

GROTHENDIECK-TYPE INEQUALITIES IN COMBINATORIAL OPTIMIZATION

SUBHASH KHOT AND ASSAF NAOR

ABSTRACT. We survey connections of the Grothendieck inequality and its variants to combinatorial optimization and computational complexity.

CONTENTS

1. Introduction	2
1.1. Assumptions from computational complexity	3
1.2. Convex and semidefinite programming	4
2. Applications of the classical Grothendieck inequality	5
2.1. Cut norm estimation	5
2.1.1. Szemerédi partitions	8
2.1.2. Frieze-Kannan matrix decomposition	10
2.1.3. Maximum acyclic subgraph	11
2.1.4. Linear equations modulo 2	14
2.2. Rounding	16
3. The Grothendieck constant of a graph	18
3.1. Algorithmic consequences	20
3.1.1. Spin glasses	20
3.1.2. Correlation clustering	20
4. Kernel clustering and the propeller conjecture	21
5. The L_p Grothendieck problem	25
6. Higher rank Grothendieck inequalities	28
7. Hardness of approximation	29
References	31

S. K. was partially supported by NSF CAREER grant CCF-0833228, NSF Expeditions grant CCF-0832795, an NSF Waterman award, and BSF grant 2008059. A. N. was partially supported by NSF Expeditions grant CCF-0832795, BSF grant 2006009, and the Packard Foundation.

1. INTRODUCTION

The Grothendieck inequality asserts that there exists a universal constant $K \in (0, \infty)$ such that for every $m, n \in \mathbb{N}$ and every $m \times n$ matrix $A = (a_{ij})$ with real entries we have

$$\max \left\{ \sum_{i=1}^m \sum_{j=1}^n a_{ij} \langle x_i, y_j \rangle : \{x_i\}_{i=1}^m, \{y_j\}_{j=1}^n \subseteq S^{m+m-1} \right\} \\ \leq K \max \left\{ \sum_{i=1}^m \sum_{j=1}^n a_{ij} \varepsilon_i \delta_j : \{\varepsilon_i\}_{i=1}^m, \{\delta_j\}_{j=1}^n \subseteq \{-1, 1\} \right\}. \quad (1)$$

Here, and in what follows, the standard scalar product on \mathbb{R}^k is denoted $\langle x, y \rangle = \sum_{i=1}^k x_i y_i$ and the Euclidean sphere in \mathbb{R}^k is denoted $S^{k-1} = \{x \in \mathbb{R}^k : \sum_{i=1}^k x_i^2 = 1\}$. We refer to [34, 56] for the simplest known proofs of the Grothendieck inequality; see Section 2.2 for a proof of (1) yielding the best known bound on K . Grothendieck proved the inequality (1) in [45], though it was stated there in a different, but equivalent, form. The formulation of the Grothendieck inequality appearing in (1) is due to Lindenstrauss and Pełczyński [83].

The Grothendieck inequality is of major importance to several areas, ranging from Banach space theory to C^* algebras and quantum information theory. We will not attempt to indicate here this wide range of applications of (1), and refer instead to [83, 114, 100, 55, 37, 34, 19, 1, 40, 33, 102, 101] and the references therein. The purpose of this survey is to focus solely on applications of the Grothendieck inequality and its variants to combinatorial optimization, and to explain their connections to computational complexity.

The infimum over those $K \in (0, \infty)$ for which (1) holds for all $m, n \in \mathbb{N}$ and all $m \times n$ matrices $A = (a_{ij})$ is called the Grothendieck constant, and is denoted K_G . Evaluating the exact value of K_G remains a long-standing open problem, posed by Grothendieck in [45]. In fact, even the second digit of K_G is currently unknown, though clearly this is of lesser importance than the issue of understanding the structure of matrices A and spherical configurations $\{x_i\}_{i=1}^m, \{y_j\}_{j=1}^n \subseteq S^{m+m-1}$ which make the inequality (1) “most difficult”. Following a series of investigations [45, 83, 107, 77, 78], the best known upper bound [21] on K_G is

$$K_G < \frac{\pi}{2 \log(1 + \sqrt{2})} = 1.782\dots, \quad (2)$$

and the best known lower bound [105] on K_G is

$$K_G \geq \frac{\pi}{2} e^{\eta_0^2} = 1.676\dots, \quad (3)$$

where $\eta_0 = 0.25573\dots$ is the unique solution of the equation

$$1 - 2\sqrt{\frac{2}{\pi}} \int_0^\eta e^{-z^2/2} dz = \frac{2}{\pi} e^{-\eta^2}.$$

In [104] the problem of estimating K_G up to an additive error of $\varepsilon \in (0, 1)$ was reduced to an optimization over a compact space, and by exhaustive search over an appropriate net it was shown that there exists an algorithm that computes K_G up to an additive error of $\varepsilon \in (0, 1)$ in time $\exp(\exp(O(1/\varepsilon^3)))$. It does not seem likely that this approach can yield computer assisted proofs of estimates such as (2) and (3), though to the best of our knowledge this has not been attempted.

In the above discussion we focused on the classical Grothendieck inequality (1). However, the literature contains several variants and extensions of (1) that have been introduced for various purposes and applications in the decades following Grothendieck’s original work. In this survey we describe some of these variants, emphasizing relatively recent developments that yielded Grothendieck-type inequalities that are a useful tool in the design of polynomial time algorithms for computing approximate solutions of computationally hard optimization problems. In doing so, we omit some important topics, including applications of the Grothendieck inequality to communication complexity and quantum information theory. While these research directions can be viewed as dealing with a type of optimization problem, they are of a different nature than the applications described here, which belong to classical optimization theory. Connections to communication complexity have already been covered in the survey of Lee and Shraibman [81]; we refer in addition to [84, 80, 85, 86] for more information on this topic. An explanation of the relation of the Grothendieck inequality to quantum mechanics is contained in Section 19 of Pisier’s survey [101], the pioneering work in this direction being that of Tsirelson [114]. An investigation of these questions from a computational complexity point of view was initiated in [28], where it was shown, for example, how to obtain a polynomial time algorithm for computing the entangled value of an XOR game based on Tsirelson’s work. We hope that the developments surrounding applications of the Grothendieck inequality in quantum information theory will eventually be surveyed separately by experts in this area. Interested readers are referred to [114, 37, 28, 1, 54, 98, 102, 61, 22, 80, 86, 106, 101]. Perhaps the most influential variants of the Grothendieck inequality are its noncommutative generalizations. The noncommutative versions in [99, 49] were conjectured by Grothendieck himself [45]; additional extensions to operator spaces are extensively discussed in Pisier’s survey [101]. We will not describe these developments here, even though we believe that they might have applications to optimization theory. Finally, multi-linear extensions of the Grothendieck inequality have also been investigated in the literature; see for example [115, 112, 20, 109] and especially Blei’s book [19]. We will not cover this research direction since its relation to classical combinatorial optimization has not (yet?) been established, though there are recent investigations of multi-linear Grothendieck inequalities in the context of quantum information theory [98, 80].

Being a mainstay of functional analysis, the Grothendieck inequality might attract to this survey readers who are not familiar with approximation algorithms and computational complexity. We wish to encourage such readers to persist beyond this introduction so that they will be exposed to, and hopefully eventually contribute to, the use of analytic tools in combinatorial optimization. For this reason we include Sections 1.1, 1.2 below; two very basic introductory sections intended to quickly provide background on computational complexity and convex programming for non-experts.

1.1. Assumptions from computational complexity. At present there are few unconditional results on the limitations of polynomial time computation. The standard practice in this field is to frame an impossibility result in computational complexity by asserting that the polynomial time solvability of a certain algorithmic task would contradict a benchmark hypothesis. We briefly describe below two key hypotheses of this type.

A graph $G = (V, E)$ is 3-colorable if there exists a partition $\{C_1, C_2, C_3\}$ of V such that for every $i \in \{1, 2, 3\}$ and $u, v \in C_i$ we have $\{u, v\} \notin E$. The $P \neq NP$ hypothesis asserts that there is no polynomial time algorithm that takes an n -vertex graph as input and

determines whether or not it is 3-colorable. We are doing an injustice to this important question by stating it this way, since it has many far-reaching equivalent formulations. We refer to [39, 108, 31] for more information, but for non-experts it suffices to keep the above simple formulation in mind.

When we say that assuming $P \neq NP$ no polynomial time algorithm can perform a certain task \mathcal{T} (e.g., evaluating the maximum of a certain function up to a predetermined error) we mean that given an algorithm ALG that performs the task \mathcal{T} one can design an algorithm ALG' that determines whether or not any input graph is 3-colorable while making at most polynomially many calls to the algorithm ALG , with at most polynomially many additional Turing machine steps. Thus, if ALG were a polynomial time algorithm then the same would be true for ALG' , contradicting the $P \neq NP$ hypothesis. Such results are called hardness results. The message that non-experts should keep in mind is that a hardness result is nothing more than the design of a new algorithm for 3-colorability, and if one accepts the $P \neq NP$ hypothesis then it implies that there must exist inputs on which ALG takes super-polynomial time to terminate.

The Unique Games Conjecture (UGC) asserts that for every $\varepsilon \in (0, 1)$ there exists a prime $p = p(\varepsilon) \in \mathbb{N}$ such that no polynomial time algorithm can perform the following task. The input is a system of m linear equations in n variables x_1, \dots, x_n , each of which has the form $x_i - x_j \equiv c_{ij} \pmod p$ (thus the input is $S \subseteq \{1, \dots, n\} \times \{1, \dots, n\}$ and $\{c_{ij}\}_{(i,j) \in S} \subseteq \mathbb{N}$). The algorithm must determine whether there exists an assignment of an integer value to each variable x_i such that at least $(1 - \varepsilon)m$ of the equations are satisfied, or whether no assignment of such values can satisfy more than εm of the equations. If neither of these possibilities occur, then an arbitrary output is allowed.

As in the case of $P \neq NP$, saying that assuming the UGC no polynomial time algorithm can perform a certain task \mathcal{T} is the same as designing a polynomial time algorithm that solves the above linear equations problem while making at most polynomially many calls to a “black box” that can perform the task \mathcal{T} . The UGC was introduced in [62], though the above formulation of it, which is equivalent to the original one, is due to [64]. The use of the UGC as a hardness hypothesis has become popular over the past decade; we refer to the survey [63] for more information on this topic.

To simplify matters (while describing all the essential ideas), we allow polynomial time algorithms to be randomized. Most (if not all) of the algorithms described here can be turned into deterministic algorithms, and corresponding hardness results can be stated equally well in the context randomized or deterministic algorithms. We will ignore these distinctions, even though they are important. Moreover, it is widely believed that in our context these distinctions do not exist, i.e., randomness does not add computational power to polynomial time algorithms; see for example the discussion of the $NP \not\subseteq BPP$ hypothesis in [11].

1.2. Convex and semidefinite programming. An important paradigm of optimization theory is that one can efficiently optimize linear functionals over compact convex sets that have a “membership oracle”. A detailed exposition of this statement is contained in [46], but for the sake of completeness we now quote the precise formulation of the results that will be used in this article.

Let $K \subseteq \mathbb{R}^n$ be a compact convex set. We are also given a point $z \in \mathbb{Q}^n$ and two radii $r, R \in (0, \infty) \cap \mathbb{Q}$ such that $B(z, r) \subseteq K \subseteq B(z, R)$, where $B(z, t) = \{x \in \mathbb{R}^n : \|x - z\|_2 \leq t\}$. In what follows, stating that an algorithm is polynomial means that we allow the running time

to grow at most polynomially in the number of bits required to represent the data (z, r, R) . Thus, if, say, $z = 0$, $r = 2^{-n}$ and $R = 2^n$ then the running time will be polynomial in the dimension n . Assume that there exists an algorithm ALG with the following properties. The input of ALG is a vector $y \in \mathbb{Q}^n$ and $\varepsilon \in (0, 1) \cap \mathbb{Q}$. The running time of ALG is polynomial in n and the number of bits required to represent the data (ε, y) . The output of ALG is the assertion that either the distance of y from K is at most ε , or that the distance of y from the complement of K is at most ε . Then there exists an algorithm ALG' that takes as input a vector $c = (c_1, \dots, c_n) \in \mathbb{Q}^n$ and $\varepsilon \in (0, 1) \cap \mathbb{Q}$ and outputs a vector $y = (y_1, \dots, y_n) \in \mathbb{R}^n$ that is at distance at most ε from K and for every $x = (x_1, \dots, x_n) \in K$ that is at distance greater than ε from the complement of K we have $\sum_{i=1}^n c_i y_i \geq \sum_{i=1}^n c_i x_i - \varepsilon$. The running time of ALG' is allowed to grow at most polynomially in n and the number of bits required to represent the data $(z, r, R, c, \varepsilon)$. This important result is due to [57]; we refer to [46] for an excellent account of this theory.

