERTICES. For each variable x in layer X_i we construct a block of vertices $V[x]$. This block contains a vertex for each subset of R_i . Throughout this section we slightly abuse notation by writing a vertex rather than the set it represents. The weight of each vertex in the block $V[x]$ is set according to $\mu_p^{R_i}$, i.e. the weight of a subset $v \subseteq R_i$ is proportional to $\mu_p^{R_i}(v) = p^{|v|}(1-p)^{|R_i \setminus v|}$ as in Definition 2.2. All blocks in the same layer have the same total weight and the total weight of each layer is $\frac{1}{\ell}$. Formally, the weight of a vertex $v \in V[x]$ where $x \in X_i$ is given by

$$\frac{1}{\ell |X_i|} \mu_p^{R_i}(v).$$

HYPEREDGES. We construct hyperedges between blocks $V[x]$ and $V[y]$ such that there exists a constraint $\pi_{x \rightarrow y}$. We connect a hyperedge between any $v_1, \dots, v_{k-1} \in V[x]$ and $u \in V[y]$ whenever $\pi_{x \rightarrow y}(\bigcap_{i=1}^{k-1} v_i) \cap u = \emptyset$.

The intuition for the hyperedges comes from the proof of completeness. In fact, the hypergraph is constructed by first deciding how to map a satisfying assignment for Φ into a subset of G 's vertices. Then, we construct the hyperedges so as to ensure that such subsets are independent sets in G .

Let $IS(G)$ denote the weight of vertices contained in the largest independent set of the hypergraph G .

Lemma 5.1 (Completeness) *If Φ is satisfiable then $IS(G) \geq p = 1 - \frac{1}{k-1} - \delta$.*

Proof: Let A be a satisfying assignment for Φ , i.e., A maps each $i \in [\ell]$ and $x \in X_i$ to an assignment in R_i such that all the constraints are satisfied. Let $\mathcal{I} \subseteq V$ contain in the block $V[x]$ all the vertices that contain the assignment $A(x)$,

$$\mathcal{I} = \bigcup_x \{v \in V[x] \mid v \ni A(x)\}.$$

We claim that \mathcal{I} is an independent set. Take any v_1, \dots, v_{k-1} in $\mathcal{I} \cap V[x]$ and a vertex u in $\mathcal{I} \cap V[y]$. The vertices v_1, \dots, v_{k-1} intersect on $A(x)$ and therefore the projection of their intersection contains $\pi_{x \rightarrow y}(A(x)) = A(y)$. Since u is in $\mathcal{I} \cap V[y]$ it must contain $A(y)$. The proof is completed by noting that inside each block, the weight of the set of all vertices that contain a specific assignment is exactly p . ■

We now turn to the soundness of the construction.

Lemma 5.2 (Soundness) *If $IS(G) \geq \varepsilon$ then Φ is satisfiable.*

Proof: Let \mathcal{I} be an independent set of weight ε . We consider the set X' of all variables x for which the weight of $\mathcal{I} \cap V[x]$ in $V[x]$ is at least $\varepsilon/2$. A simple averaging argument shows that the weight of $\bigcup_{x \in X'} V[x]$ is at least $\frac{\varepsilon}{2}$. Another averaging argument shows that in at least $\frac{\varepsilon}{4} \ell > \frac{8}{\varepsilon}$ layers, X' contains at least $\frac{\varepsilon}{4}$ fraction of the variables. Using the weak-density property of the PCP (see Definition 4.1) with the choice $\delta = \varepsilon/4$, we conclude that there exist two layers X_i and X_j such that $\frac{\varepsilon^2}{64}$ fraction of the constraints between them are constraints between variables in X' . Let us denote by X the variables in $X_i \cap X'$ and by Y the variables in $X_j \cap X'$.

