A Linear Round Lower Bound for Lovasz-Schrijver SDP Relaxations of Vertex Cover

Grant Schoenebeck* Luca Trevisan† Madhur Tulsiani

Abstract

We study semidefinite programming relaxations of Vertex Cover arising from repeated applications of the LS+ "lift-and-project" method of Lovasz and Schrijver starting from the standard linear programming relaxation.

Goemans and Kleinberg prove that after one round of LS+ the integrality gap remains arbitrarily close to 2. Charikar proves an integrality gap of 2 for a stronger relaxation that is, however, incomparable with two rounds of LS+ and is strictly weaker than the relaxation resulting from a constant number of rounds.

We prove that the integrality gap remains at least $7/6 - \varepsilon$ after $c_\varepsilon n$ rounds, where n is the number of vertices and $c_\varepsilon > 0$ is a constant that depends only on ε.

1 Introduction

Lovasz and Schrijver [**?**] describe two versions of a "lift-and-project" method that, applied to a convex programming relaxation K of a 0/1 integer linear program, produces a tighter relaxation. A weaker version of the method, denoted LS, adds auxiliary variables and linear inequalities, and the projection of the new relaxation on the original variables is denoted by $N(K)$; a stronger version, denoted LS+, adds semidefinite programming constraints as well, and the projection on the original variables is denoted by $N_+(K)$.

Lovasz and Schrijver prove that if we start from a linear programming relaxation of a 0/1 integer program with n variables, then n applications of the LS procedures are sufficient to obtain a tight relaxation where the only feasible solutions are convex combinations of integral solutions. If one starts from a linear program with poly(n) inequalities, then it is possible to optimize over the set of solutions defined by k rounds of LS or LS+ in time $O(n^{O(k)})$, which is sub-exponential for $k = o(n/\log n)$.[1]

In many interesting cases, a small constant number of applications of LS+ are sufficient to transform a simple linear programming formulation of a problem into the semidefinite programming relaxation

*This material is based upon work supported under a National Science Foundation Graduate Research Fellowship.

†This material is based upon work supported by the National Science Foundation under grant CCF 0515231 and by the US-Israel Binational Science Foundation Grant 2002246.

[1]It is also possible to optimize over feasible solutions for $N^k(K)$ and $N_+^k(K)$ in time $n^{O(k)}$ provided that a separation oracle for K is computable in time poly(n). (That is, it is not necessary for K to be a linear or semidefinite programming relaxation with a polynomial number of inequalities.)

that yields the best known polynomial-time computable approximation. For example, one round of LS+ starting from the trivial linear programming relaxation of the independent set problem gives the Lovasz Theta functions; one round of LS+ starting from a trivial linear programming relaxation of the max cut problem gives the Goemans-Williamson relaxation; and the ARV relaxation of the sparsest cut problem is no stronger than three rounds of LS+ applied to the standard linear programming relaxation of sparsest cut. (See the discussion in [**?**].)

Integrality gap results for LS+ are thus very strong unconditional negative results, as they apply to a "model of computation" that includes the best known algorithms for several problems.

Arora, Bollobas, Lovasz, and Tourlakis [**?**, **?**, **?**] prove LS round lower bounds for Vertex Cover. They show that even after $\Omega_\varepsilon(\log n)$ rounds the integrality gap is at least $2-\varepsilon$, and that even after $\Omega_\varepsilon((\log n)^2)$ rounds the integrality gap is at least $1.5-\varepsilon$.

Buresh-Oppenheim, Galesy, Hoory, Magen and Pitassi [**?**], and Alekhnovich, Arora, Tourlakis [**?**] prove $\Omega(n)$ LS+ round lower bounds for proving the unsatisfiablity of random instances of 3SAT (and, in general, kSAT with $k \geq 3$) and $\Omega_\varepsilon(n)$ round lower bounds for achieving approximation factors better than $7/8-\varepsilon$ for Max 3SAT, better than $(1-\varepsilon)\ln n$ for Set Cover, and better than $k-1-\varepsilon$ for Hypergraph Vertex Cover in k-uniform hypergraphs. They leave open the question of proving LS+ round lower bounds for approximating the Vertex Cover problem.

The standard reduction from Max 3SAT to Vertex Cover shows that if one is able to approximate Vertex Cover within a factor better than 17/16 then one can approximate Max 3SAT within a factor better than 7/8. This fact, and the $7/8-\varepsilon$ integrality gap for Max 3SAT of [**?**], however do not suffice to derive an integrality gap result for Vertex Cover. The reason is that reducing an instance of Max 3SAT to a graph, and then applying a Vertex Cover relaxation to the graph, defines a semidefinite program that is possibly tighter than the one obtained by a direct relaxation of the Max 3SAT problem.

Feige and Ofek [**?**] are able to analyse the value of the Lovasz Theta function of the graph obtained by taking a random 3SAT instance and then reducing it to an instance of Independent Set (or, equivalently, of Vertex Cover). Their result immediately implies a $17/16-\varepsilon$ integrality gap for one round of LS+, and the way in which they prove their result implies also the stronger $7/6-\varepsilon$ bound. For one round of LS+ (or, equivalently, for the function defined as number of vertices minus the Theta function) Goemans and Kleinberg [**?**] had earlier proved a $2-o(1)$ integrality gap result by using a different family of graphs. Charikar [**?**] proves a $2-o(1)$ integrality gap result for a semidefinite programming relaxation of Vertex Cover that includes additional inequalities. Charikar's relaxation is no tighter than 3 rounds of LS+, and is incomparable with the relaxation obtained after two rounds.

It was compatible with previous results that after a constant number of rounds of LS+ or after $\operatorname{poly}\log n$ rounds of LS the integrality gap for Vertex Cover could become $1+o(1)$.

Our Result

We prove that after $\Omega_\varepsilon(n)$ rounds of LS+ the integrality gap remains at least $7/6-\varepsilon$. (For a stronger reason, the lower bound applies to LS as well.)

We combine ideas from the work of Alekhnovich, Arora and Tourlakis [**?**] and Feige and Ofek [**?**]. As in [**?**], we study the instance obtained by starting from a random instance of 3XOR and

then reducing it to the independent set problem; we also define our semidefinite programming solutions in a way that is similar to [**?**] (with the difference that we need to define such solutions inside an inductive argument, while only one solution is needed in [**?**]). As in [**?**], we maintain an expansion property as an invariant in our inductive argument, and we employ an "expansion correction procedure" to make sure that the invariant is maintained. Our way of realizing the expansion correction is new.