The above statement is a key tool in optimization, as it yields a polynomial time method to compute the maximum of linear functionals on a given convex body with arbitrarily good precision. We note the following special case of this method, known as semidefinite programming. Assume that $n = k^2$ and think of \mathbb{R}^n as the space of all $k \times k$ matrices. Assume that we are given a compact convex set $K \subseteq \mathbb{R}^n$ that satisfies the above assumptions, and that for a given $k \times k$ matrix (c_{ij}) we wish to compute in polynomial time (up to a specified additive error) the maximum of $\sum_{i=1}^k \sum_{j=1}^k c_{ij} x_{ij}$ over the set of symmetric positive semidefinite matrices (x_{ij}) that belong to K . This can indeed be done, since determining whether a given symmetric matrix is (approximately) positive semidefinite is an eigenvalue computation and hence can be performed in polynomial time. The use of semidefinite programming to design approximation algorithms is by now a deep theory of fundamental importance to several areas of theoretical computer science. The Goemans-Williamson MAX-CUT algorithm [42] was a key breakthrough in this context. It is safe to say that after the discovery of this algorithm the field of approximation algorithms was transformed, and many subsequent results, including those presented in the present article, can be described as attempts to mimic the success of the Goemans-Williamson approach in other contexts.

2. APPLICATIONS OF THE CLASSICAL GROTHENDIECK INEQUALITY

The classical Grothendieck inequality (1) has applications to algorithmic questions of central interest. These applications will be described here in some detail. In Section 2.1 we discuss the cut norm estimation problem, whose relation to the Grothendieck inequality was first noted in [8]. This is a generic combinatorial optimization problem that contains well-studied questions as subproblems. Examples of its usefulness are presented in Sections 2.1.1, 2.1.2, 2.1.3, 2.1.4. Section 2.2 is devoted to the rounding problem, including the (algorithmic) method behind the proof of the best known upper bound on the Grothendieck constant.

2.1. Cut norm estimation. Let $A = (a_{ij})$ be an $m \times n$ matrix with real entries. The cut norm of A is defined as follows

$$\|A\|_{cut} = \max_{\substack{S \subseteq \{1, \dots, m\} \\ T \subseteq \{1, \dots, n\}}} \left| \sum_{\substack{i \in S \\ j \in T}} a_{ij} \right|. \quad (4)$$

We will now explain how the Grothendieck inequality can be used to obtain a polynomial time algorithm for the following problem. The input is an $m \times n$ matrix $A = (a_{ij})$ with real entries, and the goal of the algorithm is to output in polynomial time a number α that is guaranteed to satisfy

$$\|A\|_{cut} \leq \alpha \leq C\|A\|_{cut}, \quad (5)$$

where C is a (hopefully not too large) universal constant. A closely related algorithmic goal is to output in polynomial time two subsets $S_0 \subseteq \{1, \dots, m\}$ and $T_0 \subseteq \{1, \dots, n\}$ satisfying

$$\left| \sum_{\substack{i \in S_0 \\ j \in T_0}} a_{ij} \right| \geq \frac{1}{C} \|A\|_{cut}. \quad (6)$$

The link to the Grothendieck inequality is made via two simple transformations. Firstly, define an $(m+1) \times (n+1)$ matrix $B = (b_{ij})$ as follows.

$$B = \begin{pmatrix} a_{11} & a_{12} & \dots & a_{1n} & -\sum_{k=1}^n a_{1k} \\ a_{21} & a_{22} & \dots & a_{2n} & -\sum_{k=1}^n a_{2k} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ a_{m1} & a_{m2} & \dots & a_{mn} & -\sum_{k=1}^n a_{mk} \\ -\sum_{\ell=1}^m a_{\ell 1} & -\sum_{\ell=1}^m a_{\ell 2} & \dots & -\sum_{\ell=1}^m a_{\ell n} & \sum_{k=1}^n \sum_{\ell=1}^m a_{\ell k} \end{pmatrix}. \quad (7)$$

Observe that

$$\|A\|_{cut} = \|B\|_{cut}. \quad (8)$$

Indeed, for every $S \subseteq \{1, \dots, m+1\}$ and $T \subseteq \{1, \dots, n+1\}$ define $S^* \subseteq \{1, \dots, m\}$ and $T^* \subseteq \{1, \dots, n\}$ by

$$S^* = \begin{cases} S & \text{if } m+1 \notin S, \\ \{1, \dots, m\} \setminus S & \text{if } m+1 \in S, \end{cases} \quad \text{and} \quad T^* = \begin{cases} T & \text{if } n+1 \notin T, \\ \{1, \dots, n\} \setminus T & \text{if } n+1 \in T. \end{cases}$$

One checks that for all $S \subseteq \{1, \dots, m+1\}$ and $T \subseteq \{1, \dots, n+1\}$ we have

$$\left| \sum_{\substack{i \in S \\ j \in T}} b_{ij} \right| = \left| \sum_{\substack{i \in S^* \\ j \in T^*}} a_{ij} \right|,$$

implying (8). We next claim that

$$\|B\|_{cut} = \frac{1}{4} \|B\|_{\infty \rightarrow 1}, \quad (9)$$

where

$$\|B\|_{\infty \rightarrow 1} = \max \left\{ \sum_{i=1}^{m+1} \sum_{j=1}^{n+1} b_{ij} \varepsilon_i \delta_j : \{\varepsilon_i\}_{i=1}^{m+1}, \{\delta_j\}_{j=1}^{n+1} \subseteq \{-1, 1\} \right\}. \quad (10)$$

To explain this notation observe that $\|B\|_{\infty \rightarrow 1}$ is the norm of B when viewed as a linear operator from ℓ_∞^{n+1} to ℓ_1^{m+1} . Here, and in what follows, for $p \in [1, \infty]$ and $k \in \mathbb{N}$ the space ℓ_p^k is \mathbb{R}^k equipped with the ℓ_p norm $\|\cdot\|_p$, where $\|x\|_p^p = \sum_{\ell=1}^k |x_\ell|^p$ for $x = (x_1, \dots, x_k) \in \mathbb{R}^k$ (for $p = \infty$ we set as usual $\|x\|_\infty = \max_{i \in \{1, \dots, n\}} |x_i|$). Though it is important, this operator theoretic interpretation of the quantity $\|B\|_{\infty \rightarrow 1}$ will not have any role in this survey, so it may be harmlessly ignored at first reading.

The proof of (9) is simple: for $\{\varepsilon_i\}_{i=1}^{m+1}, \{\delta_j\}_{j=1}^{n+1} \subseteq \{-1, 1\}$ define $S^+, S^- \subseteq \{1, \dots, m+1\}$ and $T^+, T^- \subseteq \{1, \dots, n+1\}$ by setting $S^\pm = \{i \in \{1, \dots, m+1\} : \varepsilon_i = \pm 1\}$ and $T^\pm = \{j \in \{1, \dots, n+1\} : \delta_j = \pm 1\}$. Then

$$\sum_{i=1}^{m+1} \sum_{j=1}^{n+1} b_{ij} \varepsilon_i \delta_j = \sum_{\substack{i \in S^+ \\ j \in T^+}} b_{ij} + \sum_{\substack{i \in S^- \\ j \in T^-}} b_{ij} - \sum_{\substack{i \in S^+ \\ j \in T^-}} b_{ij} - \sum_{\substack{i \in S^- \\ j \in T^+}} b_{ij} \leq 4\|B\|_{cut}. \quad (11)$$

This shows that $\|B\|_{\infty \rightarrow 1} \leq 4\|B\|_{cut}$ (for any matrix B , actually, not just the specific choice in (7); we will use this observation later, in Section 2.1.3). In the reverse direction, given $S \subseteq \{1, \dots, m+1\}$ and $T \subseteq \{1, \dots, n+1\}$ define for $i \in \{1, \dots, m+1\}$ and $j \in \{1, \dots, n+1\}$,

$$\varepsilon_i = \begin{cases} 1 & \text{if } i \in S, \\ -1 & \text{if } i \notin S, \end{cases} \quad \text{and} \quad \delta_j = \begin{cases} 1 & \text{if } j \in T, \\ -1 & \text{if } j \notin T. \end{cases}$$

Then, since the sum of each row and each column of B vanishes,

$$\sum_{\substack{i \in S \\ j \in T}} b_{ij} = \sum_{i=1}^{m+1} \sum_{j=1}^{n+1} b_{ij} \frac{1 + \varepsilon_i}{2} \cdot \frac{1 + \delta_j}{2} = \frac{1}{4} \sum_{i=1}^{m+1} \sum_{j=1}^{n+1} b_{ij} \varepsilon_i \delta_j \leq \frac{1}{4} \|B\|_{\infty \rightarrow 1}.$$

This completes the proof of (9). We summarize the above simple transformations in the following lemma.

Lemma 2.1. *Let $A = (a_{ij})$ be an $m \times n$ matrix with real entries and let $B = (b_{ij})$ be the $(m+1) \times (n+1)$ matrix given in (7). Then*

$$\|A\|_{cut} = \frac{1}{4} \|B\|_{\infty \rightarrow 1}.$$

A consequence of Lemma 2.1 is that the problem of approximating $\|A\|_{cut}$ in polynomial time is equivalent to the problem of approximating $\|A\|_{\infty \rightarrow 1}$ in polynomial time in the sense that any algorithm for one of these problems can be used to obtain an algorithm for the other problem with the same running time (up to constant factors) and the same (multiplicative) approximation guarantee.

Given an $m \times n$ matrix $A = (a_{ij})$ consider the following quantity.

$$\text{SDP}(A) = \max \left\{ \sum_{i=1}^m \sum_{j=1}^n a_{ij} \langle x_i, y_j \rangle : \{x_i\}_{i=1}^m, \{y_j\}_{j=1}^n \subseteq S^{m+n-1} \right\}. \quad (12)$$

The maximization problem in (12) falls into the framework of semidefinite programming as discussed in Section 1.2. Therefore $\text{SDP}(A)$ can be computed in polynomial time with arbitrarily good precision. It is clear that $\text{SDP}(A) \geq \|A\|_{\infty \rightarrow 1}$, because the maximum in (12) is over a bigger set than the maximum in (10). The Grothendieck inequality says that $\text{SDP}(A) \leq K_G \|A\|_{\infty \rightarrow 1}$, so we have

$$\|A\|_{\infty \rightarrow 1} \leq \text{SDP}(A) \leq K_G \|A\|_{\infty \rightarrow 1}.$$

Thus, the polynomial time algorithm that outputs the number $\text{SDP}(A)$ is guaranteed to be within a factor of K_G of $\|A\|_{\infty \rightarrow 1}$. By Lemma 2.1, the algorithm that outputs the number $\alpha = \frac{1}{4} \text{SDP}(B)$, where the matrix B is as in (7), satisfies (5) with $C = K_G$.

Section 7 is devoted to algorithmic impossibility results. But, it is worthwhile to make at this juncture two comments regarding hardness of approximation. First of all, unless

$P = NP$, we need to introduce an error $C > 1$ in our requirement (5). This was observed in [8]: the classical MAXCUT problem from algorithmic graph theory was shown in [8] to be a special case of the problem of computing $\|A\|_{cut}$, and therefore by [51] we know that unless $P = NP$ there does not exist a polynomial time algorithm that outputs a number α satisfying (5) with C strictly smaller than $\frac{17}{16}$. In fact, by a reduction to the MAX DICUT problem one can show that C must be at least $\frac{13}{12}$, unless $P = NP$; we refer to Section 7 and [8] for more information on this topic.

Another (more striking) algorithmic impossibility result is based on the Unique Games Conjecture (UGC). Clearly the above algorithm cannot yield an approximation guarantee strictly smaller than K_G (this is the definition of K_G). In fact, it was shown in [104] that unless the UGC is false, for every $\varepsilon \in (0, 1)$ any polynomial time algorithm for estimating $\|A\|_{cut}$ whatsoever, and not only the specific algorithm described above, must make an error of at least $K_G - \varepsilon$ on some input matrix A . Thus, if we assume the UGC then the classical Grothendieck constant has a complexity theoretic interpretation: it equals the best approximation ratio of polynomial time algorithms for the cut norm problem. Note that [104] manages to prove this statement despite the fact that the value of K_G is unknown.

We have thus far ignored the issue of finding in polynomial time the subsets S_0, T_0 satisfying (6), i.e., we only explained how the Grothendieck inequality can be used for polynomial time estimation of the quantity $\|A\|_{cut}$ without actually finding efficiently subsets at which $\|A\|_{cut}$ is approximately attained. In order to do this we cannot use the Grothendieck inequality as a black box: we need to look into its proof and argue that it yields a polynomial time procedure that converts vectors $\{x_i\}_{i=1}^m, \{y_j\}_{j=1}^n \subseteq S^{m+m-1}$ into signs $\{\varepsilon_i\}_{i=1}^m, \{\delta_j\}_{j=1}^n \subseteq \{-1, 1\}$ (this is known as a rounding procedure). It is indeed possible to do so, as explained in Section 2.2. We postpone the explanation of the rounding procedure that hides behind the Grothendieck inequality in order to first give examples why one might want to efficiently compute the cut norm of a matrix.

2.1.1. Szemerédi partitions. The Szemerédi regularity lemma [111] (see also [72]) is a general and very useful structure theorem for graphs, asserting (informally) that any graph can be partitioned into a controlled number of pieces that interact with each other in a pseudo-random way. The Grothendieck inequality, via the cut norm estimation algorithm, yields a polynomial time algorithm that, when given a graph $G = (V, E)$ as input, outputs a partition of V that satisfies the conclusion of the Szemerédi regularity lemma.