We first find, for each variable $x \in X$, a short list $B(x)$ of values that are ‘‘assigned’’ to it by \mathcal{I} . Indeed consider the vertices in $\mathcal{I} \cap V[x]$. Define t to be $O(\frac{1}{\delta^3} \log \frac{1}{\varepsilon})$ with some large enough constant. Then, since $t \geq t(\frac{\varepsilon}{2}, k-1, \delta)$, Lemma 2.3 implies that there exist $k-1$ vertices in $\mathcal{I} \cap V[x]$ that intersect in at most t assignments. We denote these vertices by $v_{x,1}, \dots, v_{x,k-1}$ and their intersection by $B(x)$, where $|B(x)| \leq t$.

In the following we define an assignment to the variables in X and Y such that many of the constraints between them are satisfied. For a variable $y \in Y$ we choose the assignment that is contained in the largest number of projections of $B(x)$:

$$A(y) = \max_{a \in R_Y} |\{x \in X \mid a \in \pi_{x \rightarrow y}(B(x))\}|.$$

For a variable $x \in X$ we simply choose $A(x)$ to be a random assignment from the set $B(x)$ (or some arbitrary assignment in case $B(x)$ is empty). We will prove that indeed $A(y)$ occurs in many of the projections of $B(x)$.

Claim 5.3 *Let $y \in Y$, and let $X(y)$ be the set of all variables $x \in X$ that have a constraint with y . Then,*

$$\Pr_{x \in X(y)} [A(y) \in \pi_{x \rightarrow y}(B(x))] > \varepsilon_1 \stackrel{\text{def}}{=} \frac{1}{t \log(\frac{\varepsilon}{4}) / \log(1 - 1/(k-1)^t)}.$$

Before we prove the claim let us see how it concludes the proof. Indeed, the fraction of constraints between Y and X satisfied by A is, in expectation, at least ε_1/t . This is because for every y , the fraction of $x \in X(y)$, for which $A(y) \in \pi_{x \rightarrow y}(B(x))$ is by Claim 5.3 at least ε_1 . For such x , the probability that $A(x)$ was randomly chosen so that $\pi_{x \rightarrow y}(A(x)) = A(y)$, is at least $\frac{1}{|B(x)|} \geq \frac{1}{t}$.

Thus, there must be some assignment A that attains the expectation, and satisfies at least ε_1/t of the constraints between X and Y . This amounts to $\frac{\varepsilon_1}{t} \cdot \frac{\varepsilon^2}{64}$ fraction of the total constraints, which would imply that Φ is satisfiable, as long as

$$\frac{1}{t \log(\frac{\varepsilon}{4}) / \log(1 - 1/(k-1)^t)} \cdot \frac{1}{t} \cdot \frac{\varepsilon^2}{64} \geq \frac{1}{R^\gamma}. \quad (6)$$

Proof: (of Claim 5.3) Let $y \in Y$, $x \in X(y)$. By definition, there are no hyperedges of the form $(v_{x,1}, \dots, v_{x,k-1}, u)$ for any vertex $u \in \mathcal{I} \cap V[y]$. In other words, every vertex $u \in \mathcal{I} \cap V[y]$ must intersect $\pi_{x \rightarrow y}(B(x))$. Let x_{i_1}, \dots, x_{i_N} be the variables in X such that the constraint $\pi_{x_{i_j} \rightarrow y}$ exists, and consider the collection of sets $A_j = \pi_{x_{i_j} \rightarrow y}(B(x_{i_j}))$ for $j = 1, \dots, N$. Clearly $|A_j| \leq t$, and let D denote the maximum number of disjoint sets inside this family. We will show that D cannot be too large by estimating the maximum possible weight of the vertices in $\mathcal{I} \cap V[y]$. Indeed the vertices in $\mathcal{I} \cap V[y]$ must intersect every A_j . The constraint imposed by one A_j in isolation reduces the maximum possible weight of the vertices in $\mathcal{I} \cap V[y]$ by a factor of $1 - (1-p)^{|A_j|} \leq 1 - (1-p)^t$. Moreover, for disjoint A_j 's, these constraints are independent, so the total weight becomes at most $(1 - (1-p)^t)^D$. Since the weight of $\mathcal{I} \cap V[y]$ is at least $\frac{\varepsilon}{4}$, we have

$$\frac{\varepsilon}{4} \leq (1 - (1-p)^t)^D < \left(1 - \frac{1}{(k-1)^t}\right)^D$$

so

$$D < \log(\frac{\varepsilon}{4}) / \log(1 - 1/(k-1)^t).$$

To finish, we need the following simple claim:

