

Using the FGLSS-reduction to Prove Inapproximability Results for Minimum Vertex Cover in Hypergraphs

Oded Goldreich

Abstract. Using known results regarding PCP, we present simple proofs of the inapproximability of vertex cover for hypergraphs. Specifically, we show that

1. Approximating the size of the minimum vertex cover in $O(1)$ -regular hypergraphs to within a factor of 1.99999 is NP-hard.
2. Approximating the size of the minimum vertex cover in 4-regular hypergraphs to within a factor of 1.49999 is NP-hard.

Both results are inferior to known results (by Trevisan (2001) and Holmerin (2001)), but they are derived using much simpler proofs. Furthermore, these proofs demonstrate the applicability of the FGLSS-reduction in the context of reductions among combinatorial optimization problems.

Keywords: Complexity of approximation, combinatorial optimization problems, Vertex Cover, PCP, regular hypergraphs.

An early version of this work appeared as TR01-102 of *ECCC*. A discussion of subsequent works is deferred to Section 5.

1 Introduction

This note was inspired by a work of Dinur and Safra [5], which was new at the time this work was completed. Specifically, what we take from their work is the realization that the so-called FGLSS-reduction is actually a general paradigm that can be applied in various ways and achieve various purposes.

The FGLSS-reduction, introduced by Feige, Goldwasser, Lovász, Safra and Szegedy [7], is typically understood as a reduction from sets having certain PCP systems to approximation versions of Max-Clique (or Max Independent Set). The reduction maps inputs (either in or out of the set) to graphs that represent the pairwise consistencies among possible views of the corresponding PCP verifier. It is instructive to think of these possible verifier views as of possible partial solutions to the problem of finding an oracle that makes the verifier accept.

Dinur and Safra apply the same underlying reasoning to derive graphs that represent pairwise consistencies between partial solutions to other combinatorial problems [5]. In fact, they use two different instantiations of this reasoning. Specifically, in one case they start with the vertex-cover problem and consider the restrictions of possible vertex-covers to all possible $O(1)$ -subsets of the vertex

set. The partial solutions in this case are the vertex-covers of the subgraphs induced by all possible $O(1)$ -subsets, and pairwise consistency is defined in the natural way. Thus, we claim that in a sense, the work of Dinur and Safra [5] suggests that the FGLSS-reduction is actually a general paradigm that can be instantiated in various ways. Furthermore, the goal of applying this paradigm may vary too. In particular, the original instantiation of the FGLSS-reduction by Feige *et. al.* [7] was aimed at linking the class PCP to the complexity of approximating combinatorial optimization problems. In contrast, in the work of Dinur and Safra [5] one instantiation is aimed at deriving instances of very low “degree” (i.e., co-degree at most 2), and the other instantiation is aimed at moving the “gap location” (cf. [16] and further discussion below).

We fear that the complexity of the work of Dinur and Safra [5] may cause researchers to miss the foregoing observation (regarding the wide applicability of the FGLSS-reduction). This would be unfortunate, because we believe in the potential of that observation. In fact, this note grew out of our fascination with the foregoing observation and our attempt to find a simple illustration of it.

Our concrete results: Combining known results regarding PCP with the FGLSS-reduction, we present simple proofs of inapproximability results regarding the *minimum vertex cover problem for hypergraphs*. Specifically, we show that:

1. For every constant $\epsilon > 0$, approximating the size of the minimum vertex cover in $O(1)$ -regular hypergraphs to within a $(2 - \epsilon)$ -factor is NP-hard (see Section 3). In fact, the hypergraphs we use are $O((1/\epsilon)^{o(1)})$ -regular.

This result is inferior to Holmerin’s result [12], by which approximating vertex cover in 4-regular hypergraphs to within a $(2 - \epsilon)$ -factor is NP-hard. We also mention Trevisan’s result [17] by which, for every constant k , approximating vertex cover in k -regular hypergraphs to within a $\Omega(k^{1/19})$ -factor is NP-hard. Clearly, in terms of achieving a bigger inapproximation factor, Trevisan’s result is superior, but in terms of achieving an inapproximation result for k -regular graphs when k is small (e.g., $k < 2^{19}$) it seems that our result is better.