To make the above statements formal, we need to recall some definitions. Let $G = (V, E)$ be a graph. For every disjoint $X, Y \subseteq V$ denote the number of edges joining X and Y by $e(X, Y) = |\{(u, v) \in X \times Y : \{u, v\} \in E\}|$. Let $X, Y \subseteq V$ be disjoint and nonempty, and fix $\varepsilon, \delta \in (0, 1)$. The pair of vertex sets (X, Y) is called (ε, δ) -regular if for every $S \subseteq X$ and $T \subseteq Y$ that are not too small, the quantity $\frac{e(S, T)}{|S| \cdot |T|}$ (the density of edges between S and T) is essentially independent of the pair (S, T) itself. Formally, we require that for every $S \subseteq X$ with $|S| \geq \delta|X|$ and every $T \subseteq Y$ with $|T| \geq \delta|Y|$ we have

$$\left| \frac{e(S, T)}{|S| \cdot |T|} - \frac{e(X, Y)}{|X| \cdot |Y|} \right| \leq \varepsilon. \quad (13)$$

The almost uniformity of the numbers $\frac{e(S,T)}{|S| \cdot |T|}$ as exhibited in (13) says that the pair (X, Y) is “pseudo-random”, i.e., it is similar to a random bipartite graph where each $(x, y) \in X \times Y$ is joined by an edge independently with probability $\frac{e(X,Y)}{|X| \cdot |Y|}$.

The Szemerédi regularity lemma says that for all $\varepsilon, \delta, \eta \in (0, 1)$ and $k \in \mathbb{N}$ there exists $K = K(\varepsilon, \delta, \eta, k) \in \mathbb{N}$ such that for all $n \in \mathbb{N}$ any n -vertex graph $G = (V, E)$ can be partitioned into m -sets $S_1, \dots, S_m \subseteq V$ with the following properties

- $k \leq m \leq K$,
- $|S_i| - |S_j| \leq 1$ for all $i, j \in \{1, \dots, m\}$,
- the number of $i, j \in \{1, \dots, m\}$ with $i < j$ such that the pair (S_i, S_j) is (ε, δ) -regular is at least $(1 - \eta) \binom{m}{2}$.

Thus every graph is almost a superposition of a bounded number of pseudo-random graphs, the key point being that K is independent of n and the specific combinatorial structure of the graph in question.

It would be of interest to have a way to produce a Szemerédi partition in polynomial time with K independent of n (this is a good example of an approximation algorithm: one might care to find such a partition into the minimum possible number of pieces, but producing any partition into boundedly many pieces is already a significant achievement). Such a polynomial time algorithm was designed in [5] (see also [73]). We refer to [5, 73] for applications of algorithms for constructing Szemerédi partitions, and to [5] for a discussion of the computational complexity of this algorithmic task. We shall now explain how the Grothendieck inequality yields a different approach to this problem, which has some advantages over [5, 73] that will be described later. The argument below is due to [8].

Assume that X, Y are disjoint n -point subsets of a graph $G = (V, E)$. How can we determine in polynomial time whether or not the pair (X, Y) is close to being (ε, δ) -regular? It turns out that this is the main “bottleneck” towards our goal to construct Szemerédi partitions in polynomial time. To this end consider the following $n \times n$ matrix $A = (a_{xy})_{(x,y) \in X \times Y}$.

$$a_{xy} = \begin{cases} 1 - \frac{e(X,Y)}{|X| \cdot |Y|} & \text{if } \{x, y\} \in E, \\ -\frac{e(X,Y)}{|X| \cdot |Y|} & \text{if } \{x, y\} \notin E. \end{cases} \quad (14)$$

By the definition of A , if $S \subseteq X$ and $T \subseteq Y$ then

$$\left| \sum_{\substack{x \in S \\ y \in T}} a_{xy} \right| = |S| \cdot |T| \cdot \left| \frac{e(S,T)}{|S| \cdot |T|} - \frac{e(X,Y)}{|X| \cdot |Y|} \right|. \quad (15)$$

Hence if (X, Y) is not (ε, δ) -regular then $\|A\|_{cut} \geq \varepsilon \delta^2 n^2$. The approximate cut norm algorithm based on the Grothendieck inequality, together with the rounding procedure in Section 2.2, finds in polynomial time subsets $S \subseteq X$ and $T \subseteq Y$ such that

$$\min \left\{ n|S|, n|T|, n^2 \left| \frac{e(S,T)}{|S| \cdot |T|} - \frac{e(X,Y)}{|X| \cdot |Y|} \right| \right\} \stackrel{(15)}{\geq} \left| \sum_{\substack{x \in S \\ y \in T}} a_{xy} \right| \geq \frac{1}{K_G} \varepsilon \delta^2 n^2 \geq \frac{1}{2} \varepsilon \delta^2 n^2.$$

This establishes the following lemma.

Lemma 2.2. *There exists a polynomial time algorithm that takes as input two disjoint n -point subsets X, Y of a graph, and either decides that (X, Y) is (ε, δ) -regular or finds $S \subseteq X$ and $T \subseteq Y$ with*

$$|S|, |T| \geq \frac{1}{2} \varepsilon \delta^2 n \quad \text{and} \quad \left| \frac{e(S, T)}{|S| \cdot |T|} - \frac{e(X, Y)}{|X| \cdot |Y|} \right| \geq \frac{1}{2} \varepsilon \delta^2.$$

From Lemma 2.2 it is quite simple to design a polynomial algorithm that constructs a Szemerédi partition with bounded cardinality; compare Lemma 2.2 to Corollary 3.3 in [5] and Theorem 1.5 in [73]. We will not explain this deduction here since it is identical to the argument in [5]. We note that the quantitative bounds in Lemma 2.2 improve over the corresponding bounds in [5, 73] yielding, say, when $\varepsilon = \delta = \eta$, an algorithm with the best known bound on K as a function of ε (this bound is nevertheless still huge, as must be the case due to [44]; see also [30]). See [8] for a precise statement of these bounds. In addition, the algorithms of [5, 73] worked only in the “dense case”, i.e., when $\|A\|_{cut}$, for A as in (14), is of order n^2 , while the above algorithm does not have this requirement. This observation can be used to design the only known polynomial time algorithm for sparse versions of the Szemerédi regularity lemma [4] (see also [41]). We will not discuss the sparse version of the regularity lemma here, and refer instead to [71, 72] for a discussion of this topic. We also refer to [4] for additional applications of the Grothendieck inequality in sparse settings.

2.1.2. Frieze-Kannan matrix decomposition. The cut norm estimation problem was originally raised in the work of Frieze and Kannan [38] which introduced a method to design polynomial time approximation schemes for dense constraint satisfaction problems. The key tool for this purpose is a decomposition theorem for matrices that we now describe.

An $m \times n$ matrix $D = (d_{ij})$ is called a cut matrix if there exist subsets $S \subseteq \{1, \dots, m\}$ and $T \subseteq \{1, \dots, n\}$, and $d \in \mathbb{R}$ such that for all $(i, j) \in \{1, \dots, m\} \times \{1, \dots, n\}$ we have,

$$d_{ij} = \begin{cases} d & \text{if } (i, j) \in S \times T, \\ 0 & \text{if } (i, j) \notin S \times T. \end{cases} \quad (16)$$

Denote the matrix D defined in (16) by $CUT(S, T, d)$. In [38] it is proved that for every $\varepsilon > 0$ there exists an integer $s = O(1/\varepsilon^2)$ such that for any $m \times n$ matrix $A = (a_{ij})$ with entries bounded in absolute value by 1, there are cut matrices D_1, \dots, D_s satisfying

$$\left\| A - \sum_{k=1}^s D_k \right\|_{cut} \leq \varepsilon mn. \quad (17)$$

Moreover, these cut matrices D_1, \dots, D_s can be found in time $C(\varepsilon)(mn)^{O(1)}$. We shall now explain how this is done using the cut norm approximation algorithm of Section 2.1.

The argument is iterative. Set $A_0 = A$, and assuming that the cut matrices D_1, \dots, D_r have already been defined write $A_r = (a_{ij}(r)) = A - \sum_{k=1}^r D_k$. We are done if $\|A_r\|_{cut} \leq \varepsilon mn$, so we may assume that $\|A_r\|_{cut} > \varepsilon mn$. By the cut norm approximation algorithm we can find in polynomial time $S \subseteq \{1, \dots, m\}$ and $T \subseteq \{1, \dots, n\}$ satisfying

$$\left| \sum_{\substack{i \in S \\ j \in T}} a_{ij}(r) \right| \geq c \|A_r\|_{cut} \geq c \varepsilon mn, \quad (18)$$

where $c > 0$ is a universal constant. Set

$$d = \frac{1}{|S| \cdot |T|} \sum_{\substack{i \in S \\ j \in T}} a_{ij}(r).$$

Define $D_{r+1} = CUT(S, T, d)$ and $A_{r+1} = (a_{ij}(r+1)) = A_r - D_{r+1}$. Then by expanding the squares we have,

$$\sum_{i=1}^m \sum_{j=1}^n a_{ij}(r+1)^2 = \sum_{i=1}^m \sum_{j=1}^n a_{ij}(r)^2 - \frac{1}{|S| \cdot |T|} \left(\sum_{\substack{i \in S \\ j \in T}} a_{ij}(r) \right)^2 \stackrel{(18)}{\leq} \sum_{i=1}^m \sum_{j=1}^n a_{ij}(r)^2 - c^2 \varepsilon^2 mn.$$

It follows inductively that if we can carry out this procedure r times then

$$0 \leq \sum_{i=1}^m \sum_{j=1}^n a_{ij}(r)^2 \leq \sum_{i=1}^m \sum_{j=1}^n a_{ij}^2 - rc^2 \varepsilon^2 mn \leq mn - rc^2 \varepsilon^2 mn,$$

where we used the assumption that $|a_{ij}| \leq 1$. Therefore the above iteration must terminate after $\lceil 1/(c^2 \varepsilon^2) \rceil$ steps, yielding (17). We note that the bound $s = O(1/\varepsilon^2)$ in (17) cannot be improved [6]; see also [89, 30] for related lower bounds.

The key step in the above algorithm was finding sets S, T as in (18). In [38] an algorithm was designed that, given an $m \times n$ matrix $A = (a_{ij})$ and $\varepsilon > 0$ as input, produces in time $2^{1/\varepsilon^{O(1)}} (mn)^{O(1)}$ subsets $S \subseteq \{1, \dots, m\}$ and $T \subseteq \{1, \dots, n\}$ satisfying

$$\left| \sum_{\substack{i \in S \\ j \in T}} a_{ij} \right| \geq \|A\|_{cut} - \varepsilon mn. \quad (19)$$

The additive approximation guarantee in (19) implies (18) only if $\|A\|_{cut} \geq \varepsilon(c+1)mn$, and similarly the running time is not polynomial if, say, $\varepsilon = n^{-\Omega(1)}$. Thus the Kannan-Frieze method is relevant only to “dense” instances, while the cut norm algorithm based on the Grothendieck inequality applies equally well for all values of $\|A\|_{cut}$. This fact, combined with more work (and, necessarily, additional assumptions on the matrix A), was used in [29] to obtain a sparse version of (17): with εmn in the right hand side of (17) replaced by $\varepsilon \|A\|_{cut}$ and $s = O(1/\varepsilon^2)$ (importantly, here s is independent of m, n).

We have indicated above how the cut norm approximation problem is relevant to Kannan-Frieze matrix decompositions, but we did not indicate the uses of such decompositions since this is beyond the scope of the current survey. We refer to [38, 6, 15, 29] for a variety of applications of this methodology to combinatorial optimization problems.

2.1.3. Maximum acyclic subgraph. In the maximum acyclic subgraph problem we are given as input an n -vertex directed graph $G = (\{1, \dots, n\}, E)$. Thus E consists of a family of *ordered* pairs of distinct elements in $\{1, \dots, n\}$. We are interested in the maximum of

$$\left| \{(i, j) \in \{1, \dots, n\}^2 : \sigma(i) < \sigma(j)\} \cap E \right| - \left| \{(i, j) \in \{1, \dots, n\}^2 : \sigma(i) > \sigma(j)\} \cap E \right|$$

over all possible permutations $\sigma \in S_n$ (S_n denotes the group of permutations of $\{1, \dots, n\}$). In words, the quantity of interest is the maximum over all orderings of the vertices of the number of edges going “forward” minus the number of edges going “backward”. Note that it is trivial to get at least half of the edges to go forward by considering a random permutation,

so in essence we are measuring here the advantage of the best possible ordering over a random ordering. The best known approximation algorithm for this problem was discovered in [26] as an application of the cut norm approximation algorithm.

It is most natural to explain the algorithm of [26] for a weighted version of the maximum acyclic subgraph problem. Let $W : \{1, \dots, n\} \times \{1, \dots, n\} \rightarrow \mathbb{R}$ be skew symmetric, i.e., $W(u, v) = -W(v, u)$ for all $u, v \in \{1, \dots, n\}$. For $\sigma \in S_n$ define

$$W(\sigma) = \sum_{\substack{u, v \in \{1, \dots, n\} \\ u < v}} W(\sigma(u), \sigma(v)).$$

Thus $W(\sigma)$ is the sum of the entries of W that lie above the diagonal after the rows and columns of W have been permuted according to the permutation σ . We are interested in the quantity $M_W = \max_{\sigma \in S_n} W(\sigma)$. The case of a directed graph $G = (\{1, \dots, n\}, E)$ described above corresponds to the matrix $W(u, v) = \mathbf{1}_{\{(u, v) \in E\}} - \mathbf{1}_{\{(v, u) \in E\}}$.