2 The Lovasz-Schrijver Method

2.1 Definitions

Let $R \subseteq [0,1]^N$ be a convex relaxation of a 0/1 Integer Program. We would like to "tighten" the relaxation by adding inequalities that are valid for 0/1 solutions but that are violated by other solutions.

Ideally, we would like to say that a solution $(y_1, \ldots, y_n)$ must satisfy the conditions $y_i^2 = y_i$, because such a condition is satisfied only by 0/1 solutions. Equivalently, we could introduce n^2 new variables $Y_{i,j}$ and add the conditions (i) $Y_{i,j} = y_i \cdot y_j$ and (ii) $Y_{i,i} = y_i$. Unfortunately, condition (i) is neither linear nor convex, and so we will instead "approximate" condition (i) by enforcing a set of linear conditions that are implied by (but not equivalent to) (i).

Before getting started, we will slightly change our setting because it is more convenient to work with a *convex cone* than with a convex subset of $[0,1]^N$. Recall that a *cone* is a subset K of $\mathbb{R}^d$ such that if $\mathbf{x}, \mathbf{y} \in K$ and $\alpha, \beta \geq 0$ then $\alpha\mathbf{x} + \beta\mathbf{y} \in K$, that is, a cone is a set of vectors that is closed under non-negative linear combinations. (Note that, in particular, a cone is always convex.)

If we are interested in a convex set $R \subseteq [0,1]^N$, we first convert it into the cone $K \subseteq \mathbb{R}^{N+1}$ defined as the set of all vectors $(\lambda, \lambda y_1, \ldots, \lambda y_N)$ such that $\lambda \geq 0$ and $(y_1, \ldots, y_N) \in R$. For example, in the "cone" linear programming relaxation of the vertex cover problem on a graph $G = (V, E)$ where $V = \{1, \ldots .N\}$, a solution $(y_0, \ldots, y_N)$ is feasible if and only if

$$\begin{aligned} y_i + y_j &\geq y_0 \quad \forall (i,j) \in E \\ 0 \leq y_i &\leq y_0 \quad \forall i \in V \\ y_0 &\geq 0 \end{aligned} \qquad (VC(G))$$

and in the cone relaxation of the independent set problem $(y_0, y_1, \ldots, y_N)$ is feasible if and only if

$$\begin{aligned} y_i + y_i &\leq y_0 \quad \forall (i,j) \in E \\ 0 \leq y_i &\leq y_0 \quad \forall i \in V \\ y_0 &\geq 0 \end{aligned} \qquad (IS(G))$$

We shall define the LS+ operator N_+ such that if K is a cone then $N_+(K)$ is cone, and we define $N_+^k(K)$ as $N_+(\cdots(N_+(K))\cdots)$ applied k times. For a graph G, the relaxation of the minimum

vertex cover problem resulting from k rounds of LS+ is the result of

$$\begin{aligned}
&\max \sum_{i=1}^{N} y_i \\
&\text{subject to} \\
&\qquad (y_0, \ldots, y_N) \in N_+^k(VC(G)) \\
&\qquad y_0 = 1
\end{aligned}$$

The operator N_+ will be such that if $(1, y_1, \ldots, y_N) \in K$ and y_i are all 0/1 bits, then $(1, y_1, \ldots, y_N) \in N_+(K)$, so that it will always map a relaxation of a integral problem into another relaxation.

We now come to the formal definition.

Definition 1 (N and N_+ Operators) *If K is a cone in $\mathbb{R}^d$, then we define the set $N(K)$ (which will also be a cone in $\mathbb{R}^d$) as follows: a vector $\mathbf{y} = (y_0, \ldots, y_{d-1}) \in \mathbb{R}^d$ is in $N(K)$ if and only if there is a matrix $Y \in R^{d \times d}$ such that*

1. *Y is symmetric;*
2. *For every $i \in \{0, 1, \ldots, d-1\}$, $Y_{0,i} = Y_{i,i} = y_i$*
3. *Each row Y_i is an element of K*
4. *Each vector $Y_0 - Y_i$ is an element of K*

In such a case, Y is called the protection matrix *of* $\mathbf{y}$.

If, in addition, Y is positive semidefinite, *then $\mathbf{y} \in N_+(K)$.*

Finally, we define $N^0(K)$ and $N_+^0(K)$ as K, and $N^k(K)$ (respectively, $N_+^k(K)$) as $N(N^{k-1}(K))$ (respectively, $N_+(N_+^{k-1}(K))$).

If $\mathbf{y} = (1, y_1, \ldots, y_{d-1}) \in \{0,1\}^d$, then we can set $Y_{i,j} = y_i \cdot y_j$. Such a matrix Y is clearly positive semidefinite, and it satisfies $Y_{i.i} = y_i^2 = y_i$ if the y_i are in $\{0,1\}$. Consider now a row Y_i of Y, that is, the vector $\mathbf{r}$ such that $r_j := Y_{i,j} = y_i \cdot y_j$. Then, either $y_i = 0$, in which case $\mathbf{r} = (0, \ldots, 0)$ is in every cone, or $y_i = 1$, and $\mathbf{r} = \mathbf{y}$. Similarly, if we consider $r_j := Y_{0,j} - Y_{i,j} = (1 - y_i) \cdot y_j$ we find that it either equals the all-zero vector or it equals $\mathbf{y}$. This shows that if $\mathbf{y} = (1, y_1, \ldots, y_{d-1}) \in \{0,1\}^d$ and $\mathbf{y} \in K$, then also $\mathbf{y} \in N_+^k(K)$ for every k.

Lovasz and Schrijver prove that if $(1, y_1, \ldots, y_{d-1}) \in N_+^k(K)$, then it satisfies every linear inequality over at most k variables that is valid for 0/1 solutions. In particular, if $(1, y_1, \ldots, y_{d-1}) \in N_+^{d-1}(K)$, then $(y_1, \ldots, y_{d-1})$ must be a convex combination of 0/1 solutions.

We will be interested in the integrality gap of LS+ relaxations of vertex cover. It will be easier, however, to reason about the independent set problem. The settings are equivalent.

Lemma 2 *Let $G = (V, E)$ be a graph and $V = 1, \ldots, N$. Then, for every $k \geq 0$, $(y_0, y_1, \ldots, y_N) \in N_+^k(VC(G))$ if and only if $(y_0, y_0 - y_1, \ldots, y_0 - y_N) \in N^k(IS(G))$.*

PROOF: We prove it by induction, with the $k = 0$ base case being clear. If $(y_0, y_1, \ldots, y_N) \in N_+^{k+1}(VC(G))$ then there is a protection matrix Y that is symmetric, positive semidefinite, and such that $Y_{i,i} = Y_{0,i} = y_i$, and such that the vectors Y_i and $Y_0 - Y_i$ are in $N^k(VC(G))$. Since Y is positive semidefinite, there must be vectors $\mathbf{b}_0, \ldots, \mathbf{b}_N$ such that $Y_{i,j} = \mathbf{b}_i \cdot \mathbf{b}_j$.