2. For every constant $\epsilon > 0$, approximating the size of the minimum vertex cover in 4-regular hypergraphs to within a $(1.5 - \epsilon)$ -factor is NP-hard (see Section 4).

Again, this result is inferior to Holmerin’s result [12].

We mention that our work was done independently of Holmerin’s work [12], but after the publication of Trevisan’s work [17].

2 Preliminaries

This section contains a review of the notion of a vertex cover in a hypergraph and the notion of free-bit complexity. We also recall the FGLSS-reduction and discuss its relation to the vertex cover problem in graphs.

Vertex covers in hypergraphs. A k -regular hypergraph is a pair (V, E) such that E is a collection of k -subsets (called hyper-edges) of V ; that is, for every $e \in E$ it holds that $e \subseteq V$ and $|e| = k$. For a k -regular hypergraph $H = (V, E)$ and $C \subseteq V$, we say that C is a vertex cover of H if for every $e \in E$ it holds that $e \cap C \neq \emptyset$.

Free-bit complexity and the class \mathcal{FPCP} . We assume that the reader is familiar with the basic PCP-terminology (cf. [1–3] and [8, Sec. 2.4]). (For sake of simplicity we consider non-adaptive verifiers.) We say that the free-bit complexity of a PCP system is bounded by $f : \mathbb{N} \rightarrow \mathbb{R}$ if on every input x and any possible random-pad ω used by the verifier, there are at most $2^{f(|x|)}$ possible sequence of answers that the verifier may accept (on input x and random-pad ω). Clearly, the free-bit complexity of a PCP system is bounded by the number of queries it makes, but the former may be much lower. Free-bit complexity is a key parameter in the FGLSS-reduction. For functions $c, s : \mathbb{N} \rightarrow [0, 1]$, $r : \mathbb{N} \rightarrow \mathbb{N}$ and $f : \mathbb{N} \rightarrow \mathbb{R}$, we denote by $\mathcal{FPCP}_{c,s}[r, f]$ the class of sets having PCP systems of completeness bound c , soundness bound s , randomness complexity r and free-bit complexity f . In particular, for every input x in the set, there exist an oracle that makes the verifier accept with probability at least $c(|x|)$, whereas for every input x not in the set and every oracle the verifier accepts with probability at most $s(|x|)$.

The FGLSS-graph. For $S \in \mathcal{FPCP}_{c,s}[r, f]$, the FGLSS-reduction maps x to a graph G_x having $2^{r(|x|)}$ layers, each having at most $2^{f(|x|)}$ vertices. The vertices represent possible views of the verifier, where the $N \stackrel{\text{def}}{=} 2^{r(|x|)}$ layers correspond to all possible choices of the random-tape and the vertices in each layer correspond to the up-to $2^{f(|x|)}$ possible sequences of answers that the verifier may accept. The edges represent *inconsistencies* among these views. In particular, each layer consists of a clique (because only one sequence of answers is possible for a fixed random-tape and a fixed oracle). If the random-tapes $\omega_1, \omega_2 \in \{0, 1\}^{r(|x|)}$ both lead the verifier to make the same query q (and both answers are acceptable), then the corresponding layers will have edges between vertices encoding views in which different answers are given to query q . In the case that $x \in S$ the graph G_x will have an independent set of size $c(|x|) \cdot N$, whereas in the case that $x \notin S$ the maximum independent set in G_x has size at most $s(|x|) \cdot N$. Thus, the inapproximability factor for the maximum independent set problem shown by such a reduction is $c(|x|)/s(|x|)$, and the fact the maximum independent set is always at most a $2^{-f(|x|)}$ fraction of the size of G_x does not effect the gap. However, inapproximability factor for the minimum vertex cover shown by such a reduction is

$$\frac{2^{f(|x|)} \cdot N - s(|x|) \cdot N}{2^{f(|x|)} \cdot N - c(|x|) \cdot N} = \frac{2^{f(|x|)} - s(|x|)}{2^{f(|x|)} - c(|x|)} < \frac{2^{f(|x|)}}{2^{f(|x|)} - 1}. \quad (1)$$