Theorem 2.3 ([26]). *The exists a polynomial time algorithm that takes as input an $n \times n$ skew symmetric $W : \{1, \dots, n\} \times \{1, \dots, n\} \rightarrow \mathbb{R}$ and outputs a permutation $\sigma \in S_n$ satisfying¹*

$$W(\sigma) \gtrsim \frac{M_W}{\log n}.$$

Proof. The proof below is a slight variant of the reasoning of [26]. By the cut norm approximation algorithm one can find in polynomial time two subsets $S, T \subseteq \{1, \dots, n\}$ satisfying

$$\sum_{\substack{u \in S \\ v \in T}} W(u, v) \geq c \|W\|_{cut}, \quad (20)$$

where $c \in (0, \infty)$ is a universal constant. Note that we do not need to take the absolute value of the left hand side of (20) because W is skew symmetric. Observe also that since W is skew symmetric we have $\sum_{u, v \in S \cap T} W(u, v) = 0$ and therefore

$$\sum_{\substack{u \in S \\ v \in T}} W(u, v) = \sum_{\substack{u \in S \setminus T \\ v \in T \setminus S}} W(u, v) + \sum_{\substack{u \in S \setminus T \\ v \in S \cap T}} W(u, v) + \sum_{\substack{u \in S \cap T \\ v \in T \setminus S}} W(u, v).$$

By replacing the pair of subsets (S, T) by one of $\{(S \setminus T, T \setminus S), (S \setminus T, S \cap T), (S \cap T, T \setminus S)\}$, and replacing the constant c in (20) by $c/3$, we may assume without loss of generality that (20) holds with S and T disjoint. Denote $R = \{1, \dots, n\} \setminus (S \cup T)$ and write $S = \{s_1, \dots, s_{|S|}\}$, $T = \{t_1, \dots, t_{|T|}\}$ and $R = \{r_1, \dots, r_{|R|}\}$, where $s_1 < \dots < s_{|S|}$, $t_1 < \dots < t_{|T|}$ and $r_1 < \dots < r_{|R|}$.

Define two permutations $\sigma^1, \sigma^2 \in S_n$ as follows.

$$\sigma^1(u) = \begin{cases} s_u & \text{if } u \in \{1, \dots, |S|\}, \\ t_{u-|S|} & \text{if } u \in \{|S| + 1, \dots, |S| + |T|\}, \\ r_{u-|S|-|T|} & \text{if } u \in \{|S| + |T| + 1, \dots, n\}, \end{cases}$$

and

$$\sigma^2(u) = \begin{cases} r_{|R|-u+1} & \text{if } u \in \{1, \dots, |R|\}, \\ s_{|R|+|S|-u+1} & \text{if } u \in \{|R| + 1, \dots, |R| + |S|\}, \\ t_{n-u+1} & \text{if } u \in \{|R| + |S| + 1, \dots, n\}. \end{cases}$$

¹Here, and in what follows, the relations \gtrsim, \lesssim indicate the corresponding inequalities up to an absolute factor. The relation \asymp stands for $\gtrsim \wedge \lesssim$.

In words, σ^1 orders $\{1, \dots, n\}$ by starting with the elements of S in increasing order, then the elements of T in increasing order, and finally the elements of R in increasing order. At the same time, σ^2 orders $\{1, \dots, n\}$ by starting with the elements of R in decreasing order, then the elements of S in decreasing order, and finally the elements of T in decreasing order. The quantity $W(\sigma^1) + W(\sigma^2)$ consists of a sum of terms of the form $W(u, v)$ for $u, v \in \{1, \dots, n\}$, where if $(u, v) \in (S \times S) \cup (T \times T) \cup (R \times \{1, \dots, n\})$ then both $W(u, v)$ and $W(v, u)$ appear exactly once in this sum, and if $(u, v) \in S \times T$ then $W(u, v)$ appears twice in this sum and $W(v, u)$ does not appear in this sum at all. Therefore, using the fact that W is skew symmetric we have the following identity.

$$W(\sigma^1) + W(\sigma^2) = 2 \sum_{\substack{u \in S \\ v \in T}} W(u, v).$$

It follows that for some $\ell \in \{1, 2\}$ we have

$$M(\sigma^\ell) \geq \sum_{\substack{u \in S \\ v \in T}} W(u, v) \stackrel{(20)}{\geq} c \|W\|_{cut}.$$

The output of the algorithm will be the permutation σ^ℓ , so it suffices to prove that

$$\|W\|_{cut} \gtrsim \frac{M_W}{\log n}. \quad (21)$$

We will prove below that

$$\|W\|_{cut} \gtrsim \frac{1}{\log n} \sum_{\substack{u, v \in \{1, \dots, n\} \\ u < v}} W(u, v). \quad (22)$$

Inequality (21) follows by applying (22) to $W'(u, v) = W(\sigma(u), \sigma(v))$ for every $\sigma \in S_n$.

To prove (22) first note that $\|W\|_{cut} \geq \frac{1}{4} \|W\|_{\infty \rightarrow 1}$; we have already proved this inequality as a consequence of the simple identity (11). Moreover, we have

$$\|W\|_{\infty \rightarrow 1} \gtrsim \max \left\{ \sum_{u=1}^n \sum_{v=1}^n W(u, v) \sin(\alpha_u - \beta_v) : \{\alpha_u\}_{u=1}^n, \{\beta_v\}_{v=1}^n \subseteq \mathbb{R} \right\}. \quad (23)$$

Inequality (23) is a special case of (1) with the choice of vectors $x_u = (\sin \alpha_u, \cos \alpha_u) \in \mathbb{R}^2$ and $y_v = (\cos \beta_v, -\sin \beta_v) \in \mathbb{R}^2$. We note that this two-dimensional version of the Grothendieck inequality is trivial with the constant in the right hand side of (23) being $\frac{1}{2}$, and it is shown in [78] that the best constant in the right hand side of (23) is actually $\frac{1}{\sqrt{2}}$.

For every $\theta_1, \dots, \theta_n \in \mathbb{R}$, an application of (23) when $\alpha_u = \beta_u = \theta_u$ and $\alpha_u = \beta_u = -\theta_u$ yields the inequality

$$\|W\|_{cut} \gtrsim \left| \sum_{u=1}^n \sum_{v=1}^n W(u, v) \sin(\theta_u - \theta_v) \right| = 2 \left| \sum_{\substack{u, v \in \{1, \dots, n\} \\ u < v}} W(u, v) \sin(\theta_u - \theta_v) \right|, \quad (24)$$

where for the equality in (24) we used the fact that W is skew symmetric. Consequently, for every $k \in \mathbb{N}$ we have

$$\|W\|_{cut} \gtrsim \left| \sum_{\substack{u,v \in \{1, \dots, n\} \\ u < v}} W(u, v) \sin\left(\frac{\pi(v-u)k}{n}\right) \right|. \quad (25)$$

By the standard orthogonality relation for the sine function, for every $u, v \in \{1, \dots, n\}$ such that $u < v$ we have

$$\frac{2}{n} \sum_{k=1}^{n-1} \sum_{\ell=1}^{n-1} \sin\left(\frac{\pi(v-u)k}{n}\right) \sin\left(\frac{\pi k \ell}{n}\right) = 1. \quad (26)$$

Readers who are unfamiliar with (26) are referred to its derivation in the appendix of [26]; it can be proved by substituting $\sin(\pi(v-u)k/n) = (e^{i\pi(v-u)k/n} - e^{-i\pi(v-u)k/n})/(2i)$ and $\sin(\pi k \ell/n) = (e^{i\pi k \ell/n} - e^{-i\pi k \ell/n})/(2i)$ into the left hand side of (26) and computing the resulting geometric sums explicitly. Now,

$$\begin{aligned} \sum_{\substack{u,v \in \{1, \dots, n\} \\ u < v}} W(u, v) &\stackrel{(26)}{=} \frac{2}{n} \sum_{\substack{u,v \in \{1, \dots, n\} \\ u < v}} W(u, v) \sum_{k=1}^{n-1} \sum_{\ell=1}^{n-1} \sin\left(\frac{\pi(v-u)k}{n}\right) \sin\left(\frac{\pi k \ell}{n}\right) \\ &\leq \frac{2}{n} \sum_{k=1}^{n-1} \left| \sum_{\ell=1}^{n-1} \sin\left(\frac{\pi k \ell}{n}\right) \right| \cdot \left| \sum_{\substack{u,v \in \{1, \dots, n\} \\ u < v}} W(u, v) \sin\left(\frac{\pi(v-u)k}{n}\right) \right| \\ &\stackrel{(25)}{\lesssim} \frac{\sum_{k=1}^{n-1} \left| \sum_{\ell=1}^{n-1} \sin\left(\frac{\pi k \ell}{n}\right) \right|}{n} \|W\|_{cut}. \end{aligned}$$

Hence, the desired inequality (22) will follow from $\sum_{k=1}^{n-1} \left| \sum_{\ell=1}^{n-1} \sin(\pi k \ell/n) \right| \lesssim n \log n$. To establish this estimate observe that by writing $\sin(\pi k \ell/n) = (e^{i\pi k \ell/n} - e^{-i\pi k \ell/n})/(2i)$ and computing geometric sums explicitly, one sees that $\sum_{\ell=1}^{n-1} \sin(\pi k \ell/n) = 0$ if k is even and $\sum_{\ell=1}^{n-1} \sin(\pi k \ell/n) = \cot(\pi k/(2n))$ if k is odd (see the appendix of [26] for the details of this computation). Hence, since $\cot(\theta) < 1/\theta$ for every $\theta \in (0, \pi/2)$, we have

$$\sum_{k=1}^{n-1} \left| \sum_{\ell=1}^{n-1} \sin\left(\frac{\pi k \ell}{n}\right) \right| = \sum_{j=0}^{\lfloor \frac{n-1}{2} \rfloor} \cot\left(\frac{\pi(2j+1)}{2n}\right) \leq \frac{2n}{\pi} \sum_{j=0}^{\lfloor \frac{n-1}{2} \rfloor} \frac{1}{2j+1} \lesssim n \log n \quad \square$$

2.1.4. Linear equations modulo 2. Consider a system \mathcal{E} of N linear equations modulo 2 in n Boolean variables z_1, \dots, z_n such that in each equation appear only three distinct variables. Let $\text{MAXSAT}(\mathcal{E})$ be the maximum number of equations in \mathcal{E} that can be satisfied simultaneously. A random $\{0, 1\}$ assignment of these variables satisfies in expectation $N/2$ equations, so it is natural to ask for a polynomial time approximation algorithm to the quantity $\text{MAXSAT}(\mathcal{E}) - N/2$. We describe below the best known [65] approximation algorithm for this problem, which uses the Grothendieck inequality in a crucial way. The approximation guarantee thus obtained is $O(\sqrt{n/\log n})$. While this allows for a large error, it is shown in [52] that for every $\varepsilon \in (0, 1)$ if there were a polynomial time algorithm that approximates $\text{MAXSAT}(\mathcal{E}) - N/2$ to within a factor of $2^{(\log n)^{1-\varepsilon}}$ in time $2^{(\log n)^{O(1)}}$ then there would be an

algorithm for 3-colorability that runs in time $2^{(\log n)^{O(1)}}$, a conclusion which is widely believed to be impossible.

Let \mathcal{E} be a system of linear equations as described above. Write $a_{ijk} = 1$ if the equation $z_i + z_j + z_k = 0$ is in the system \mathcal{E} . Similarly write $a_{ijk} = -1$ if the equation $z_i + z_j + z_k = 1$ is in \mathcal{E} . Finally, write $a_{ijk} = 0$ if no equation in \mathcal{E} corresponds to $z_i + z_j + z_k$. Assume that the assignment $(z_1, \dots, z_n) \in \{0, 1\}^n$ satisfies m of the equations in \mathcal{E} . Then

$$\sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n a_{ijk} (-1)^{z_i+z_j+z_k} = m - (N - m) = 2 \left(m - \frac{N}{2} \right).$$

It follows that

$$\max \left\{ \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n a_{ijk} \varepsilon_i \varepsilon_j \varepsilon_k : \{\varepsilon_i\}_{i=1}^n \subseteq \{-1, 1\} \right\} = 2 \left(\text{MAXSAT}(\mathcal{E}) - \frac{N}{2} \right) \stackrel{\text{def}}{=} M. \quad (27)$$

We will now present a randomized polynomial algorithm that outputs a number $\alpha \in \mathbb{R}$ which satisfies with probability at least $\frac{2}{3}$,

$$\frac{1}{20K_G} \sqrt{\frac{\log n}{n}} M \leq \alpha \leq M. \quad (28)$$

Fix $m \in \mathbb{N}$ that will be determined later. Choose $\varepsilon^1, \dots, \varepsilon^m \in \{-1, 1\}^n$ independently and uniformly at random and consider the following random variable.

$$\alpha = \frac{1}{10K_G} \max_{\ell \in \{1, \dots, m\}} \max \left\{ \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n a_{ijk} \varepsilon_i^\ell \langle y_j, z_k \rangle : \{y_j\}_{j=1}^n, \{z_k\}_{k=1}^n \subseteq S^{2n-1} \right\}. \quad (29)$$

By the Grothendieck inequality we know that

$$\alpha \leq \frac{1}{10} \max \left\{ \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n a_{ijk} \varepsilon_i \delta_j \zeta_k : \{\varepsilon_i\}_{i=1}^n, \{\delta_j\}_{j=1}^n, \{\zeta_k\}_{k=1}^n \subseteq \{-1, 1\} \right\} \leq M. \quad (30)$$

The final step in (30) follows from an elementary decoupling argument; see [65, Lem. 2.1].