Consider now the vectors $\mathbf{c}_0, \ldots, \mathbf{c}_N$ defined as follows: $\mathbf{c}_0 := \mathbf{b}_0$ and $\mathbf{c}_i := \mathbf{b}_0 - \mathbf{b}_i$ for $i > 0$. Define the matrix Z as $Z_{i,j} := \mathbf{c}_i \cdot \mathbf{c}_j$. Thus the matrix Z is symmetric and positive semidefinite. We will argue that Z is a protection matrix by showing that the vector $\mathbf{z} := (y_0, y_0 - y_1, \ldots, y_0 - y_N) \in N_+^{k+1}(IS(G))$.

First, we see that $Z_{0,0} = Y_{0,0} = y_0$ and that, for $i > 0$,

$$Z_{i,i} = (\mathbf{b}_0 - \mathbf{b}_i) \cdot (\mathbf{b}_0 - \mathbf{b}_i) = Y_{0,0} - 2Y_{0,i} + Y_{i,i} = y_0 - y_i$$

Consider now the row vector Z_i, which is equal to $(r_0, \ldots, r_N)$ where

$$r_0 = \mathbf{b}_0 \cdot (\mathbf{b}_0 - \mathbf{b}_i) = y_0 - y_i$$

and, for $j > 0$,

$$r_j = (\mathbf{b}_0 - \mathbf{b}_j) \cdot (\mathbf{b}_0 - \mathbf{b}_i) = y_0 - y_j - y_i + Y_{i,j}$$

We need to show $(r_0, \ldots, r_N) \in N_+^k(IS(G))$ which, by the inductive hypothesis, is equivalent to $(r_0, r_0 - r_1, \ldots, r_0 - r_N) \in N_+^k(VC(G))$. But $(r_0, r_0 - r_1, \ldots, r_0 - r_N) = Y_0 - Y_i$ which belongs to $N_+^k(VC(G))$ by our assumption that Y is a protection matrix for $\mathbf{y}$. The other conditions are similarly verified. □

2.2 The Prover-Adversary Game

As done in previous work on Lovasz-Schrijver relaxations, in order to prove that a certain vector belongs to $N_+^k(IS(G))$, it will be convenient to formulate the argument in terms of a prover-adversary game.

A *prover* $\mathcal{P}$ is an algorithm that, on an input vector $(y_0, \ldots, y_N)$, either fails or outputs a matrix $Y \in \mathbb{R}^{(N+1)\times(N+1)}$ and a set of vectors $O \subseteq \mathbb{R}^{N+1}$ such that

1. Y is positive semidefinite
2. $Y_{i,i} = Y_{0,i} = y_i$
3. Each vector Y_i and $Y_0 - Y_i$ is a non-negative linear combination of vectors of O
4. Each element of O is in $IS(G)$

Consider now the following game played by a prover against another party called the *adversary*. We start from a vector $\mathbf{y} = (y_0, \ldots, y_N)$, and the prover, on input $\mathbf{y}$, outputs Y and O as before. Then the adversary chooses a vector $\mathbf{z} \in O$, and the prover, on input $\mathbf{z}$, outputs a matrix Y' and a set O', and so on. The adversary *wins* when the prover fails.

Lemma 3 *Suppose that there is a prover such that, starting from a vector $\mathbf{y} \in IS(G)$, every adversary strategy requires at least $k+1$ moves to win. Then $\mathbf{y} \in N^k(IS(G))$.*

PROOF: We proceed by induction on k, with $k = 0$ being the simple base case. Suppose that, for every adversary, it takes at least $k+1$ moves to win, and let Y and O be the output of the prover on input $\mathbf{y}$. Then, for every element $\mathbf{z} \in O$, and every prover strategy, it takes at least k moves to win starting from $\mathbf{z}$. By inductive, hypothesis, each element of O is in $N_+^{k-1}(IS(G))$, and since $N_+^{k-1}(IS(G))$ is closed under non-negative linear combinations, the vectors Y_i and $Y_0 - Y_i$ are all in $N_+^{k-1}(IS(G))$, and so Y is a protection matrix that shows that $\mathbf{y}$ is in $N_+^k(IS(G))$, $\square$

3 Overview of Our Result

Let φ be an instance of 3XOR, that is, a collection of linear equations mod 2 over variables $x_1, \ldots, x_n$ such that each equation is over exactly 3 variables. We denote by $\text{opt}(\varphi)$ the largest number of simultaneously satisfiable equations in φ.

Given a 3XOR instance φ with m equation, we define the FGLSS graph G_φ of φ as follows: G_φ has $4m$ vertices, one for each equation of φ and for each assignment to the three variables that satisfies the equation. We think of each vertex as being labeled by a partial assignment to three variables. Two vertices u and v are connected if and only if the partial assignments that label u and v are inconsistent. For example, for each equation, the four vertices corresponding to that equation form a clique. It is easy to see that $\text{opt}(\varphi)$ is precisely the size of the largest independent set of G_φ. Note that, in particular, the independent set size of G_φ is at most $N/4$, where $N = 4m$ is the number of vertices.

We say that φ is (k, c)-expanding if every set S of at most k equations in φ involves at least $c|S|$ distinct variables.

Our main result is that if φ is highly expanding, then even after a large number of rounds of Lovasz-Schrijver, the optimum of the relaxation is $N/4$, the largest possible value.

Lemma 4 (Main) *Let φ be a $(k, 1.95)$-expanding instance of 3XOR such that any two clauses share at most one variable, and let G_φ be its FGLSS graph.*

Then $(1, \frac{1}{4}, \ldots, \frac{1}{4})$ is in $N_+^{(k-4)/44}(IS(G_\varphi))$.

Our integrality gap result follows from the well known fact that there are highly expanding instances of 3XOR where it is impossible to satisfy significantly more than half of the equations.

Lemma 5 *For every $c < 2$ and $\varepsilon > 0$ there are $\alpha, \beta > 0$ such that for every n there is an instance φ of 3XOR with n variables and $m = \beta n$ equations such that*

- *No more than $(1/2 + \varepsilon)m$ equations are simultaneously satisfiable;*
- *Any two clauses share at most one variable*
- *φ is $(\alpha n, c)$-expanding.*

The rest of the paper is devoted to the proof of Lemma 4. We prove it in Section 4 by describing a prover strategy that survives for at least $(k-4)/44$ rounds. Proofs of variants of Lemma 5 have

appeared before in the literature, for example in [**?**, **?**]; we give a proof in the Appendix for the sake of self-containment.