Claim 5.4 *Let A_1, \dots, A_N be a collection of N sets, each of size at most $T \geq 1$. If there are not more than D pairwise disjoint sets in this collection, then there is an element that is contained in at least $\frac{N}{TD}$ sets.*

Proof: Take any maximal collection of pairwise disjoint sets and let A' denote their union. Notice that $|A'| \leq TD$ and that any set not in this collection must contain one of the elements of A' . Hence, there must exist an element of A' that is contained in at least $1 + \frac{N-D}{TD} \geq N/TD$ sets. ■

Claim 5.4 implies that there exists an assignment for y that is contained in at least

$$\frac{N}{t \cdot \log(\frac{\varepsilon}{4}) / \log(1 - 1/(k-1)^t)}$$

of the A_j 's. ■

This completes the soundness proof. (Lemma 5.2) ■

Theorem 1.1 now follows from Lemmas 5.1 and 5.2, and our main multilayered PCP construction from Theorem 4.2, since the ratio between the sizes of the vertex cover in the yes and no cases is at least $\frac{1-\varepsilon}{1-p} = (1-\varepsilon)/(\frac{1}{k-1} + \delta)$ which can be made arbitrarily close to $k-1$. ■

As before, let us compute some parameters of the reduction, which will be useful in the next section. We start by choosing R so that Equation (6) is satisfied. Since $\frac{1}{(k-1)^t}$ is small, we estimate $\log(1-x) \approx -x$ and get

$$R \geq \Omega \left(\frac{1}{\varepsilon^2} \log \frac{1}{\varepsilon} \cdot t^2 \cdot (k-1)^t \right)^{1/\gamma}. \quad (7)$$

By plugging in the value for t , we obtain that choosing

$$R = 2^{O(\delta^{-3} \log k \log \frac{1}{\varepsilon})}$$

suffices. The number of variables in the multilayered PCP is $n^{O(\ell \log R)}$ which, using $\ell = O(1/\varepsilon^2)$, is at most

$$n^{(\frac{\log k}{\varepsilon \delta})^{O(1)}}.$$

Hence, the number of vertices in the hypergraph is at most

$$n^{(\frac{\log k}{\varepsilon \delta})^{O(1)}} \cdot 2^{R^{O(\ell)}} \leq n^{(\frac{\log k}{\varepsilon \delta})^{O(1)}} \cdot 2^{2^{(\frac{\log k}{\varepsilon \delta})^{O(1)}}}.$$

Since the number of constraints each variable participates in is at most $R^{O(\ell)}$, the degree of the hypergraph is at most

$$R^{O(\ell)} \cdot 2^{O(kR^{O(\ell)})} \leq 2^{2^{(\frac{\log k}{\varepsilon \delta})^{O(1)}}}.$$

Finally, the running time of the reduction is polynomial in the number of hyperedges which is at most

$$2^{2^{(\frac{\log k}{\varepsilon \delta})^{O(1)}}} \cdot n^{(\frac{\log k}{\varepsilon \delta})^{O(1)}}. \quad (8)$$

6 Non-constant values of k

In this section we outline what happens to our construction when we increase k beyond a constant. This should be contrasted with the general hypergraph vertex cover problem in which k is unbounded (i.e., the set cover problem). It is well known that, if N is the number of hyperedges, the greedy algorithm achieves a factor $\ln N$ approximation, and no polynomial time algorithm can achieve approximation factor $(1 - o(1)) \ln N$, unless $\text{NP} \subseteq \text{DTIME}(n^{O(\log \log n)})$ [10].

Below, we discuss in turn the extension of both the $k-1$ and $k/2$ hardness bounds from Sections 5 and 3 respectively to the case when k is super-constant. We note that in this context the $k/2$ result is not subsumed by the $(k-1)$ result since it works for much larger values of k (as a function of the number of hyperedges in the hypergraph). We discuss the $(k-1)$ result first since this reduction and the calculations at the end of Section 5 are hopefully fresh in the reader's mind at this point of reading.