We claim that

$$\Pr \left[\alpha \geq \frac{1}{20K_G} \sqrt{\frac{\log n}{n}} M \right] \geq 1 - e^{-cm/\sqrt[4]{n}}. \quad (31)$$

Once (31) is established, it would follow that for $m \asymp \sqrt[4]{n}$ we have $\alpha \geq \frac{1}{20K_G} \sqrt{\frac{\log n}{n}} M$ with probability at least $\frac{2}{3}$. This combined with (30) would complete the proof of (28) since α as defined in (29) can be computed in polynomial time, being the maximum of $O(\sqrt[4]{n})$ semidefinite programs.

To check (31) let $\|\cdot\|$ be the norm on \mathbb{R}^n defined for every $x = (x_1, \dots, x_n) \in \mathbb{R}^n$ by

$$\|x\| = \max \left\{ \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n a_{ijk} x_i \langle y_j, z_k \rangle : \{y_j\}_{j=1}^n, \{z_k\}_{k=1}^n \subseteq S^{2n-1} \right\}.$$

Define $K = \{x \in \mathbb{R}^n : \|x\| \leq 1\}$ and let $K^\circ = \{w \in \mathbb{R}^n : \sup_{x \in K} \langle x, w \rangle \leq 1\}$ be the polar of K . Then $\max\{\|w\|_1 : w \in K^\circ\} = \max\{\|x\| : \|x\|_\infty \leq 1\} \geq M$, where the first equality is straightforward duality and the final inequality is a consequence of the definition

of $\|\cdot\|$ and M . It follows that there exists $w \in K^\circ$ with $\|w\|_1 \geq M$. Hence, recalling that $\alpha = \frac{1}{10K_G} \max_{\ell \in \{1, \dots, m\}} \|\varepsilon^\ell\|$, we have

$$\begin{aligned} \Pr \left[\alpha \geq \frac{1}{20K_G} \sqrt{\frac{\log n}{n}} M \right] &\stackrel{(29)}{=} 1 - \prod_{\ell=1}^m \Pr \left[\|\varepsilon^\ell\| < \frac{1}{2} \sqrt{\frac{\log n}{n}} M \right] \\ &\geq 1 - \left(\Pr \left[\sum_{i=1}^n \varepsilon_i^1 w_i < \frac{1}{2} \sqrt{\frac{\log n}{n}} \sum_{i=1}^n |w_i| \right] \right)^m. \end{aligned}$$

In order to prove (31) it therefore suffices to prove that if ε is chosen uniformly at random from $\{-1, 1\}^n$ and $a \in \mathbb{R}^n$ satisfies $\|a\|_1 = 1$ then $\Pr \left[\sum_{i=1}^n \varepsilon_i a_i \geq \sqrt{\log n / (4n)} \right] \geq 1 - c / \sqrt[4]{n}$, where $c \in (0, \infty)$ is a universal constant. This probabilistic estimate for i.i.d. Bernoulli sums can be proved directly; see [65, Lem. 3.2].

2.2. Rounding. Let $A = (a_{ij})$ be an $m \times n$ matrix. In Section 2.1 we described a polynomial time algorithm for approximating $\|A\|_{cut}$ and $\|A\|_{\infty \rightarrow 1}$. For applications it is also important to find in polynomial time signs $\varepsilon_1, \dots, \varepsilon_m, \delta_1, \dots, \delta_n \in \{-1, 1\}$ for which $\sum_{i=1}^m \sum_{j=1}^n a_{ij} \varepsilon_i \delta_j$ is at least a constant multiple of $\|A\|_{\infty \rightarrow 1}$. This amounts to a ‘‘rounding problem’’: we need to find a procedure that, given vectors $x_1, \dots, x_m, y_1, \dots, y_n \in S^{m+n-1}$, produces signs $\varepsilon_1, \dots, \varepsilon_m, \delta_1, \dots, \delta_n \in \{-1, 1\}$ whose existence is ensured by the Grothendieck inequality, i.e., $\sum_{i=1}^m \sum_{j=1}^n a_{ij} \varepsilon_i \delta_j$ is at least a constant multiple of $\sum_{i=1}^m \sum_{j=1}^n a_{ij} \langle x_i, y_j \rangle$. For this purpose one needs to examine proofs of the Grothendieck inequality, as done in [8]. We will now describe the rounding procedure that gives the best known approximation guarantee. This procedure yields a randomized algorithm that produces the desired signs; it is also possible to obtain a deterministic algorithm, as explained in [8].

The argument below is based on a clever two-step rounding method due to Krivine [77]. Fix $k \in \mathbb{N}$ and assume that we are given two centrally symmetric measurable partitions of \mathbb{R}^k , or equivalently two odd measurable functions $f, g : \mathbb{R}^k \rightarrow \{-1, 1\}$. Let $G_1, G_2 \in \mathbb{R}^k$ be independent random vectors that are distributed according to the standard Gaussian measure on \mathbb{R}^k , i.e., the measure with density $x \mapsto e^{-\|x\|_2^2/2} / (2\pi)^{k/2}$. For $t \in (-1, 1)$ define

$$\begin{aligned} H_{f,g}(t) &\stackrel{\text{def}}{=} \mathbb{E} \left[f \left(\frac{1}{\sqrt{2}} G_1 \right) g \left(\frac{t}{\sqrt{2}} G_1 + \frac{\sqrt{1-t^2}}{\sqrt{2}} G_2 \right) \right] \\ &= \frac{1}{\pi^k (1-t^2)^{k/2}} \int_{\mathbb{R}^k} \int_{\mathbb{R}^k} f(x) g(y) \exp \left(\frac{-\|x\|_2^2 - \|y\|_2^2 + 2t \langle x, y \rangle}{1-t^2} \right) dx dy. \end{aligned} \quad (32)$$

Then $H_{f,g}$ extends to an analytic function on the strip $\{z \in \mathbb{C} : \Re(z) \in (-1, 1)\}$. The pair of functions $\{f, g\}$ is called a Krivine rounding scheme if $H_{f,g}$ is invertible on a neighborhood of the origin, and if we consider the Taylor expansion $H_{f,g}^{-1}(z) = \sum_{j=0}^{\infty} a_{2j+1} z^{2j+1}$ then there exists $c = c(f, g) \in (0, \infty)$ satisfying $\sum_{j=0}^{\infty} |a_{2j+1}| c^{2j+1} = 1$.

For (f, g) as above and unit vectors $\{x_i\}_{i=1}^m, \{y_j\}_{j=1}^n \subseteq S^{m+n-1}$, one can find new unit vectors $\{u_i\}_{i=1}^m, \{v_j\}_{j=1}^n \subseteq S^{m+n-1}$ satisfying the identities

$$\forall (i, j) \in \{1, \dots, m\} \times \{1, \dots, n\}, \quad \langle u_i, v_j \rangle = H_{f,g}^{-1}(c(f, g) \langle x_i, y_j \rangle). \quad (33)$$

We refer to [21] for the proof that $\{u_i\}_{i=1}^m, \{v_j\}_{j=1}^n$ exist. This existence proof is not via an efficient algorithm, but as explained in [8], once we know that they exist the new vectors

can be computed efficiently provided $H_{f,g}^{-1}$ can be computed efficiently; this simply amounts to computing a Cholesky decomposition or, alternatively, solving a semidefinite program corresponding to (33). This completes the first (preprocessing) step of a generalized Krivine rounding procedure. The next step is to apply a random projection to the new vectors thus obtained, as in Grothendieck’s original proof [45] or the Goemans-Williamson algorithm [42].

Let $G : \mathbb{R}^{m+n} \rightarrow \mathbb{R}^k$ be a random $k \times (m+n)$ matrix whose entries are i.i.d. standard Gaussian random variables. Define random signs $\{\varepsilon_i\}_{i=1}^m, \{\delta_j\}_{j=1}^n \subseteq \{-1, 1\}$ by

$$\forall (i, j) \in \{1, \dots, m\} \times \{1, \dots, n\}, \quad \varepsilon_i \stackrel{\text{def}}{=} f\left(\frac{1}{\sqrt{2}}Gu_i\right) \quad \text{and} \quad \delta_j \stackrel{\text{def}}{=} g\left(\frac{1}{\sqrt{2}}Gv_j\right). \quad (34)$$

Now,

$$\mathbb{E} \left[\sum_{i=1}^m \sum_{j=1}^n a_{ij} \varepsilon_i \delta_j \right] \stackrel{(*)}{=} \mathbb{E} \left[\sum_{i=1}^m \sum_{j=1}^n a_{ij} H_{f,g}(\langle u_i, v_j \rangle) \right] \stackrel{(33)}{=} c(f, g) \sum_{i=1}^m \sum_{j=1}^n a_{ij} \langle x_i, y_j \rangle, \quad (35)$$

where $(*)$ follows by rotation invariance from (34) and (32). The identity (35) yields the desired polynomial time randomized rounding algorithm, provided one can bound $c(f, g)$ from below. It also gives a systematic way to bound the Grothendieck constant from above: for every Krivine rounding scheme $f, g : \mathbb{R}^k \rightarrow \{-1, 1\}$ we have $K_G \leq 1/c(f, g)$. Krivine used this reasoning to obtain the bound $K_G \leq \pi / (2 \log(1 + \sqrt{2}))$ by considering the case $k = 1$ and $f_0(x) = g_0(x) = \text{sign}(x)$. One checks that $\{f_0, g_0\}$ is a Krivine rounding scheme with $H_{f_0, g_0}(t) = \frac{2}{\pi} \arcsin(t)$ (Grothendieck’s identity) and $c(f_0, g_0) = \frac{2}{\pi} \log(1 + \sqrt{2})$.

Since the goal of the above discussion is to round vectors $\{x_i\}_{i=1}^m, \{y_j\}_{j=1}^n \subseteq S^{m+n-1}$ to signs $\{\varepsilon_i\}_{i=1}^m, \{\delta_j\}_{j=1}^n \subseteq \{-1, 1\}$, it seems natural to expect that the best possible Krivine rounding scheme occurs when $k = 1$ and $f(x) = g(x) = \text{sign}(x)$. If true, this would imply that $K_G = \pi / (2 \log(1 + \sqrt{2}))$; a long-standing conjecture of Krivine [77]. Over the years additional evidence supporting Krivine’s conjecture was discovered, and a natural analytic conjecture was made in [76] as a step towards proving it. We will not discuss these topics here since in [21] it was shown that actually $K_G \leq \pi / (2 \log(1 + \sqrt{2})) - \varepsilon_0$ for some effective constant $\varepsilon_0 > 0$.

It is known [21, Lem. 2.4] that among all *one dimensional* Krivine rounding schemes $f, g : \mathbb{R} \rightarrow \{-1, 1\}$ we indeed have $c(f, g) \leq \frac{2}{\pi} \log(1 + \sqrt{2})$, i.e., it does not pay off to take partitions of \mathbb{R} which are more complicated than the half-line partitions. Somewhat unexpectedly, it was shown in [21] that a certain two dimensional Krivine rounding scheme $f, g : \mathbb{R}^2 \rightarrow \{-1, 1\}$ satisfies $c(f, g) > \frac{2}{\pi} \log(1 + \sqrt{2})$. The proof of [21] uses a Krivine rounding scheme $f, g : \mathbb{R}^2 \rightarrow \{-1, 1\}$ when $f = g$ corresponds to the partition of \mathbb{R}^2 as the sub-graph and super-graph of the polynomial $y = c(x^5 - 10x^3 + 15x)$, where $c > 0$ is an appropriately chosen constant. This partition is depicted in Figure 1.

As explained in [21, Sec. 3], there is a natural guess for the “best” two dimensional Krivine rounding scheme based on a certain numerical computation which we will not discuss here. For this (conjectural) scheme we have $f \neq g$, and the planar partition corresponding to f is depicted in Figure 2. Of course, once Krivine’s conjecture has been disproved and the usefulness of higher dimensional rounding schemes has been established, there is no reason to expect that the situation won’t improve as we consider k -dimensional Krivine rounding schemes for $k \geq 3$. A positive solution to an analytic question presented in [21] might even lead to an exact computation of K_G ; see [21, Sec. 3] for the details.




FIGURE 1. The partition of \mathbb{R}^2 used in [21] to show that K_G is smaller than Krivine’s bound; the shaded regions are separated by the graph $y = c(x^5 - 10x^3 + 15x)$.



FIGURE 2. The “tiger partition” restricted to the square $[-20, 20]^2$. This is the conjectured [21] optimal partition of \mathbb{R}^2 for the purpose of Krivine-type rounding.