This is the reason that, while the FGLSS-reduction allows to establish quite optimal inapproximability factors for the maximum independent set problem,

it failed so far to establish optimal inapproximability factors for the minimum vertex cover problem (although, it was used by Hastad [10] in deriving the $7/6$ hardness factor by using Eq. (1) with $f = 2$, $c \approx 1$ and $s = 1/2$). In a sense, the gap between the size of the maximum independent set of G_x for $x \in S$ versus for $x \notin S$ is at the “right” location for establishing inapproximability factors for the maximum independent set problem, but is at the “wrong” location for establishing inapproximability factors for the minimum vertex cover problem. Thus, what we do below is “move the gap location”: Specifically, in Section 3, we take a maximum independent set gap of $c2^{-f}$ versus $s2^{-f}$ (which means a minimum vertex cover gap of $1 - c2^{-f}$ versus $1 - s2^{-f}$), and transform it into a minimum vertex cover gap of $(2 - c) \cdot 2^{-f}$ versus $(2 - s) \cdot 2^{-f}$.

3 A $2 - \epsilon$ Hardness Factor for $O(1)$ -Regular Hypergraphs

We start with the usual FGLSS-graph, denoted G , derived from the FGLSS-reduction as applied to input x of a $\mathcal{FPCP}_{1-\epsilon,s}[\log, f]$ scheme (for a set in \mathcal{NP}). For simplicity, think of f as being a constant such that 2^f is an integer. Without loss of generality, each layer of G has $\ell = 2^f$ vertices.

We now apply the “FGLSS paradigm” by considering vertex-covers of G , and their projection on each layer. Such projections (or “partial assignments”) have either ℓ or $\ell - 1$ vertices. We focus on the good vertex covers, having exactly $\ell - 1$ vertices in (almost) each layer. Thus, for each $(\ell - 1)$ -subset of each layer, we introduce a vertex in the hypergraph, to be denoted H . We also introduce hyper-edges so to reflect the inconsistencies of the various partial (i.e. layer-projected) vertex covers of G . This construction, presented next, will provide a correspondance between vertex covers of G and vertex covers of H .

The construction of the hypergraph H . For each layer $L = (v_1, \dots, v_\ell)$ in G , we introduced a corresponding layer in H containing ℓ vertices such that each H -vertex corresponds to an $(\ell - 1)$ -subset of L ; that is, we introduce ℓ vertices that correspond to $L \setminus \{v_1\}, \dots, L \setminus \{v_\ell\}$. For each pair of layers $L' = (v'_1, \dots, v'_\ell)$ and $L'' = (v''_1, \dots, v''_\ell)$, if (v'_i, v''_j) is an edge in G , then we introduce the $2 \cdot (\ell - 1)$ -hyperedge containing the H -vertices that correspond to the subsets $\{L' \setminus \{v'_k\} : k \neq i\}$ and $\{L'' \setminus \{v''_k\} : k \neq j\}$; that is, the hyper-edge consists of all the H -vertices of these two layers except for the two H -vertices that correspond to the subsets $L' \setminus \{v'_i\}$ and $L'' \setminus \{v''_j\}$. In addition, for each layer in H , we introduce an ℓ -size hyper-edge containing all ℓ vertices of that layer.

To get rid of the non-regularity of this construction, we augment each layer with a sets of $\ell - 2$ auxiliary vertices, and replace the abovementioned ℓ -size hyper-edge by a hyper-edge containing all vertices of that layer (i.e., the original ℓ vertices as well as the $\ell - 2$ auxiliary vertices). We refer to these hyper-edges as intra-layer ones. This completes the construction of H .