6.1 The $(k-1-\varepsilon)$ inapproximability bound

Theorem 6.1 *There exists some $c > 0$ such that unless $\text{NP} \subseteq \text{DTIME}(n^{O(\log \log n)})$, there is no polynomial time algorithm for approximating Ek -Vertex-Cover for $k \leq (\log \log N)^{1/c}$ to within a factor of $k-1.01$, where N is the number of hyperedges in the k -uniform hypergraph.*

Proof: Recall the calculation at the end of Section 5, specifically the bound on number of hyperedges from Equation (8). The produced hypergraph G has

$$N = 2^{2^{(\frac{\log k}{\varepsilon\delta})^{O(1)}}} \cdot n^{(\frac{\log k}{\varepsilon\delta})^{O(1)}}$$

hyperedges, when starting from a 3SAT instance of size n . Furthermore, the reduction runs in time polynomial in the number of hyperedges.

The gap in vertex cover sizes equals

$$\frac{1 - \varepsilon}{1/(k-1) + \delta} \geq (k-1)(1-\varepsilon)(1-k\delta) \geq (k-1.01)$$

if we set $\varepsilon = 1/(200k)$ and $\delta = 1/(200k^2)$. For these choices of ε, δ , the number of hyperedges N is at least $2^{2^{k^{O(1)}}}$, so the largest value of k we can set is $k = (\log \log n)^{1/c'}$ for a large enough constant c' . For this choice of k , the number of hyperedges N and the runtime of the reduction can both be bounded above by $n^{O(\log \log n)}$. For large enough n , $\log \log N \leq 2 \log \log n$ for this choice of parameters, so we get a result that works for $k \leq (\log \log N)^{1/c}$ for some fixed constant c . ■

6.2 The $(k/2 - \varepsilon)$ inapproximability bound

We can obtain a similar result corresponding to the factor $k/2$ hardness result, as stated below. Note that this result works for k up to $(\log N)^{1/b}$ for some fixed constant $b > 0$.

Theorem 6.2 *There exists some $b > 0$ such that unless $\text{NP} \subseteq \text{DTIME}(n^{O(\log \log n)})$, there is no polynomial time algorithm for approximating Ek -Vertex-Cover for $k \leq (\log N)^{1/b}$ to within a factor of $\lfloor \frac{k}{2} \rfloor - 0.01$, where N is the number of hyperedges in the k -uniform hypergraph.*

Proof: Recall the calculation at the end of Section 3, specifically Equations (4) and (5). The produced hypergraph G has $n_0 \leq n^{O(\log(1/\varepsilon\delta))} 2^{(1/\varepsilon\delta)^{O(1)}}$ vertices and $N = n_0 2^{k/(\varepsilon\delta)^{O(1)}}$ hyperedges, when starting from a 3SAT instance of size n . Furthermore, the reduction runs in time polynomial in the number of hyperedges.

By plugging in $k = (\log n)^{1/b'}$ for a large enough constant $b' > 0$, $\varepsilon = 1/(100k)$ and $\delta = 1/(100k^2)$, we get $\log(1/\varepsilon\delta) \leq \log \log n$, and $2^{k/(\varepsilon\delta)^{O(1)}} \leq n$. Thus, the reduction runs in time $n^{O(\log \log n)}$, produces a hypergraph with $N = n^{O(\log \log n)}$ hyperedges, with a gap in vertex cover sizes (assuming k is even for convenience) of

$$\frac{1 - \varepsilon}{2/k + \delta} \geq \frac{k}{2}(1 - \varepsilon)(1 - k\delta) \geq \frac{k}{2} - \frac{1}{100}$$

by our choices of ε, δ . Since for large enough n , $\log n \geq \sqrt{\log N}$ for the above choice of parameters, we get a result that works for $k \leq (\log N)^{1/b}$ for some fixed constant b . ■

7 Acknowledgements

We would like to thank Noga Alon for his help with s -wise t -intersecting families. We would also like to thank Luca Trevisan and the anonymous referees for useful comments.