3. THE GROTHENDIECK CONSTANT OF A GRAPH

Fix $n \in \mathbb{N}$ and let $G = (\{1, \dots, n\}, E)$ be a graph on the vertices $\{1, \dots, n\}$. We assume throughout that G does not contain any self loops, i.e., $E \subseteq \{S \subseteq \{1, \dots, n\} : |S| = 2\}$. Following [7], define the Grothendieck constant of G , denoted $K(G)$, to be the smallest constant $K \in (0, \infty)$ such that every $n \times n$ matrix (a_{ij}) satisfies

$$\max_{x_1, \dots, x_n \in S^{n-1}} \sum_{\substack{i, j \in \{1, \dots, n\} \\ \{i, j\} \in E}} a_{ij} \langle x_i, x_j \rangle \leq K \max_{\varepsilon_1, \dots, \varepsilon_n \in \{-1, 1\}} \sum_{\substack{i, j \in \{1, \dots, n\} \\ \{i, j\} \in E}} a_{ij} \varepsilon_i \varepsilon_j. \quad (36)$$

Inequality (36) is an extension of the Grothendieck inequality since (1) is the special case of (36) when G is a bipartite graph. Thus

$$K_G = \sup_{n \in \mathbb{N}} \{K(G) : G \text{ is an } n\text{-vertex bipartite graph}\}. \quad (37)$$

The opposite extreme of bipartite graphs is $G = K_n$, the n -vertex complete graph. In this case (36) boils down to the following inequality

$$\max_{x_1, \dots, x_n \in S^{n-1}} \sum_{\substack{i, j \in \{1, \dots, n\} \\ i \neq j}} a_{ij} \langle x_i, x_j \rangle \leq K(K_n) \max_{\varepsilon_1, \dots, \varepsilon_n \in \{-1, 1\}} \sum_{\substack{i, j \in \{1, \dots, n\} \\ i \neq j}} a_{ij} \varepsilon_i \varepsilon_j. \quad (38)$$

It turns out that $K(K_n) \asymp \log n$. The estimate $K(K_n) \lesssim \log n$ was proved in [94, 91, 60, 27]. In fact, as shown in [7, Thm. 3.7], the following stronger inequality holds true for every $n \times n$

matrix (a_{ij}) ; it implies that $K(K_n) \lesssim \log n$ by the Cauchy-Schwartz inequality.

$$\begin{aligned} & \max_{x_1, \dots, x_n \in S^{n-1}} \sum_{\substack{i, j \in \{1, \dots, n\} \\ i \neq j}} a_{ij} \langle x_i, x_j \rangle \\ & \lesssim \log \left(\frac{\sum_{i \in \{1, \dots, n\}} \sum_{j \in \{1, \dots, n\} \setminus \{i\}} |a_{ij}|}{\sqrt{\sum_{i \in \{1, \dots, n\}} \sum_{j \in \{1, \dots, n\} \setminus \{i\}} a_{ij}^2}} \right) \max_{\varepsilon_1, \dots, \varepsilon_n \in \{-1, 1\}} \sum_{\substack{i, j \in \{1, \dots, n\} \\ i \neq j}} a_{ij} \varepsilon_i \varepsilon_j. \end{aligned}$$

The matching lower bound $K(K_n) \gtrsim \log n$ is due to [7], improving over a result of [60].

How can we interpolate between the two extremes (37) and (38)? The Grothendieck constant $K(G)$ depends on the combinatorial structure of the graph G , but at present our understanding of this dependence is incomplete. The following general bounds are known.

$$\log \omega \lesssim K(G) \lesssim \log \vartheta, \quad (39)$$

and

$$K(G) \leq \frac{\pi}{2 \log \left(\frac{1 + \sqrt{(\vartheta-1)^2 + 1}}{\vartheta-1} \right)}, \quad (40)$$

where (39) is due to [7] and (40) is due to [23]. Here ω is the clique number of G , i.e., the largest $k \in \{2, \dots, n\}$ such that there exists $S \subseteq \{1, \dots, n\}$ of cardinality k satisfying $\{i, j\} \in E$ for all distinct $i, j \in S$, and

$$\vartheta = \min \left\{ \max_{i \in \{1, \dots, n\}} \frac{1}{\langle x_i, y \rangle^2} : x_1, \dots, x_n, y \in S^n \wedge \forall \{i, j\} \in E, \langle x_i, x_j \rangle = 0 \right\}. \quad (41)$$

The parameter ϑ is known as the Lovász theta function of the complement of G ; an important graph parameter that was introduced in [87]. We refer to [59] and [7, Thm. 3.5] for alternative characterizations of ϑ . It suffices to say here that it was shown in [87] that $\vartheta \leq \chi$, where χ is the chromatic number of G , i.e., the smallest integer k such that there exists a partition $\{A_1, \dots, A_k\}$ of $\{1, \dots, n\}$ such that $\{i, j\} \notin E$ for all $(i, j) \in \bigcup_{\ell=1}^k A_\ell \times A_\ell$. Note that the upper bound in (39) is superior to (40) when ϑ is large, but when $\vartheta = 2$ the bound (40) implies Krivine's classical bound [77] $K_G \leq \pi / (2 \log(1 + \sqrt{2}))$.

The upper and lower bounds in (39) are known to match up to absolute constants for a variety of graph classes. Several such sharp Grothendieck-type inequalities are presented in Sections 5.2 and 5.3 of [7]. For example, as explained in [7], it follows from (39), combined with combinatorial results of [87, 9], that for every $n \times n \times n$ 3-tensor (a_{ijk}) we have

$$\max_{\{x_{ij}\}_{i,j=1}^n \subseteq S^{n^2-1}} \sum_{\substack{i, j, k \in \{1, \dots, n\} \\ i \neq j \neq k}} a_{ijk} \langle x_{ij}, x_{jk} \rangle \lesssim \max_{\{\varepsilon_{ij}\}_{i,j=1}^n \subseteq \{-1, 1\}} \sum_{\substack{i, j, k \in \{1, \dots, n\} \\ i \neq j \neq k}} a_{ijk} \varepsilon_{ij} \varepsilon_{jk}.$$

While (39) is often a satisfactory asymptotic evaluation of $K(G)$, this isn't always the case. In particular, it is unknown whether $K(G)$ can be bounded from below by a function of ϑ that tends to ∞ as $\vartheta \rightarrow \infty$. An instance in which (39) is not sharp is the case of Erdős-Rényi [36] random graphs $G(n, 1/2)$. For such graphs we have $\omega \asymp \log n$ almost surely as $n \rightarrow \infty$; see [90] and [10, Sec. 4.5]. At the same time, for $G(n, 1/2)$ we have [58] $\vartheta \asymp \sqrt{n}$ almost surely as $n \rightarrow \infty$. Thus (39) becomes in this case the rather weak estimate $\log \log n \lesssim K(G(n, 1/2)) \lesssim \log n$. It turns out [3] that $K(G(n, 1/2)) \asymp \log n$ almost surely as

$n \rightarrow \infty$; we refer to [3] for additional computations of this type of the Grothendieck constant of random and psuedo-random graphs. An explicit evaluation of the Grothendieck constant of certain graph families can be found in [79]; for example, if G is a graph of girth g that is not a forest and does not admit K_5 as a minor then $K(G) = \frac{g \cos(\pi/g)}{g-2}$.

3.1. Algorithmic consequences. Other than being a natural variant of the Grothendieck inequality, and hence of intrinsic mathematical interest, (36) has ramifications to discrete optimization problems, which we now describe.

3.1.1. Spin glasses. Perhaps the most natural interpretation of (36) is in the context of solid state physics, specifically the problem of efficient computation of ground states of Ising spin glasses. The graph G represents the interaction pattern of n particles; thus $\{i, j\} \notin E$ if and only if the particles i and j cannot interact with each other. Let a_{ij} be the magnitude of the interaction of i and j (the sign of a_{ij} corresponds to attraction/repulsion). In the Ising model each particle $i \in \{1, \dots, n\}$ has a spin $\varepsilon_i \in \{-1, 1\}$ and the total energy of the system is given by the quantity $-\sum_{\{i,j\} \in E} a_{ij} \varepsilon_i \varepsilon_j$. A spin configuration $(\varepsilon_1, \dots, \varepsilon_n) \in \{-1, 1\}^n$ is called a ground state if it minimizes the total energy. Thus the problem of finding a ground state is precisely that of computing the maximum appearing in the right hand side of (36). For more information on this topic see [88, pp. 352–355].

Physical systems seek to settle at a ground state, and therefore it is natural to ask whether it is computationally efficient (i.e., polynomial time computable) to find such a ground state, at least approximately. Such questions have been studied in the physics literature for several decades; see [18, 16, 13, 22]. In particular, it was shown in [16] that if G is a planar graph then one can find a ground state in polynomial time, but in [13] it was shown that when G is the three dimensional grid then this computational task is NP-hard.

Since the quantity in the left hand side of (36) is a semidefinite program and therefore can be computed in polynomial time with arbitrarily good precision, a good bound on $K(G)$ yields a polynomial time algorithm that computes the energy of a ground state with correspondingly good approximation guarantee. Moreover, as explained in [7], the proof of the upper bound in (39) yields a polynomial time algorithm that finds a spin configuration $(\sigma_1, \dots, \sigma_n) \in \{-1, 1\}^n$ for which

$$\sum_{\substack{i,j \in \{1, \dots, n\} \\ \{i,j\} \in E}} a_{ij} \sigma_i \sigma_j \gtrsim \frac{1}{\log \vartheta} \cdot \max_{\{\varepsilon_i\}_{i=1}^n \subseteq \{-1, 1\}} \sum_{\substack{i,j \in \{1, \dots, n\} \\ \{i,j\} \in E}} a_{ij} \varepsilon_i \varepsilon_j. \quad (42)$$

An analogous polynomial time algorithm corresponds to the bound (40). These algorithms yield the best known efficient methods for computing a ground state of Ising spin glasses on a variety of interaction graphs.

3.1.2. Correlation clustering. A different interpretation of (36) yields the best known polynomial time approximation algorithm for the correlation clustering problem [14, 25]; this connection is due to [27]. Interpret the graph $G = (\{1, \dots, n\}, E)$ as the “similarity/dissimilarity graph” for the items $\{1, \dots, n\}$, in the following sense. For $\{i, j\} \in E$ we are given a sign $a_{ij} \in \{-1, 1\}$ which has the following meaning: if $a_{ij} = 1$ then i and j are deemed to be similar, and if $a_{ij} = -1$ then i and j are deemed to be different. If $\{i, j\} \notin E$ then we do not express any judgement on the similarity or dissimilarity of i and j .

Assume that A_1, \dots, A_k is a partition (or “clustering”) of $\{1, \dots, n\}$. An agreement between this clustering and our similarity/dissimilarity judgements is a pair $i, j \in \{1, \dots, n\}$ such that $a_{ij} = 1$ and $i, j \in A_r$ for some $r \in \{1, \dots, k\}$ or $a_{ij} = -1$ and $i \in A_r, j \in A_s$ for distinct $r, s \in \{1, \dots, k\}$. A disagreement between this clustering and our similarity/dissimilarity judgements is a pair $i, j \in \{1, \dots, n\}$ such that $a_{ij} = 1$ and $i \in A_r, j \in A_s$ for distinct $r, s \in \{1, \dots, k\}$ or $a_{ij} = -1$ and $i, j \in A_r$ for some $r \in \{1, \dots, k\}$. Our goal is to cluster the items while encouraging agreements and penalizing disagreements. Thus, we wish to find a clustering of $\{1, \dots, n\}$ into an unspecified number of clusters which maximizes the total number of agreements minus the total number of disagreements.

It was proved in [27] that the case of clustering into two parts is the “bottleneck” for this problem: if there were a polynomial time algorithm that finds a clustering into two parts for which the total number of agreements minus the total number of disagreements is at least a fraction $\alpha \in (0, 1)$ of the maximum possible (over all bi-partitions) total number of agreements minus the total number of disagreements, then one could find in polynomial time a clustering which is at least a fraction $\alpha/(2 + \alpha)$ of the analogous maximum that is defined without specifying the number of clusters.

One checks that the problem of finding a partition into two clusters that maximizes the total number of agreements minus the total number of disagreements is the same as the problem of computing the maximum in the right hand side of (36). Thus the upper bound in (39) yields a polynomial time algorithm for correlation clustering with approximation guarantee $O(\log \vartheta)$, which is the best known approximation algorithm for this problem. Note that when G is the complete graph then the approximation ratio is $O(\log n)$. As will be explained in Section 7, it is known [69] that for every $\gamma \in (0, 1/6)$, if there were a polynomial time algorithm for correlation clustering that yields an approximation guarantee of $(\log n)^\gamma$ then there would be an algorithm for 3-colorability that runs in time $2^{(\log n)^{O(1)}}$, a conclusion which is widely believed to be impossible.

4. KERNEL CLUSTERING AND THE PROPELLER CONJECTURE

Here we describe a large class of Grothendieck-type inequalities that is motivated by algorithmic applications to a combinatorial optimization problem called Kernel Clustering. This problem originates in machine learning [110], and its only known rigorous approximation algorithms follow from Grothendieck inequalities (these algorithms are sharp assuming the UGC). We will first describe the inequalities and then the algorithmic application.

Consider the special case of the Grothendieck inequality (1) where $A = (a_{ij})$ is an $n \times n$ positive semidefinite matrix. In this case we may assume without loss of generality that in (1) $x_i = y_i$ and $\varepsilon_i = \delta_i$ for every $i \in \{1, \dots, n\}$ since this holds for the maxima on either side of (1) (see also the explanation in [8, Sec. 5.2]). It follows from [45, 107] (see also [95]) that for every $n \times n$ symmetric positive semidefinite matrix $A = (a_{ij})$ we have

$$\max_{x_1, \dots, x_n \in S^{n-1}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle \leq \frac{\pi}{2} \cdot \max_{\varepsilon_1, \dots, \varepsilon_n \in \{-1, 1\}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \varepsilon_i \varepsilon_j, \quad (43)$$

and that $\frac{\pi}{2}$ is the best possible constant in (43).