Motivation to the analysis. Consider a generic vertex cover, C , of G , and let S denote the set of all vertices of H that correspond to the $(\ell - 1)$ -subsets of

C . Note that C contains ℓ vertices of some layer of G if and only if S contains all vertices of the corresponding layer in H , and in this case all edges (resp., hyper-edges) adjacent to this layer are covered. Thus, we focus on layers of G that contain $\ell - 1$ vertices of C , and note that (for each such layer) S contains a single vertex of H that resides in the corresponding layer. Let $L' = (v'_1, \dots, v'_\ell)$ and $L'' = (v''_1, \dots, v''_\ell)$ be two such layers of G , and let v'_i and v''_j denote the two vertices that are missing from C (which implies that (v'_i, v''_j) is not an edge in G). Then, $L' \setminus \{v'_i\}$ and $L'' \setminus \{v''_j\}$ are in S , and they cover all the hyper-edges that connect L' and L'' , because $\{L' \setminus \{v'_k\} : k \neq i\} \cup \{L'' \setminus \{v''_k\} : k \neq j\}$ is not a hyper-edge in H .

The actual analysis. Fixing any input x , we consider the corresponding FGLSS-graph $G = G_x$, and the hypergraph $H = H_x$ derived from G by following the above construction. Let N denote the number of layers in G (and H).

Claim 3.1 *If x is a yes-instance, then the hypergraph H_x has a vertex-cover of size at most $(1 + \epsilon) \cdot N$.*

Proof: Since x is a yes-instance, the graph $G = G_x$ has an independent set (IS) of size at least $(1 - \epsilon) \cdot N$. Consider this IS or actually the corresponding vertex-cover (i.e., VC) of G . Call a layer in G good if it has $\ell - 1$ vertices in this VC, and note that at least $(1 - \epsilon) \cdot N$ layers are good. We create a vertex-cover for $H = H_x$ as follows. For each good layer, place in C the corresponding H -vertex; that is, the H -vertex corresponding to the $(\ell - 1)$ -subset (of this layer in G) that is in the VC of G . For the rest of the layers (i.e., the non-good layers), place in C any two H -vertices of each (non-good) layer.

In total we placed in C at most $(1 - \epsilon)N + 2\epsilon N = (1 + \epsilon)N$ vertices. We show that C is a vertex cover of H by considering all possible hyper-edges, bearing in mind the correspondence between layers of G and layers of H .

- Each intra-layer hyper-edge of H (which consists of all vertices of that layer) is definitely covered, because we placed in C at least one H -vertex from each layer.
- Each hyper-edge connecting H -vertices from two good layers is covered. This is shown by considering the edge, denoted (u, v) , of G that is “responsible” for the introduction of each hyper-edge (in H).¹ Since we started with a vertex cover of G , either u or v must be in that cover. Suppose, without loss of generality, that u is in the VC of G . Then, we must have placed in C one of the H -vertices that corresponds to a $(\ell - 1)$ -subset that contains u . But, then, this H -vertex covers the said hyper-edge (because, by construction, the latter hyper-edge contains all $(\ell - 1)$ -subsets that contain u).
- Each hyper-edge that contains H -vertices from at least one non-good layer is covered, because we placed in C two H -vertices from each non-good layer,

¹ A hyper-edge that corresponds to layers L' and L'' has the form $\{L' \setminus \{w\} : w \neq u\} \cup \{L'' \setminus \{w\} : w \neq v\}$, where $u \in L'$ and $v \in L''$. Furthermore, (u, v) must be an edge in G .

whereas each hyper-edge containing H -vertices of some layer contains all but at most one vertex of that layer.