References

- [1] 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.
- [2] S. Arora and S. Safra. Probabilistic checking of proofs : A new characterization of NP. *Journal of the ACM*, 45(1) : 70-122, 1998.
- [3] M. Bellare, O. Goldreich and M. Sudan. Free bits, PCPs and non-approximability. *SIAM Journal on Computing*, 27(3):804-915, June 1998.
- [4] H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Ann. Math. Statistics*, 23:493–507, 1952.
- [5] J. Chuzhoy, S. Guha, E. Halperin, S. Khanna, G. Kortsarz and J. Naor. Asymmetric K-Center is $\log^* n$ Hard to Approximate. *Proceedings of the 36th ACM Symposium on Theory of Computing (STOC)*, 2004.
- [6] P. Crescenzi, R. Silvestri, and L. Trevisan. On weighted vs unweighted versions of combinatorial optimization problems. *Inform. and Comput.*, 167(1):10–26, 2001. ISTCS'96 (Jerusalem).
- [7] I. Dinur, V. Guruswami and S. Khot. Vertex cover on k -uniform hypergraphs is hard to approximate within factor $(k - 3 - \epsilon)$. *Electronic Colloquium on Computational Complexity, Technical Report TR02-027*, 2002.
- [8] I. Dinur, O. Regev and C. Smyth. The hardness of 3-uniform hypergraph coloring. *In Proc. of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 33-42.
- [9] I. Dinur and S. Safra. The importance of being biased. *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 33-42, May 2002.
- [10] U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4): 634-652 (1998).
- [11] U. Feige, S. Goldwasser, L. Lovász, S. Safra and M. Szegedy. Interactive proofs and the hardness of approximating cliques. *Journal of the ACM*, 43(2):268-292, March 1996.
- [12] O. Goldreich. Using the FGLSS-reduction to prove inapproximability results for minimum vertex cover in hypergraphs. *ECCC Technical Report TR01-102*, December 2001.
- [13] V. Guruswami, J. Håstad and M. Sudan. Hardness of approximate hypergraph coloring. *SIAM J. Comput.*, 31(6): 1663-1686 (2002). Preliminary version in *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 149-158, November 2000.
- [14] E. Halperin. Improved approximation algorithms for the vertex cover problem in graphs and hypergraphs. *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 329-337, January 2000.
- [15] R. L. Graham, M. Grötschel, and L. Lovász, editors. *Handbook of combinatorics. Vol. 1, 2*. Elsevier Science B.V., Amsterdam, 1995.

- [16] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182(1):105-142, 1999.
- [17] J. Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798-859, July 2001.
- [18] J. Holmerin. Vertex cover on 4-uniform hypergraphs is hard to approximate within $2 - \epsilon$. *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC)*, pages 544-552, May 2002.
- [19] J. Holmerin. Improved inapproximability results for vertex cover on k-uniform hypergraphs. *Proc. of the 29th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 1005-1016, July 2002.
- [20] D. S. Johnson. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9:256–278, 1974.
- [21] S. Khot. Hardness results for coloring 3-colorable 3-uniform hypergraphs. In *Proc. of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 23-32, November 2002.
- [22] S. Khot and O. Regev. Vertex cover might be hard to approximate to within $2 - \epsilon$. In *Proc. of 18th IEEE Annual Conference on Computational Complexity (CCC)*, pages 379–386, 2003.
- [23] L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13:383–390, 1975.
- [24] R. Raz. A parallel repetition theorem. *SIAM J. of Computing*, vol 27(3), 763-803, 1998.
- [25] R. Raz and S. Safra. A Sub-Constant Error-Probability Low-Degree Test, and a Sub-Constant Error-Probability PCP Characterization of NP. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 475-484, 1997.
- [26] A. Samorodnitsky and L. Trevisan. A PCP characterization of NP with optimal amortized query complexity. In *Proc. of the 32nd Annual ACM Symposium on Theory of Computing*, pages 191-199, 2000.
- [27] L. Trevisan. Non-approximability results for optimization problems on bounded degree instances. In *Proc. of the Annual ACM 33rd Symposium on Theory of Computing*, pages 453-461, July 2001.