A natural variant of (43) is to replace the numbers $-1, 1$ by general vectors $v_1, \dots, v_k \in \mathbb{R}^k$, namely one might ask for the smallest constant $K \in (0, \infty)$ such that for every symmetric

positive semidefinite $n \times n$ matrix (a_{ij}) we have:

$$\max_{x_1, \dots, x_n \in S^{n-1}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle \leq K \max_{u_1, \dots, u_n \in \{v_1, \dots, v_k\}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle u_i, u_j \rangle. \quad (44)$$

The best constant K in (44) can be characterized as follows. Let $B = (b_{ij} = \langle v_i, v_j \rangle)$ be the Gram matrix of v_1, \dots, v_k . Let $C(B)$ be the maximum over all partitions $\{A_1, \dots, A_k\}$ of \mathbb{R}^{k-1} into measurable sets of the quantity $\sum_{i=1}^k \sum_{j=1}^k b_{ij} \langle z_i, z_j \rangle$, where for $i \in \{1, \dots, k\}$ the vector $z_i \in \mathbb{R}^{k-1}$ is the Gaussian moment of A_i , i.e.,

$$z_i = \frac{1}{(2\pi)^{(k-1)/2}} \int_{A_i} x e^{-\|x\|_2^2/2} dx.$$

It was proved in [67] that (44) holds with $K = 1/C(B)$ and that this constant is sharp.

Inequality (44) with $K = 1/C(B)$ is proved via the following *rounding procedure*. Fix unit vectors $x_1, \dots, x_n \in S^{n-1}$. Let $G = (g_{ij})$ be a $(k-1) \times n$ random matrix whose entries are i.i.d. standard Gaussian random variables. Let $A_1, \dots, A_k \subseteq \mathbb{R}^{k-1}$ be a measurable partition of \mathbb{R}^{k-1} at which $C(B)$ is attained (for a proof that the maximum defining $C(B)$ is indeed attained, see [67]). Define a random choice of $u_i \in \{v_1, \dots, v_k\}$ by setting $u_i = v_\ell$ for the unique $\ell \in \{1, \dots, k\}$ such that $Gx_i \in A_\ell$. The fact that (44) holds with $K = 1/C(B)$ is a consequence of the following fact, whose proof we skip (the full details are in [67]).

$$\mathbb{E} \left[\sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle u_i, u_j \rangle \right] \geq C(B) \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle. \quad (45)$$

Determining the partition of \mathbb{R}^{k-1} that achieves the value $C(B)$ is a nontrivial problem in general, even in the special case when $B = I_k$ is the $k \times k$ identity matrix. Note that in this case one desires a partition $\{A_1, \dots, A_k\}$ of \mathbb{R}^{k-1} into measurable sets so as to maximize the following quantity.

$$\sum_{i=1}^k \left\| \frac{1}{(2\pi)^{(k-1)/2}} \int_{A_i} x e^{-\|x\|_2^2/2} dx \right\|_2^2.$$

As shown in [66, 67], the optimal partition is given by simplicial cones centered at the origin. When $B = I_2$ we have $C(I_2) = \frac{1}{\pi}$, and the optimal partition of \mathbb{R} into two cones is the positive and the negative axes. When $B = I_3$ it was shown in [66] that $C(I_3) = \frac{9}{8\pi}$, and the optimal partition of \mathbb{R}^2 into three cones is the *propeller* partition, i.e., into three cones with angular measure 120° each.

Though it might be surprising at first sight, the authors posed in [66] the *propeller conjecture*: for any $k \geq 4$, the optimal partition of \mathbb{R}^{k-1} into k parts is $\mathcal{P} \times \mathbb{R}^{k-3}$ where \mathcal{P} is the propeller partition of \mathbb{R}^2 . In other words, even if one is allowed to use k parts, the propeller conjecture asserts that the best partition consists of only three nonempty parts. Recently, this conjecture was solved positively [53] for $k = 4$, i.e., for partitions of \mathbb{R}^3 into four measurable parts. The proof of [53] reduces the problem to a concrete finite set of numerical inequalities which are then verified with full rigor in a computer-assisted fashion. Note that this is the first nontrivial (surprising?) case of the propeller conjecture, i.e., this is the first case in which we indeed drop one of the four allowed parts in the optimal partition.

We now describe an application of (44) to the Kernel Clustering problem; a general framework for clustering massive statistical data so as to uncover a certain hypothesized structure [110]. The problem is defined as follows. Let $A = (a_{ij})$ be an $n \times n$ symmetric positive semidefinite matrix which is usually normalized to be centered, i.e., $\sum_{i=1}^n \sum_{j=1}^n a_{ij} = 0$. The matrix A is often thought of as the correlation matrix of random variables (X_1, \dots, X_n) that measure attributes of certain empirical data, i.e., $a_{ij} = \mathbb{E}[X_i X_j]$. We are also given another symmetric positive semidefinite $k \times k$ matrix $B = (b_{ij})$ which functions as a hypothesis, or test matrix. Think of n as huge and k as a small constant. The goal is to cluster A so as to obtain a smaller matrix which most resembles B . Formally, we wish to find a partition $\{S_1, \dots, S_k\}$ of $\{1, \dots, n\}$ so that if we write $c_{ij} = \sum_{(p,q) \in S_i \times S_j} a_{pq}$ then the resulting clustered version of A has the maximum correlation $\sum_{i=1}^k \sum_{j=1}^k c_{ij} b_{ij}$ with the hypothesis matrix B . In words, we form a $k \times k$ matrix $C = (c_{ij})$ by summing the entries of A over the blocks induced by the given partition, and we wish to produce in this way a matrix that is most correlated with B . Equivalently, the goal is to evaluate the number:

$$\mathbf{Clust}(A|B) = \max_{\sigma: \{1, \dots, n\} \rightarrow \{1, \dots, k\}} \sum_{i=1}^k \sum_{j=1}^k a_{ij} b_{\sigma(i)\sigma(j)}. \quad (46)$$

The strength of this generic clustering framework is based in part on the flexibility of adapting the matrix B to the problem at hand. Various particular choices of B lead to well studied optimization problems, while other specialized choices of B are based on statistical hypotheses which have been applied with some empirical success. We refer to [110, 66] for additional background and a discussion of specific examples.

In [66] it was shown that there exists a randomized polynomial time algorithm that takes as input two positive semidefinite matrices A, B and outputs a number α that satisfies $\mathbf{Clust}(A|B) \leq \mathbb{E}[\alpha] \leq (1 + \frac{3\pi}{2}) \mathbf{Clust}(A|B)$. There is no reason to believe that the approximation factor of $1 + \frac{3\pi}{2}$ is sharp, but nevertheless prior to this result, which is based on (44), no constant factor polynomial time approximation algorithm for this problem was known.

Sharper results can be obtained if we assume that the input matrices are normalized appropriately. Specifically, assume that $k \geq 3$ and restrict only to inputs A that are centered, i.e., $\sum_{i=1}^n \sum_{j=1}^n a_{ij} = 0$, and inputs B that are either the identity matrix I_k , or satisfy $\sum_{i=1}^k \sum_{j=1}^k b_{ij} = 0$ (B is centered as well) and $b_{ii} = 1$ for all $i \in \{1, \dots, k\}$ (B is “spherical”). Under these assumptions the output of the algorithm of [66] satisfies $\mathbf{Clust}(A|B) \leq \mathbb{E}[\alpha] \leq \frac{8\pi}{9} (1 - \frac{1}{k}) \mathbf{Clust}(A|B)$. Moreover, it was shown in [66] that assuming the propeller conjecture and the UGC, no polynomial time algorithm can achieve an approximation guarantee that is strictly smaller than $\frac{8\pi}{9} (1 - \frac{1}{k})$ (for input matrices normalized as above). Since the propeller conjecture is known to hold true for $k = 3$ [66] and $k = 4$ [53], we know that the UGC hardness threshold for the above problem is exactly $\frac{16\pi}{27}$ when $k = 3$ and $\frac{2\pi}{3}$ when $k = 4$.

A finer, and perhaps more natural, analysis of the kernel clustering problem can be obtained if we fix the matrix B and let the input be only the matrix A , with the goal being, as before, to approximate the quantity $\mathbf{Clust}(A|B)$ in polynomial time. Since B is symmetric and positive semidefinite we can find vectors $v_1, \dots, v_k \in \mathbb{R}^k$ such that B is their Gram matrix, i.e., $b_{ij} = \langle v_i, v_j \rangle$ for all $i, j \in \{1, \dots, k\}$. Let $R(B)$ be the smallest possible radius of a Euclidean ball in \mathbb{R}^k which contains $\{v_1, \dots, v_k\}$ and let $w(B)$ be the center of this ball.

We note that both $R(B)$ and $w(B)$ can be efficiently computed by solving an appropriate semidefinite program. Let $C(B)$ be the parameter defined above.

It is shown in [67] that for every fixed symmetric positive semidefinite $k \times k$ matrix B there exists a randomized polynomial time algorithm which given an $n \times n$ symmetric positive semidefinite centered matrix A , outputs a number $\text{Alg}(A)$ such that

$$\mathbf{Clust}(A|B) \leq \mathbb{E}[\text{Alg}(A)] \leq \frac{R(B)^2}{C(B)} \mathbf{Clust}(A|B).$$

As we will explain in Section 7, assuming the UGC no polynomial time algorithm can achieve an approximation guaranty strictly smaller than $R(B)^2/C(B)$.

The algorithm of [67] uses semidefinite programming to compute the value

$$\begin{aligned} \text{SDP}(A|B) &= \max \left\{ \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle : x_1, \dots, x_n \in \mathbb{R}^n \wedge \|x_i\|_2 \leq 1 \forall i \in \{1, \dots, n\} \right\} \\ &= \max \left\{ \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle : x_1, \dots, x_n \in S^{n-1} \right\}, \end{aligned} \quad (47)$$

where the last equality in (47) holds since the function $(x_1, \dots, x_n) \mapsto \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle$ is convex (by virtue of the fact that A is positive semidefinite). We claim that

$$\frac{\mathbf{Clust}(A|B)}{R(B)^2} \leq \text{SDP}(A|B) \leq \frac{\mathbf{Clust}(A|B)}{C(B)}, \quad (48)$$

which implies that if we output the number $R(B)^2 \text{SDP}(A|B)$ we will obtain a polynomial time algorithm which approximates $\mathbf{Clust}(A|B)$ up to a factor of $\frac{R(B)^2}{C(B)}$. To verify (48) let $x_1^*, \dots, x_n^* \in S^{n-1}$ and $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$ be such that

$$\text{SDP}(A|B) = \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i^*, x_j^* \rangle \quad \text{and} \quad \mathbf{Clust}(A|B) = \sum_{i=1}^n \sum_{j=1}^n a_{ij} b_{\sigma^*(i)\sigma^*(j)}.$$

Write $(a_{ij})_{i,j=1}^n = (\langle u_i, u_j \rangle)_{i,j=1}^n$ for some $u_1, \dots, u_n \in \mathbb{R}^n$. The assumption that A is centered means that $\sum_{i=1}^n u_i = 0$. The rightmost inequality in (48) is just the Grothendieck inequality (44). The leftmost inequality in (48) follows from the fact that $\frac{v_{\sigma^*(i)} - w(B)}{R(B)}$ has norm at most 1 for all $i \in \{1, \dots, n\}$. Indeed, these norm bounds imply that

$$\begin{aligned} \text{SDP}(A|B) &\geq \sum_{i=1}^n \sum_{j=1}^n a_{ij} \left\langle \frac{v_{\sigma^*(i)} - w(B)}{R(B)}, \frac{v_{\sigma^*(j)} - w(B)}{R(B)} \right\rangle \\ &= \frac{1}{R(B)^2} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle v_{\sigma^*(i)}, v_{\sigma^*(j)} \rangle \\ &\quad - \frac{2}{R(B)^2} \sum_{i=1}^n \langle w(B), v_{\sigma^*(i)} \rangle \left\langle u_i, \sum_{j=1}^n u_j \right\rangle + \frac{\|w(B)\|_2^2}{R(B)^2} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \\ &= \frac{\mathbf{Clust}(A|B)}{R(B)^2}. \end{aligned}$$

This completes the proof that the above algorithm approximates efficiently the number $\text{Clust}(A|B)$, but does not address the issue of how to efficiently compute an assignment $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, k\}$ for which the induced clustering of A has the required value. The issue here is to find efficiently a conical simplicial partition A_1, \dots, A_k of \mathbb{R}^{k-1} at which $C(B)$ is attained. Such a partition exists and may be assumed to be hardwired into the description of the algorithm. Alternately, the partition that achieves $C(B)$ up to a desired degree of accuracy can be found by brute-force for fixed k (or $k = k(n)$ growing sufficiently slowly as a function of n); see [67]. For large values of k the problem of computing $C(B)$ efficiently remains open.