The claim follows. \blacksquare

Claim 3.2 *If x is a no-instance, then every vertex-cover of the hypergraph H_x has size at least $(2 - s(|x|)) \cdot N$.*

Proof: Consider any vertex cover C of H . Note that due to the intra-layer hyper-edges, C must contain at least one vertex in each layer. Furthermore, without loss of generality, C contains only original vertices (rather than the $\ell-2$ auxiliary vertices added to each layer). Denote by C' the set of layers that have a *single* vertex in C . Then, $|C| \geq |C'| + 2(N - |C'|) = 2N - |C'|$. The claim follows by proving that $|C'| \leq sN$, where $s \stackrel{\text{def}}{=} s(|x|)$.

Suppose, towards the contradiction, that $|C'| > sN$. We consider the set of G -vertices, denoted I , that correspond to the (single) H -vertices in these layers; that is, for layer L (in C') such that C contains the H -vertex (which corresponds to) $L \setminus \{v\}$, place $v \in G$ in I . We show that I is an independent set in G (and so derive a contradiction to $G = G_x$ not having an independent set of size greater than sN , because x is a no-instance). Specifically, for every $u, v \in I$, we show that (u, v) cannot be an edge in G . Suppose (u, v) is an edge in G , then the corresponding hyper-edge in H cannot be covered by C ; that is, the hyper-edge $\{L \setminus \{w\} : w \neq u\} \cup \{L' \setminus \{w\} : w \neq v\}$ (which must be introduced due to the edge (u, v)) cannot be covered by the H -vertices that correspond to the $(\ell-1)$ -subsets $L \setminus \{u\}$ and $L' \setminus \{v\}$. The claim follows. \blacksquare

Conclusion: Starting from a $\mathcal{FPCP}_{1-\epsilon, s}[\log, f]$ system for \mathcal{NP} , we have shown that the minimum vertex-cover in $(2^{f+1} - 2)$ -regular hypergraphs is NP-hard to approximate to a $(2-s)/(1+\epsilon)$ -factor. Now, if we start with any $\mathcal{FPCP}_{1, s}[\log, f]$ for \mathcal{NP} , with $s \approx 0$, then we get a hardness result for a factor of $2 - s \approx 2$. Any $\mathcal{NP} \subseteq \mathcal{PCP}[\log, O(1)]$ result (starting from [1]) will do for this purpose, because a straightforward error-reduction will yield $\mathcal{NP} \subseteq \mathcal{FPCP}_{1, s}[\log, O(1)]$, for any $s > 0$. The (amortized) free-bit complexity only effects the growth of the hyper-edge size as a function of the deviation of the hardness-factor from 2. Specifically, if we start with an “amortized free-bit complexity zero” result (i.e., $\mathcal{NP} \subseteq \mathcal{FPCP}_{1, s}[\log, o(\log_2(1/s))]$ for every $s > 0$), then we get a factor of $2 - s$ hardness for $(1/s)^{o(1)}$ -regular hypergraphs. That is, starting with Hastad’s first such result [9] (or from the simplest one currently known [11]), and applying the foregoing reasoning, we obtain our first little result:

Theorem 3.3 *For every constant $\epsilon > 0$, approximating the size of the minimum vertex cover in $(1/\epsilon)^{o(1)}$ -regular hypergraphs to within a $(2-\epsilon)$ -factor is NP-hard.*

Alternatively, if we start with Hastad’s “maxLIN3 result” [10] (i.e., the result $\mathcal{NP} \subseteq \mathcal{FPCP}_{1-\epsilon, 0.5}[\log, 2]$ for every $\epsilon > 0$), then we get a hardness factor of $(2 - 0.5)/(1 + \epsilon) \approx 1.5$ for 6-regular hypergraphs. Below we show that the same hardness factor holds also for 4-regular hypergraphs (by starting with the same “maxLIN3 result” [10] but capitalizing on an additional feature of it).

4 A 1.5 – ϵ Hardness Factor for 4-Regular Hypergraphs

We start with the FGLSS-graph derived from applying the FGLSS-reduction to Hastad’s “maxLIN3 system” [10]; that is, the $\mathcal{FPCP}_{1-\epsilon,0.5}[\log, 2]$ system for \mathcal{NP} ($\forall \epsilon > 0$). The key observation is that, *in this system*, for any two queries, all four answer pairs are possible (as accepting configurations).² This observation is relied upon when establishing (below) simple structural properties of the derived FGLSS-graph.