The two lemmas combine to give our lower bound.

Theorem 6 *For every $\varepsilon > 0$ there is a $c_\varepsilon > 0$ such that for infinitely many t there is a graph G with t vertices such that the ratio between the minimum vertex cover size of G and the optimimum of $N^{c_\varepsilon t}(VC(G))$ is at least $7/6 - \varepsilon$.*

Proof: Using Lemma 5, construct an instance φ of 3XOR with n clauses and $O_\varepsilon(m)$ equations such that (i) no more than an $1/2+\varepsilon$ fraction of equations can be simultaneously satisfied; (ii) any two clauses share at most one variable; and (iii) φ is $(\Omega_\varepsilon(n), 1.95)$-expanding.

The minimum vertex size in the graph G_φ is at least $4m - (1/2 + \varepsilon)m$, but, by Lemma 4, the solution $(1, 3/4, \cdots, 3/4)$ is feasible for $N^{\Omega_\varepsilon(n)}(VC(G_\varphi))$, and so the optimum of $N^{\Omega_\varepsilon(n)}(VC(G_\varphi))$ is at most $3m$. □

4 The Prover Algorithm

For the sake of this section, we refer to a fixed formula φ with n variables $X = \{x_1, \ldots, x_n\}$ and m clauses which is $(k, 1.95)$-expanding and such that two clauses share at most one variable. The graph G_φ has $4m$ vertices. Recall that each vertex v of G_φ corresponds to an equation C of φ and to an assignment of values to the three variables of C that satisfies C. (In the following, if v is one of the vertices corresponding to an equation C, we call C the *equation of* v.)

4.1 Some Intuition

Suppose that φ were satisfiable, and let D be a distribution over satisfying assignments for φ. Then define the vector $\mathbf{y} = \mathbf{y}(D)$ as follows: $y_0 = 1$ and

$$y_v := \mathbf{Pr}_{a \in D}[a \text{ agrees with } v]$$

We claim that this solution is in $N_+^k(IS(G))$ for all k. This follows from the fact that it is a convex combination of 0/1 solutions, but it is instructive to construct the protection matrix for $\mathbf{y}$. Define the matrix Y such that $Y_{0,i} = Y_{i,0} = y_i$ and

$$Y_{u,v} := \mathbf{Pr}_{a \in D}[a \text{ agrees with } v \text{ and with } u]$$

It is easy to see that this matrix is positive semidefinite.

Consider now the u-th row of Y. If $y_u = 0$, then the row is the all-zero vector. Otherwise, it is y_0 times the vector $\mathbf{z} := (1, Y_{1,u}/y_0, \ldots, Y_{N,u}/y_0)$. Observe that, for $v \neq 0$,

$$z_v = \frac{\mathbf{Pr}_{a \in D}[a \text{ agrees with } v \text{ and with } u]}{\mathbf{Pr}_{a \in D}[a \text{ agrees with } u]} = \mathbf{Pr}_{a \in D}[a \text{ agrees with } v | a \text{ agrees with } u]$$

This is the vector $\mathbf{y}(D|u)$ where $(D|u)$ is the distribution D conditioned on assignments that agree with u.

Consider now the vector $Y_0 - Y_u$. If $y_u = 1$, then this is the all-zero vector. Otherwise, it is $(y_0 - y_u)$ times the vector $\mathbf{z} := (1, (y_1 - Y_{1,u})/(1-y_u), \ldots, (y_n - Y_{n,u})/(1-y_u))$. We have

$$z_v = \frac{\mathbf{Pr}_{a \in D}[a \text{ agrees with } v \text{ but not with } u]}{\mathbf{Pr}_{a \in D}[a \text{ does not agree with } u]} = \mathbf{Pr}_{a \in D}[a \text{ agrees with } v | a \text{ does not agree with } u]$$

And this is the same as $\mathbf{y}(D|\neg u)$, where $(D|\neg u)$ is the distribution D conditioned on assignments that do not agree with u. Note, also, that $\mathbf{y}(D|\neg u)$ can be realized as a convex combination of vectors $\mathbf{y}(D|w)$, where w ranges over the other vertices that correspond to satisfying assignments for the equation of u.

These observations suggest the following prover algorithm: on input a vector of the form $\mathbf{y}(D)$, output a matrix Y as above, and then set

$$O := \{\mathbf{y}(D|u) : u \in V \text{ and } \mathbf{Pr}_{a \in D}[a \text{ consistent with } u] > 0\}$$

To prove Lemma 4, we need to find a prover strategy that suceeds for a large number of rounds starting from the vector $(1, \frac{1}{4}, \ldots, \frac{1}{4})$. The above prover strategy would work if there is a distribution over satisfying assignments for φ such that, for each equation C, each of the four satisfying assignments for C occurs with probability $\frac{1}{4}$ in the distribution. Since we want to prove an integrality gap, however, we will need to work with highly unsatisfiable instances, and so no such distribution exists.

In our proof, we essentially proceed by pretending that such a distribution exists. Every time we "look" at certain equations and have certain conditions, we refer to the distribution that is uniform over all assignments that satisfy the equations and meet the conditions; this will mean that, for example, when defining the matrix Y we will refer to different distributions when filling up different entries. If the instance is highly expanding, however, it will take several rounds for the adversary to make the prover fail. This is because if there is an adversary strategy that makes the prover fail after k rounds, we can find a non-expanding subset of the formula of size $O(k)$.

4.2 Fractional Solutions and Protection Matrices Based on Partial Assignments

All the fractional solutions and protection matrices produced by the prover algorithm have a special structure and are based on *partial assignments* to the variables of φ. Before describing the prover algorithm, we will describe such solutions and matrices, and prove various facts about them.

A *partial assignment* $\alpha \subseteq X \times \{0,1\}$ is a set of assignments of values to some of the variables of φ such that each variable is given at most one value. For example, $\{(x_3, 0), (x_5, 1)\}$ is a partial assignment. A partial assignment α *contradicts* an equation of φ if it assigns values to all the three variables of the equations, and such values do not satisfy the equation; a partial assignment is *consistent* with φ if it contradicts none of the equations of φ.

If α is a consistent partial assignment, then the *restriction of φ to α*, denoted $\varphi_{|\alpha}$, is the set of equations that we obtain by applying the assignments of values to variables prescribed by α. (We remove the equations in whch all variables are assigned, and that are reduced to $0 = 0$.)

$\varphi_{|\alpha}$ contains some equations with three variables, some equations with two variables and some equations with one variable (as we said, we remove the equations with zero variables). If an

equation has two variables, we say those variables are α-*equivalent*. Note that α-equivalence is an equivalence relation, and so the variables in X not fixed by α are split into a collection of equivalence classes.

We make the following observations.

Claim 7 *If $\varphi_{|\alpha}$ is $(2, 1.51)$-expanding, then*

1. *Each equivalence class contains at most two variables;*
2. *If an equation contains three variables, then those variables belong to three distinct classes.*

The first part of the claim follows from the fact that, under the expansion assumption, all equations of size two are over disjoint sets of variables (otherwise, two equations of size two with a variable in common would form a set of two equations with only three occurring variables). The second part of the claim follows from the first part and from the assumption that in φ (and, for a stronger reason, in $\varphi_{|\alpha}$) two clauses can share at most one variable.

Definition 8 (Good partial assignment) *A partial assignment α is* good *for a formula φ if: (i) it is consistent with φ; (ii) $\varphi_{|\alpha}$ has no equation with only one variable; (iii) the set of all equations of $\varphi_{|\alpha}$ with two variables is satisfiable.*

The third condition seems very strong, but it is implied by expansion.

Claim 9 *Suppose that α is consistent with φ and that $\varphi_{|\alpha}$ is $(2, 1.51)$-expanding. Then α is a good partial assignment.*

PROOF: $\varphi_{|\alpha}$ cannot contain an equation with a single variable, otherwise it would not even be $(1, 1.1)$-expanding. Furthermore, any pair of equations with two variables cannot have any variable in common (otherwise we would have two equations involving only 3 variables), and so it is trivial to simultaneously satisfy all the size-2 equations. □

Definition 10 (α-Consistent Assignment) *If α is a good partial assignment for φ, then we say that an assignment $r \in \{0,1\}^n$ is α-consistent if it agrees with α and if it satisfies all the equations with two variables in $\varphi_{|\alpha}$.*

Definition 11 (Fractional solution associated to a good partial assignment) *Let α be a good partial assignment for φ. We describe the following fractional solution $\mathbf{y} = \mathbf{y}(\alpha)$ of the independent set problem in G_φ: $y(\alpha)_0 := 1$, and for every vertex v*

$$y(\alpha)_v := \mathbf{Pr}_{r \in \{0,1\}^n}[\ r \text{ agrees with } v \mid r \text{ agrees with } \alpha \text{ and satisfies } C\ ]$$

where C is the equation of v.

Another way of thinking of $\mathbf{y}(\alpha)$ is to remove from G_φ all the vertices that are inconsistent with α, and then, for each equation, split equally among the surviving vertices for that equation a total weight of 1. Note that, in $\mathbf{y}(\alpha)$ each entry is either 1, or 1/2, or 1/4 or 0.

Claim 12 *Let α be a good partial assignment for φ. Then $\mathbf{y}(\alpha)$ is a feasible solution in the cone $IS(G)$.*

PROOF: If two vertices u and v are connected in G, then there is a variable x such that u and v assign different values to x. If $\mathbf{y}(\alpha)$ assigns non-zero weight to both u and v, then it means that x is not assigned a value by α, and so both $y(\alpha)_u$ and $y(\alpha)_v$ are at most $1/2$. □

We also define the following "semidefinite solution."

Definition 13 (Protection Matrix Associated to a Partial Assignment) *Let α be a good partial assignment. To every vertex v we associate a $(d+1)$-dimensional vector $\mathbf{b}_v = \mathbf{b}_v(\alpha)$, where d is the number of equivalence classes in the set of variables of $\varphi_{|\alpha}$. When two variables are α-equivalent, we choose a representative. (Recall the earlier discussion about variables being α-equivalent.)*

- *If v is inconsistent with α, then we simply have $\mathbf{b}_v := (0, \ldots, 0)$.*
- *If α assigns values to all the variables of v (consistently with v), then $\mathbf{b}_v = (1, 0, \ldots, 0)$.*
- *If the equation of v has only two free variables in $\varphi_{|\alpha}$, they are in the same class, say the i-th class, and one of them is the representative. Then $\mathbf{b}_v = (1, 0, \ldots, 0, \pm\frac{1}{2}, 0, \ldots, 0)$, where the only non-zero entries are the 0th and the ith. The ith entry is $1/2$ if v requires the representative of the ith class to be 1; the ith entry is $-1/2$ otherwise.*
- *If the equation of v has three free variables in $\varphi_{|\alpha}$, then they are in three distinct classes, say the ith, the jth and the hth. Then $\mathbf{b}_v = (\frac{1}{4}, 0, \ldots, \pm\frac{1}{4}, \ldots \pm\frac{1}{4}, \ldots, \pm\frac{1}{4}, \ldots, 0)$, where the only nonzero entries are the 0th, the ith, the jth and the hth. The ith entry is $1/4$ if v requires the representative of the ith class to be 1, and $-1/4$ otherwise, and similarly for the other classes.*
- *Finally, let $\mathbf{b}_0(\alpha) = (1, 0 \ldots, 0)$.*

Define the matrix $Y(\alpha)$ as

$$Y_{u,v}(\alpha) := \mathbf{b}_u(\alpha) \cdot \mathbf{b}_v(\alpha) \tag{1}$$

Note that, by definition, $Y(\alpha)$ is positive semidefinite.

The matrix has the following equivalent characterization

Claim 14 *Let α be a good partial assignment such that $\varphi_{|\alpha}$ is $(4, 1.51)$-expanding. Then, for two vertices u, v, let C_1, C_2 be their equations; we have:*

$$Y_{u,v}(\alpha) = \mathbf{Pr}_{r \in \{0,1\}^n}[r \text{ agrees with } u \text{ and } v \mid r \text{ satisfies } C_1, C_2,\ r \text{ is } \alpha\text{-consistent }]$$

Furthermore, $Y_{0,u}(\alpha) = y_u(\alpha)$.

PROOF: To simplify notation we will omit the dependency on α.

If u and v correspond to two distinct assignments for the same equation, then it is easy to see that $Y_{u,v} = 0$.