As before, there will be a correspondence between the vertex set of G and the vertex set of H . Here it is actually simpler to just identify the two sets. So it just remains to specify the hyper-edges of H . Again, we place (intra-layer) hyper-edges between all (i.e., four) vertices of each layer. As for the construction of inter-layer hyper-edges, we consider three cases regarding each pair of layers:

1. *The trivial case*: If there are no edges between these two layers in G , then there would be no hyper-edges between these layers in H . This case corresponds to the case that these two layers correspond to two random-tapes that induce two query sets with empty intersection.
2. *The interesting case* is when these two layers correspond to two random-tapes that induce two query sets having a *single* query, denoted q , in common. Relying on the property of the starting PCP system, it follows that both answers are possible to this query and that each possible answer is represented by two vertices in each corresponding layer. Accordingly, we denote the vertices of the first layer by $u_1^0, u_2^0, u_1^1, u_2^1$, where u_i^b is the i th configuration in this layer in which query q is answered by the bit b . Similarly, denote the vertices of the second layer by $v_1^0, v_2^0, v_1^1, v_2^1$. (We stress that this notation is used only for determining the hyper-edges between the current pair of layers, and when considering a different pair of layers a different notation may be applicable.) In this case we introduce the two hyper-edges $\{u_1^0, u_2^0, v_1^1, v_2^1\}$ and $\{u_1^1, u_2^1, v_1^0, v_2^0\}$.

Intuition: Note that the edges in G [*sic*] between these two layers are two $K_{2,2}$ ’s (i.e., for each $b \in \{0, 1\}$, between the two u_i^b ’s on one side and the two v_i^{1-b} ’s on the other side). These two $K_{2,2}$ ’s enforce that if some u_i^b is in some IS, then v_j^{1-b} is not in the IS. For a H -VC having a single vertex in each layer, the (two) hyper-edges will have the same effect.

3. *The annoying case* is when these two layers (correspond to two random-tapes that induce two query sets that) have *two* or more queries in common. In this case, we label the vertices in these two layers according to these two answers; that is, we denote the four vertices of the first layer by $u_{0,0}, u_{0,1}, u_{1,0}, u_{1,1}$, where $u_{a,b}$ is the unique configuration in this layer in which these two queries are answered by a and b , respectively. Similarly, denote the vertices of the second layer by $v_{0,0}, v_{0,1}, v_{1,0}, v_{1,1}$. (Again, this notation is used only for determining the hyper-edges between the current pair of layers.) In this case,

² Recall that the number of queries is typically higher than the free-bit complexity. Indeed, the aforementioned system makes *three queries* and has free-bit complexity *two*.

we introduce four hyper-edges between these two layers, each has one vertex of the first layer and the three “non-matching” vertices of the second layer; that is, the hyper-edges are $\{u_{a,b}, v_{a,1-b}, v_{1-a,b}, v_{1-a,1-b}\}$, for $a, b \in \{0, 1\}$. Intuition: The pair $(u_{a,b}, v_{a',b'})$ is an edge in G if and only if either $a \neq a'$ or $b \neq b'$. Similarly, the pair $(u_{a,b}, v_{a',b'})$ participates in an hyper-edge of H if and only if either $a \neq a'$ or $b \neq b'$.

This completes the construction. Note that $H = H_x$ is a 4-regular hypergraph.

Claim 4.1 *If x is a yes-instance, then the hypergraph H_x has a vertex-cover of size at most $(1 + 3\epsilon) \cdot N$, where N denotes the number of layers.*

Proof: Since x is a yes-instance, the graph $G = G_x$ has an independent set (IS) of size $(1 - \epsilon)N$. Consider such an IS, denoted I . Call a layer in G **good** if it has a vertex in I , and note that at least $(1 - \epsilon)N$ layers are good. Augment I by the set of all vertices residing in non-good layers. In total we took at most $(1 - \epsilon)N + 4\epsilon N = (1 + 3\epsilon)N$ vertices. We show that these vertices cover all hyper-edges of H .