If the equation of u and the equation of v have variables in disjoint classes, then $Y_{u,v} = \mathbf{b}_u \cdot \mathbf{b}_v = \mathbf{b}_{u,0}\mathbf{b}_{v,0}$, where

$$\mathbf{b}_{u,0} = \mathbf{Pr}_{r\in\{0,1\}^n}[r \text{ agrees with } u \mid r \text{ is } \alpha\text{-consistent }]$$

and

$$\mathbf{b}_{v,0} = \mathbf{Pr}_{r\in\{0,1\}^n}[r \text{ agrees with } v \mid r \text{ is } \alpha\text{-consistent }]$$

and, using independence

$$\mathbf{b}_{u,0}\mathbf{b}_{v,0} = \mathbf{Pr}_{r\in\{0,1\}^n}[r \text{ agrees with } u \text{ and } v \mid r \text{ is } \alpha\text{-consistent }]$$

If the equation of u and the equation of v each share precisely one variable from the same class i, then either both equations must involve three variables, or one equation involves two variables and the second involves two variables from the same class. In either case we have $Y_{u,v} = \mathbf{b}_{u,0}\mathbf{b}_{v,0} + \mathbf{b}_{u,i}\mathbf{b}_{v,i}$. In the first case, if the label of u and the label of v assign consistent values to the variable(s) in class i, then $Y_{u,v} = 1/8$, otherwise $Y_{u,v} = 0$, in accordance with the claim. In the second case, if the label of u and the label of v assign consistent values to the variable(s) in class i, then $Y_{u,v} = 1/4$, otherwise $Y_{u,v} = 0$, again, in accordance with the claim.

Finally, it is impossible for two distinct equations to have each two variables in common classes. Otherwise, we would have four equations involving at most six variables and contradict expansion. □

The matrix has also the following useful property.

Claim 15 *For a vertex v, let S denote the set of vertices corresponding to the equation of v which are consistent with α.*

Then

$$Y_0 - Y_v = \sum_{v'\in S-\{v\}} Y_{v'}$$

PROOF: The claim follows from the fact that

$$\sum_{v'\in S} \mathbf{b}(\alpha)_{v'} = \mathbf{b}(\alpha)_0 = (1, 0, \ldots, 0)$$

a fact that can be established by a simple cases analysis:

- If S contains only one element, then that element must be v, and it must be the case that $\mathbf{b}_v(\alpha) = (1, 0, \ldots, 0)$.
- If $S = \{v, v'\}$ contains two elements, then v and v' have $1/2$ in the first coordinate and then one has $1/2$ and another has $-1/2$ in the coordinate corresponding to the equivalence class of the two unassigned variables in the equation.

- If $S = \{v, v_1, v_2, v_3\}$ has four elements, then each one has $1/4$ in the first coordinates and then they have $\pm 1/4$ entries in the three coordinates corresponding to the three classes of the variables occurring in the equation. Each variable is given value zero in 2 vertices and value one in 2 vertices, so the entries in these three coordinates all cancel out.

□

4.3 Expansion and Satisfiability

Let α be a good partial assignments for φ, let C be an equation whose three variables are not assigned in α, and v be one of the vertices in G_φ corresponding to C. For the sake of this subsection, we think of v as being itself a partial assignment.

We define the "closure" of $\alpha \cup v$ as the output of the following algorithm

- $\beta := \alpha \cup v$;
- while φ_β has at least an equation with only one variable, of the form $x_i = b$
  - $\beta := \beta \cup \{(x_i, b)\}$
- return β

If φ_α is highly expanding, then the above algorithm terminates almost immediately, and it outputs an assignment β such that every small enough subset of the equations of φ_β are mutually satisfiable.

Lemma 16 *Suppose that φ_α is a $(k, 1.9)$-expanding instance of 3XOR let v be a vertex of G_φ that is not inconsistent with α. Let β be the closure of $\alpha \cup v$. Then β is a consistent partial assignment, and it fixes at most one variable not fixed in $\alpha \cup v$.*

PROOF: First, we note that $\varphi_{|\alpha \cup v}$ has at most one equation with only one variable. (Otherwise we would have three equations with a total of only 5 variables in $\varphi_{|\alpha}$.)

Let α' be $\alpha \cup v$ possibly extended to assign a value to the only equation of size one in $\varphi_{|\alpha \cup v}$ so that the equation is satisfied.

Then α' is a consistent partial assignment for φ such that $\varphi_{|\alpha'}$ has no equation of size one. (Otherwise, if $\varphi_{|\alpha'}$ had an equation of size one, then there would be three equations with five variables in $\varphi_{|\alpha}$.) We conclude that $\beta = \alpha'$ and the lemma follows. □

Lemma 17 (Satisfiability of Subsets of Expanding Instances) *Suppose that φ_α is a $(k, 1.9)$-expanding instance of 3XOR, let v be a vertex of G_φ corresponding to an equation involving variables not assigned by α. Let β be the closure of $\alpha \cup v$.*

Let S be any subset of at most $k - 2$ equations of $\varphi_{|\beta}$. Then there is assignment that satisfies all the equations of S. Furthermore, for every equation C in S and every assignment to the variables of C that satisfies C, it is possible to extend such an assignment to an assignment that satisfies all the equations in S.

PROOF: Recall that the difference between $\varphi_{|\beta}$ and $\varphi_{|\alpha}$ is that $\varphi_{|\beta}$ has either one fewer equation and at most three fewer variables than $\varphi_{|\alpha}$, or two fewer equations and at most four fewer variables than φ_α. (Depending on whether the closure algorithm performs zero steps or one step.)

Let C be an equation in $\varphi_{|\beta}$, let a be an assignment to the free variables in C that satisfies C, and let S be a smallest set of equations in φ_β such that S cannot be satisfied by an extension of a.

Suppose towards a contradiction that S contains at most $k-3$ equations.

Observe that, in $\varphi_{|\beta\cup a}$, every variable that occurs in S must occur in at least two equations of S, otherwise we would be violating minimality.

We will need to consider a few cases.

1. S cannot contain just a single equation C_1, because C_1 must have at least two variables in $\varphi_{|\beta}$, and it can share at most one variable with C.

2. Also, S cannot contain just two equations C_1 and C_2, because, for this to happen, C_1 and C_2 can have, between them, at most one variable not occurring in C, so that C, C_1 and C_2 are three clauses involving at most 4 variables in $\varphi_{|\beta}$; this leads to having either 4 equations involving at most 7 variables in $\varphi_{|\alpha}$, or to 5 equations involving at most 8 variables. In either case, we contradict the expansion assumption.