- The intra-layer hyper-edges are definitely covered (since we took at least one vertex from each layer).
- Each hyper-edge connecting vertices from two good layers is covered. This is shown by considering each of the two non-trivial cases (in the construction). In the interesting case, I (having a single vertex in each good layer) must have a single vertex in *each* $K_{2,2}$. But then this vertex covers the corresponding hyper-edge. In the annoying case, I (having a single vertex in each good layer) must contain vertices with matching labels in these two layers. But then these two vertices cover all four hyper-edges, because each hyper-edge contains a (single) vertex of each label.
- Hyper-edges containing H -vertices from non-good layers are covered trivially (because we took all vertices of each non-good layer).

The claim follows. ■

Claim 4.2 *If x is a no-instance, then every vertex-cover of the hypergraph H_x has size at least $1.5 \cdot N$.*

Proof: Consider a cover C of H . Note that (due to the intra-layer hyper-edges) C must contain at least one vertex in each layer. Denote by C' the set of layers that have a single vertex in C . Then, $|C| \geq |C'| + 2(N - |C'|)$. The claim follows by proving that $|C'| \leq 0.5N$. Suppose, towards the contradiction, that $|C'| > 0.5N$. Consider the set of vertices, denoted I , that correspond to these layers (i.e., for a layer in C' consider the layer's vertex that is in C). We show that I is an independent set in G (and so we derive contradiction).

Suppose (towards the contradiction) that $u, v \in I$ and (u, v) is an edge in G . In the interesting case, this (i.e., (u, v) being an edge in G) means that u and v are in the same hyper-edge in H , and being the only vertices in C that are

in these layers, no vertex covers the other (vertex-disjoint) hyper-edge between these layers. In the annoying case, this (i.e., (u, v) being an edge in G) means that u and v do not have the same label and one of the four hyper-edges in H cannot be covered by them; specifically, without loss of generality, suppose that u is in the first layer, then neither $v = v_{a,b}$ nor $u \neq u_{a,b}$ covers the hyper-edge $\{u_{a,b}, v_{a,1-b}, v_{1-a,b}, v_{1-a,1-b}\}$. ■

Conclusion: Starting from the abovementioned $\mathcal{NP} \subseteq \mathcal{FPCP}_{1-\epsilon, 0.5}[\log, 2]$ result of Hastad [10], we have shown that the minimum vertex-cover in 4-regular hypergraphs is NP-hard to approximate to a factor of $1.5/(1+3\epsilon)$. Let us state this as our second little result:

Theorem 4.3 *For every constant $\epsilon > 0$, approximating the size of the minimum vertex cover in 4-regular hypergraphs to within a $(1.5 - \epsilon)$ -factor is NP-hard.*

5 Subsequent Work

As hinted in the introduction, our motivation in this memo was to draw attention to the wide applicability of the FGLSS-reduction, and the specific results obtained were merely a good excuse to do so. Recall that all our results are inferior to Holmerin’s independently achieved result [12], by which approximating vertex cover in 4-regular hypergraphs to within a $(2 - \epsilon)$ -factor is NP-hard. Thus, the fact that also the latter result was subsequently improved is not relevant to the main motivation of the current work. Nevertheless, we briefly review some of the related results that appear after the current work was completed, differentiating between what was known already in 2001 and what is known in 2010.

Original postscript (2001). Following this work, Holmerin has applied related FGLSS-type reductions to different PCP systems and obtained improved inapproximability results for vertex cover in hypergraphs [14]. Specifically, for every constant $\epsilon > 0$, he showed that:

1. Approximating the size of the minimum vertex cover in k -regular hypergraphs to within a factor of $\Omega(k^{1-\epsilon})$ is NP-hard.
2. Approximating the size of the minimum vertex cover in 3-regular hypergraphs to within a factor of $1.5 - \epsilon$ is NP-hard.