3. Consider now the case $|S| = 3$. We note that no equation in S can have three free variables in $\varphi_{|\beta\cup a}$, because then one of those three variables would not appear in the other two equations. Thus, each equation has at most two variables, each variable must occur in at least two equations, and so we have at most three variables occuring in S in $\varphi_{|\beta\cup a}$. In $\varphi_{|\alpha}$, this corresponds to either 5 clauses involving at most 9 variables, or 6 clauses involving at most 10 variables, and we again violate expansion.

4. If $|S| = 4$, then we consider two cases. If each equation in S has three free variables in $\varphi_{|\beta\cup a}$, then there can be at most 6 variables occurring in S, and we have a set of 4 equations in $\varphi_{|\alpha}$ involving only 6 variables.

   If some of the equations in S have less than three free variables, then at most a total of 5 variables can occur S in $\varphi_{|\beta\cup a}$. This means that we can find either 6 equations in φ_α involving at most 11 variables, or 7 equations involving at most 12 variables.

5. If $|S| \geq 5$, then at most $1.5 \cdot |S|$ variables can occur in S in $\varphi_{|\beta\cup a}$, and so we find either $|S|+2$ equations in $\varphi_{|\alpha}$ involving at most $\lfloor 1.5 \cdot |S| \rfloor + 6$ variables, or $|S|+3$ equations involving at most $\lfloor 1.5 \cdot |S| \rfloor + 7$ variables. Either situation violates expansion if $|S| \geq 5$.

□

4.4 Expansion-Correction

We will make use of the following simple fact.

Lemma 18 *Let ψ be an instance of 3XOR, and k be an integer. Then there is a subset $|S|$ of at most k equations such that:*

- *The instance $\psi - S$ is a $(k - |S|, 1.9)$-expanding instance of 3XOR;*
- *There is at most a total of $1.9|S|$ variables occurring in the equations in S.*

PROOF: Take a largest set S of equations in ψ such that $|S| \leq k$ and at most $1.9|S|$ variables occur in S. (Note that, possibly, S is the empty set.)

Suppose towards a contradiction that $\psi - S$ is not $(k - |S|, 1.9)$-expanding. Then there is a set T of equations in $\psi - S$ such that $|T| \leq k - |S|$ and at most $1.9|T|$ variables occur in T. Then the union of S and T and observe that it contradicts te maximality assumption about S. □

4.5 The Output of the Prover Algorithm

The prover algorithm takes in input a vector $\mathbf{y} = \mathbf{y}(\alpha)$ such that α is a consistent partial assignment and $\varphi_{|\alpha}$ is $(k, 1.9)$-expanding, $k \geq 4$. The output is a positive semidefinite matrix Y that is a protection matrix for $\mathbf{y}$ and a set of vectors $O \subseteq \mathbb{R}^{1+4m}$ such that each column Y_v of Y and each difference $Y_0 - Y_v$ are positive linear combinations of elements of O.

As we will see, each element of O is itself a vector of the form $\mathbf{y}(\beta)$, where β is an extension of α.

4.5.1 The Positive Semidefinite Matrix

The matrix Y is the matrix $Y(\alpha)$ defined in (1). By definition, Y is positive semidefinite. It also follows from the definition that $Y_{0,v} = Y_{v,v} = \mathbf{y}_v$.

4.5.2 The Set of Vectors

Because of Claim 15, each vector $Y_0 - Y_v$ is a non-negative linear combination of vectors Y_u, and so it is enough to prove that Y_v can be obtained as a non-negative combination of O.

In order to define the set O, we will define a set O_v for each vertex of the graph, and show that Y_v is a positive linear combination of elements of O_v. We will then define O to be the union of the sets O_v.

Let us fix a vertex v.

Let β be the closure of $\alpha \cup v$.

Let us now find a set $|S|$ of equations of $\varphi_{|\beta}$ as in Lemma 18 with parameter $k - 3$. The equations in S (if any), are simultaneously satisfiable by Lemma 17. Let A be the set of assigments that satisfy all equations in S. Define

$$O_v = \{\mathbf{y}(\beta \cup a) | a \in A\}$$

Lemma 19 *The vector Y_v is a non-negative combination of elements of O_v.*

PROOF: We will argue that

$$Y_v = Y_{0,v} \cdot \frac{1}{|A|} \sum_{a \in A} \mathbf{y}(\beta \cup a)$$

As a warm-up, we note that $Y_{v,v} = Y_{0,v}$ (as observed before) and that $\mathbf{y}_v(\beta \cup a) = 1$ for every a (because β already sets all the variables of v consistently with the label of v).

Let us now consider $u \neq v$, and let C be the equation of v. Recall that $Y_{u,v}$ has the following probabilistic interpretation:

$$Y_{u,v} = \mathbf{Pr}_{r\in\{0,1\}^n}[\ r \text{ agrees with } u \text{ and } v \mid r \text{ preserves } \alpha\text{-consistency, } r \text{ satisfies } C, C']$$

where C is the equation of u and C' is the equation of v.

We can also derive a probabilistic interpretation of the right-hand side of the equation we wish to prove

$$\frac{1}{|A|}\sum_{a\in A}\mathbf{y}(\beta\cup a)_u =$$

$$= \mathbf{Pr}_{a\in A, r\in\{0,1\}^n}[\ r \text{ agrees with } u \mid r \text{ satisfies } C \text{ and agrees with } \beta\cup a\ ] \tag{2}$$

Now we claim that the probability (2) is precisely the same as

$$\mathbf{Pr}_{r\in\{0,1\}^n}[\ r \text{ agrees with } u \mid r \text{ satisfies } C \text{ and agrees with } \beta\ ] \tag{3}$$

This is clear if the clauses in S and C share no variable outside β, because the conditioning on a has no influence on the event we are considering.

If C shares some, but not all, of its variables outside β with the clauses in S, then a random element a of A assigns uniform and independent values to such variables. This is because A is an affine space, and so if the above were not true, then A would force an affine dependency among a strict subset of the variables of C outside β; this would mean that there is a satisfying assignment for C that is inconsistent with each assignment in A, that is, there is a satisfying assignment for C that is inconsistent with S (in $\varphi_{|\beta}$), thus violating Lemma 17.) If C shares all its variables with the clauses of S, then a random a in A must assign to the variables of C a random satisfying assignment. (Otherwise, we would again conclude that there is a satisfying assignment for C that is inconsistent with S in $\varphi_{|\beta}$.)