Additional postscript (2010). The results reported in the original postscript were further improved by subsequent works, culminating in the following two results:

1. For every constant $\epsilon > 0$, approximating the size of the minimum vertex cover in k -regular hypergraphs to within a factor of $k - 1 - \epsilon$ is NP-hard [4].
2. Assuming the Unique Game Conjecture (UGC), for every constant $\epsilon > 0$ and every integer $k \geq 2$, it is hard to approximate the size of the minimum vertex cover in k -regular hypergraphs to within a factor of $k - \epsilon$ [15].³

³ Indeed, the case $k = 2$ drew most attention.

Acknowledgments

We are grateful to Johan Hastad for referring us to the works of Trevisan [17] and Holmerin [12].

References

1. S. Arora, C. Lund, R. Motwani, M. Sudan and M. Szegedy. Proof Verification and Intractability of Approximation Problems. *JACM*, Vol. 45, pages 501–555, 1998. Preliminary version in *33rd FOCS*, 1992.
2. S. Arora and S. Safra. Probabilistic Checkable Proofs: A New Characterization of NP. *JACM*, Vol. 45, pages 70–122, 1998. Preliminary version in *33rd FOCS*, 1992.
3. M. Bellare, O. Goldreich and M. Sudan. Free Bits, PCPs and Non-Approximability – Towards Tight Results. *SICOMP*, Vol. 27, No. 3, pages 804–915, 1998.
4. I. Dinur, V. Guruswami, S. Khot, and O. Regev. A New Multilayered PCP and the Hardness of Hypergraph Vertex Cover. *SICOMP*, Vol. 34, No. 5, pages 1129–1146, 2005.
5. I. Dinur and S. Safra. The Importance of Being Biased. Manuscript, 2001. See also [6].
6. I. Dinur and S. Safra. The importance of being biased. In *34th STOC*, pages 33–42, 2002.
7. U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy. Approximating Clique is almost NP-complete. *JACM*, Vol. 43, pages 268–292, 1996. Preliminary version in *32nd FOCS*, 1991.
8. O. Goldreich. *Modern Cryptography, Probabilistic Proofs and Pseudorandomness. Algorithms and Combinatorics series (Vol. 17)*, Springer, 1999.
9. J. Hastad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, Vol. 182, pages 105–142, 1999. Preliminary versions in *28th STOC* (1996) and *37th FOCS* (1996).
10. J. Hastad. Some optimal in-approximability results. In *29th STOC*, pages 1–10, 1997.
11. J. Hastad and S. Khot. Query efficient PCPs with Perfect Completeness. In *42nd FOCS*, pages 610–619, 2001.
12. J. Holmerin. Vertex Cover on 4-regular Hypergraphs is Hard to Approximate within $2 - \epsilon$. TR01-094, *ECCC*, 2001. See also [13].
13. J. Holmerin. Vertex Cover on 4-regular Hypergraphs is Hard to Approximate within $2 - \epsilon$. In *34th STOC*, pages 544–552, 2002.
14. J. Holmerin. Improved Inapproximability Results for Vertex Cover on k -regular Hypergraphs. In *29th ICALP*, pages 1005-1016, 2002.
15. S. Khot and O. Regev. Vertex Cover Might be Hard to Approximate to within $2 - \epsilon$. *JCSS*, Vol 74 (3), pages 335–349, 2008. Preliminary version in *18th Conf. on Comput. Complex.*, 2003.
16. E. Petrank. The Hardness of Approximations: Gap Location. *Computational Complexity*, Vol. 4, pages 133–157, 1994.
17. L. Trevisan. Non-approximability Results for Optimization Problems on Bounded-Degree Instances. In *33rd STOC*, pages 453–461, 2001.