The next step is to observe that, thus,

$$\frac{1}{|A|}\sum_{a\in A}\mathbf{y}(\beta\cup a)$$

is the same as the probability that a random extension of α is consistent with u conditioned on: (i) being consistent with v; (ii) satisfying the equation of u; (iii) preserving α-consistency. If we multiply by $Y_{0,v}$, what we get is the probability that a random extension of α is consistent with the labels of u and v conditioned on: (i) satisfying the equation of v; (ii) satisfying the equation of u; (iii) preserving α-consistency. And this is just the definition of $Y_{u,v}$. □

4.6 Putting Everything Together

Let φ be a $(k, 1.95)$-expanding instance of 3XOR, and suppose that there is an adversary strategy that makes the game terminate after r steps.

The game begins with the solution $(1, 1/4, \ldots, 1/4)$, which is $\mathbf{y}(\emptyset)$, and, at each round, the prover picks a solution of the form $\mathbf{y}(\alpha)$ for a partial assignment α. The game ends when $\varphi_{|\alpha}$ is not a $(4, 1.9)$-expanding instance of 3XOR.

Let us denote by $\mathbf{y}(\emptyset)$, $\mathbf{y}(\alpha_1)$, $\ldots$, $\mathbf{y}(\alpha_r)$ the solutions chosen by the adversary. Note that α_i is an extension of α_{i-1} in which the variables occuring in a set S_i of clauses have been fixed, in addition to the the variables occuring in one or two clauses (call this set T_i). We also have that S_i contains at most $1.9|S_i|$ variables that do not occur in T_j $j \leq i$ or in S_j, $j \leq i-1$. By the properties of the expansion correction, $\mathbf{y}(\alpha_i)$ is $(k - \sum_{j \leq i} |S_j| + |T_j|, 1.9)$-expanding.

When the game terminates, we have

$$k \geq \sum_i |S_i| + |T_i| \geq k - 4$$

Let t be total number of variables occurring in the S_i and T_i. We have

$$t \geq 1.95(\sum_i |S_i| + |T_i|)$$

because of the expansion in φ. But we also have

$$t \leq 3|T_i| + 1.9 \sum_i |S_i|$$

so

$$\sum_i |S_i| \leq 21 \sum_i |T_i|$$

and

$$k \leq 4 + 22 \sum_i |T_i| \leq 4 + 44r$$

which gives $r \geq k/44 - 1/11$.

A Appendix: Proof of Lemma 5

Let $\beta = \frac{\ln 2}{2\varepsilon^2} + 1$. We will first show that with probability $1 - o(1)$, only $1/2 + \varepsilon$ of the clauses are satisfied. We will then show that such a formula is a $(\alpha n, c)$ expander with probability $1 - o(1)$. Finally, we will show that the probability such a randomly chosen formula has no two clauses which share two variables is at least some constant. Using a union bound, the proof follows.

Fix an assignment to n variables. Now if we choose, $m = \beta n$ clauses at random, the probability that more than a $1/2 + \varepsilon$ fraction of them are satisfied is at most $\exp(-2\varepsilon^2 m) = \exp(-2\varepsilon^2 \beta n)$. To get this, we use the Chernoff Bound that says

$$\Pr[X \geq \mathbb{E}[X] + \lambda] \leq \exp(-2\lambda^2/m)$$

where X is the number of satisfied clauses, $\mathbb{E}[X] = m/2$, $\lambda = \varepsilon m$. Picking a random formula and random assignment, the probability that more than a $1/2+\varepsilon$ fraction of the clauses are satisfied is $\exp(-2\varepsilon^2\beta n)$. Taking a union bound over all assignments, we get

$$\begin{aligned}\Pr[\text{any assignment satisfies} \geq (1/2+\varepsilon)m \text{ clauses}] &\leq \exp(-2\varepsilon^2\beta n)\cdot 2^n \\ &= \exp(n(\ln 2 - 2\varepsilon^2\beta)) = \exp(-2\varepsilon^2 n)\end{aligned}$$

by our choice of β.

Now we bound the probability that for a random formula φ, H_φ is *not* a $(\alpha n, c)$-expander. The probability that there is a set of k clauses containing a total of fewer than ck variables can be upper bounded as $(O(1)k/n)^{(2-c)k}$ and so, as we will later show, it can be made $o(1)$, even after summing over all k from 1 to αn, for a proper choice of α. We can upper bound the probability that there is a set of k clauses containing a total of fewer than ck variables by

$$\binom{n}{ck}\cdot\binom{\binom{ck}{3}}{k}\cdot k!\cdot\binom{m}{k}\cdot\binom{n}{3}^{-k}$$

where $\binom{n}{ck}$ is the choice of the variables, $\binom{\binom{ck}{3}}{k}$ is the choice of the k clauses constructed out of such variables, $k!\cdot\binom{m}{k}$ is a choice of where to put such clauses in our ordered sequence of m clauses, and $\binom{n}{3}^{-k}$ is the probability that such clauses were generated as prescribed.

Using $\binom{N}{K} < (eN/K)^K$, $k! < k^k$, $m = \beta n$, and we simplify to obtain the upper bound $(O(k/n))^{(2-c)k}$. Next, we look at

$$\sum_{k=1}^{\alpha n}(O(k/n))^{(2-c)k} = \sum_{k=1}^{\ln^2 n}(O(k/n))^{(2-c)k} + \sum_{k=\ln^2 n+1}^{\alpha n}(O(k/n))^{(2-c)k}$$

Now $\sum_{k=1}^{\infty} t^k = t/(1-t) \leq 2t$ if $t \leq 1/2$. So

$$\sum_{k=1}^{\ln^2 n}(O(k/n))^{(2-c)k} \leq \sum_{k=1}^{\infty}((O(\ln^2 n/n))^{2-c})^k \leq 2(O(\ln^2 n/n))^{2-c})$$

for sufficiently large n, which is $o(1)$. Also,

$$\sum_{k=\ln^2 n+1}^{\alpha n}(O(k/n))^{(2-c)k} \leq (O(\alpha))^{(2-c)\ln^2 n}\sum_{k=1}^{\infty}(O(\alpha)^{2-c})^k \leq 2(O(\alpha))^{(2-c)\ln^2 n}$$

for sufficiently small α. Also, this is $o(1)$ for a small enough α.

Finally, the probability that there are no two clauses sharing two variables must be at least $\prod_{k=1,\ldots,m}(1-O(k)/n^2)$ because when we choose the kth clause, by wanting it to not share two variables with another previously chosen clause, we are forbidding O(k) pairs of variables to occur together. Each such pair happens to be in the clause with probability $O(1/n^2)$. Now we use that for small enough x, $1-x > \exp(-O(x))$ the probability is at least $\exp(-O((\sum_{k=1,\ldots,m} k)/n^2)) = \exp(-O(m^2/n^2)) = \exp(-O(\beta^2))$ which is some positive constant.

References