5. THE L_p GROTHENDIECK PROBLEM

Fix $p \in [1, \infty]$ and consider the following algorithmic problem. The input is an $n \times n$ matrix $A = (a_{ij})$ whose diagonal entries vanish, and the goal is to compute (or estimate) in polynomial time the quantity

$$M_p(A) = \max_{\substack{t_1, \dots, t_n \in \mathbb{R} \\ \sum_{k=1}^n |t_k|^p \leq 1}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} t_i t_j = \max_{\substack{t_1, \dots, t_n \in \mathbb{R} \\ \sum_{k=1}^n |t_k|^p = 1}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} t_i t_j. \quad (49)$$

The second equality in (49) follows from a straightforward convexity argument since the diagonal entries of A vanish. Some of the results described below hold true without the vanishing diagonal assumption, but we will tacitly make this assumption here since the second equality in (49) makes the problem become purely combinatorial when $p = \infty$. Specifically, if $G = (\{1, \dots, n\}, E)$ is the complete graph then $M_\infty(A) = \max_{\varepsilon_1, \dots, \varepsilon_n \in \{-1, 1\}} \sum_{\{i, j\} \in E} a_{ij} \varepsilon_i \varepsilon_j$. The results described in Section 3 therefore imply that there is a polynomial time algorithm that approximates $M_\infty(A)$ up to a $O(\log n)$ factor, and that it is computationally hard to achieve an approximation guarantee smaller than $(\log n)^\gamma$ for all $\gamma \in (0, 1/6)$.

There are values of p for which the above problem can be solved in polynomial time. When $p = 2$ the quantity $M_2(A)$ is the largest eigenvalue of A , and hence can be computed in polynomial time [43, 82]. When $p = 1$ it was shown in [2] that it is possible to approximate $M_1(A)$ up to a factor of $1 + \varepsilon$ in time $n^{O(1/\varepsilon)}$. It is also shown in [2] that the problem of $(1 + \varepsilon)$ -approximately computing $M_1(A)$ is $W[1]$ complete; we refer to [35] for the definition of this type of hardness result and just say here that it indicates that a running time of $c(\varepsilon)n^{O(1)}$ is impossible.

The algorithm of [2] proceeds by showing that for every $m \in \mathbb{N}$ there exist $y_1, \dots, y_n \in \frac{1}{m}\mathbb{Z}$ with $\sum_{i=1}^n |y_i| \leq 1$ and $\sum_{i=1}^n \sum_{j=1}^n a_{ij} y_i y_j \geq (1 - \frac{1}{m}) M_1(A)$. The number of such vectors y is $1 + \sum_{k=1}^m \sum_{\ell=1}^k 2^\ell \binom{n}{\ell} \binom{k-1}{\ell-1} \leq 4n^m$. An exhaustive search over all such vectors will then approximate $M_1(A)$ to within a factor of $m/(m-1)$ in time $O(n^m)$. To prove the existence of y fix $t_1, \dots, t_n \in \mathbb{R}$ with $\sum_{k=1}^n |t_k| = 1$ and $\sum_{i=1}^n \sum_{j=1}^n a_{ij} t_i t_j = M_1(A)$. Let $X \in \mathbb{R}^n$ be a random vector given by $\Pr[X = \text{sign}(t_j)e_j] = |t_j|$ for every $j \in \{1, \dots, n\}$. Here e_1, \dots, e_n is the standard basis of \mathbb{R}^n . Let $\{X_s = (X_{s1}, \dots, X_{sn})\}_{s=1}^m$ be independent copies of X and set $Y = (Y_1, \dots, Y_n) = \frac{1}{m} \sum_{s=1}^m X_s$. Note that if $s, t \in \{1, \dots, m\}$ are distinct then for all $i, j \in \{1, \dots, n\}$ we have $\mathbb{E}[X_{si} X_{tj}] = \text{sign}(t_i) \text{sign}(t_j) |t_i| \cdot |t_j| = t_i t_j$. Also, for every $s \in \{1, \dots, m\}$ and every distinct $i, j \in \{1, \dots, n\}$ we have $X_{si} X_{sj} = 0$. Since the diagonal

entries of A vanish it follows that

$$\mathbb{E} \left[\sum_{i=1}^n \sum_{j=1}^n a_{ij} Y_i Y_j \right] = \frac{1}{m^2} \sum_{\substack{s,t \in \{1, \dots, m\} \\ s \neq t}} \sum_{\substack{i,j \in \{1, \dots, n\} \\ i \neq j}} a_{ij} \mathbb{E} [X_{si} X_{tj}] = \left(1 - \frac{1}{m}\right) M_1(A). \quad (50)$$

Noting that the vector Y has ℓ_1 norm at most 1 and all of its entries are integer multiples of $1/m$, it follows from (50) that with positive probability Y will have the desired properties.

How can we interpolate between the above results for $p \in \{1, 2, \infty\}$? It turns out that there is a satisfactory answer for $p \in (2, \infty)$ but the range $p \in (1, 2)$ remains a mystery. To explain this write $\gamma_p = (\mathbb{E} [|G|^p])^{1/p}$, where G is a standard Gaussian random variable. One computes that

$$\gamma_p = \sqrt{2} \left(\frac{\Gamma(\frac{p+1}{2})}{\sqrt{\pi}} \right)^{1/p}. \quad (51)$$

Also, Stirling's formula implies that $\gamma_p^2 = \frac{p}{e} + O(1)$ as $p \rightarrow \infty$. It follows from [92, 48] that for every fixed $p \in [2, \infty)$ there exists a polynomial time algorithm that approximates $M_p(A)$ to within a factor of γ_p^2 , and that for every $\varepsilon \in (0, 1)$ the existence of a polynomial time algorithm that approximates $M_p(A)$ to within a factor $\gamma_p^2 - \varepsilon$ would imply that $P = NP$. These results improve over the earlier work [70] which designed a polynomial time algorithm for $M_p(A)$ whose approximation guarantee is $(1 + o(1))\gamma_p^2$ as $p \rightarrow \infty$, and which proved a $\gamma_p^2 - \varepsilon$ hardness results assuming the UGC rather than $P \neq NP$.

The following Grothendieck-type inequality was proved in [92] and independently in [48]. For every $n \times n$ matrix $A = (a_{ij})$ and every $p \in [2, \infty)$ we have

$$\max_{\substack{x_1, \dots, x_n \in \mathbb{R}^n \\ \sum_{k=1}^n \|x_k\|_2^p \leq 1}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle \leq \gamma_p^2 \max_{\substack{t_1, \dots, t_n \in \mathbb{R} \\ \sum_{k=1}^n |t_k|^p \leq 1}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} t_i t_j. \quad (52)$$

The constant γ_p^2 in (52) is sharp. The validity of (52) implies that $M_p(A)$ can be computed in polynomial time to within a factor γ_p^2 . This follows since the left hand side of (52) is the maximum of $\sum_{i=1}^n \sum_{j=1}^n a_{ij} X_{ij}$, which is a linear functional in the variables (X_{ij}) , given the constraint that (X_{ij}) is a symmetric positive semidefinite matrix and $\sum_{i=1}^n X_{ii}^{p/2} \leq 1$. The latter constraint is convex since $p \geq 2$, and therefore this problem falls into the framework of convex programming that was described in Section 1.2. Thus the left hand side of (52) can be computed in polynomial time with arbitrarily good precision.

Choosing the specific value $p = 3$ in order to illustrate the current satisfactory state of affairs concretely, the NP -hardness threshold of computing $\max_{\sum_{i=1}^n |x_i|^3 \leq 1} \sum_{i=1}^n \sum_{j=1}^n a_{ij} x_i x_j$ equals $2/\sqrt[3]{\pi}$. Such a sharp NP -hardness result (with transcendental hardness ratio) is quite remarkable, since it shows that the geometric algorithm presented above probably yields the best possible approximation guarantee even when one allows any polynomial time algorithm whatsoever. Results of this type have been known to hold under the UGC, but this NP -hardness result of [48] seems to be the first time that such an algorithm for a simple to state problem was shown to be optimal assuming $P \neq NP$.

When $p \in [1, 2]$ one can easily show [92] that

$$\max_{\substack{x_1, \dots, x_n \in \mathbb{R}^n \\ \sum_{k=1}^n \|x_k\|_2^p \leq 1}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle = \max_{\substack{t_1, \dots, t_n \in \mathbb{R} \\ \sum_{k=1}^n |t_k|^p \leq 1}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} t_i t_j. \quad (53)$$

While the identity (53) seems to indicate the problem of computing $M_p(A)$ in polynomial time might be easy for $p \in (1, 2)$, the above argument fails since the constraint $\sum_{i=1}^n X_{ii}^{p/2} \leq 1$ is no longer convex. This is reflected by the fact that despite (53) the problem of $(1 + \varepsilon)$ -approximately computing $M_1(A)$ is $W[1]$ complete [2]. It remains open whether for $p \in (1, 2)$ one can approximate $M_p(A)$ in polynomial time up to a factor $O(1)$, and no hardness of approximation result is known for this problem as well.

Remark 5.1. If $p \in [2, \infty]$ then for positive semidefinite matrices (a_{ij}) the constant γ_p^2 in the right hand side of (52) can be improved [92] to $\gamma_{p^*}^{-2}$, where here and in what follows $p^* = p/(p-1)$. For $p = \infty$ this estimate coincides with the classical bound [45, 107] that we have already encountered in (43), and it is sharp in the entire range $p \in [2, \infty]$. Moreover, this bound shows that there exists a polynomial time algorithm that takes as input a positive semidefinite matrix A and outputs a number that is guaranteed to be within a factor $\gamma_{p^*}^{-2}$ of $M_p(A)$. Conversely, the existence of a polynomial time algorithm for this problem whose approximation guarantee is strictly smaller than $\gamma_{p^*}^{-2}$ would contradict the UGC [92].

Remark 5.2. The bilinear variant of (52) is an immediate consequence of the Grothendieck inequality (1). Specifically, assume that $p, q \in [1, \infty]$ and $x_1, \dots, x_m, y_1, \dots, y_n \in \mathbb{R}^{m+n}$ satisfy $\sum_{i=1}^m \|x_i\|_2^p \leq 1$ and $\sum_{j=1}^n \|y_j\|_2^q \leq 1$. Write $\alpha_i = \|x_i\|_2$ and $\beta_j = \|y_j\|_2$. For an $m \times n$ matrix (a_{ij}) the Grothendieck inequality provides $\varepsilon_1, \dots, \varepsilon_m, \delta_1, \dots, \delta_n \in \{-1, 1\}$ such that $\sum_{i=1}^m \sum_{j=1}^n a_{ij} \langle x_i, y_j \rangle \leq K_G \sum_{i=1}^m \sum_{j=1}^n a_{ij} \alpha_i \beta_j \varepsilon_i \delta_j$. This establishes the following inequality.

$$\max_{\substack{\{x_i\}_{i=1}^m, \{y_j\}_{j=1}^n \subseteq \mathbb{R}^{m+n} \\ \sum_{i=1}^m \|x_i\|_2^p \leq 1 \\ \sum_{j=1}^n \|y_j\|_2^q \leq 1}} \sum_{i=1}^m \sum_{j=1}^n a_{ij} \langle x_i, y_j \rangle \leq K_G \cdot \max_{\substack{\{s_i\}_{i=1}^m, \{t_j\}_{j=1}^n \subseteq \mathbb{R} \\ \sum_{i=1}^m |s_i|^p \leq 1 \\ \sum_{j=1}^n |t_j|^q \leq 1}} \sum_{i=1}^m \sum_{j=1}^n a_{ij} s_i t_j. \quad (54)$$

Observe that the maximum on the right hand side of (54) is $\|A\|_{p \rightarrow q^*}$; the operator norm of A acting as a linear operator from $(\mathbb{R}^m, \|\cdot\|_p)$ to $(\mathbb{R}^n, \|\cdot\|_{q^*})$. Moreover, if $p, q \geq 2$ then the left hand side of (54) can be computed in polynomial time. Thus, for $p \geq 2 \geq r \geq 1$, the generalized Grothendieck inequality (54) yields a polynomial time algorithm that takes as input an $m \times n$ matrix $A = (a_{ij})$ and outputs a number that is guaranteed to be within a factor K_G of $\|A\|_{p \rightarrow r}$. This algorithmic task has been previously studied in [96] (see also [93, Sec. 4.3.2]), where for $p \geq 2 \geq r \geq 1$ a polynomial time algorithm was designed that approximates $\|A\|_{p \rightarrow r}$ up to a factor $3\pi / (6\sqrt{3} - 2\pi) \in [2.293, 2.294]$. The above argument yields the approximation factor $K_G < 1.783$ as a formal consequence of the Grothendieck inequality. The complexity of the problem of approximating $\|A\|_{p \rightarrow r}$ has been studied in [17], where it is shown that if either $p \geq r > 2$ or $2 > p \geq r$ then it is NP -hard to approximate $\|A\|_{p \rightarrow r}$ up to any constant factor, and unless 3-colorability can be solved in time $2^{(\log n)^{O(1)}}$, for any $\varepsilon \in (0, 1)$ no polynomial time algorithm can approximate $\|A\|_{p \rightarrow r}$ up to $2^{(\log n)^{1-\varepsilon}}$.

Remark 5.3. Let $K \subseteq \mathbb{R}^n$ be a compact and convex set which is invariant under reflections with respect to the coordinate hyperplanes. Denote by C_K the smallest $C \in (0, \infty)$ such

that for every $n \times n$ matrix (a_{ij}) we have

$$\max_{\substack{x_1, \dots, x_n \in \mathbb{R}^n \\ (\|x_1\|_2, \dots, \|x_n\|_2) \in K}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \langle x_i, x_j \rangle \leq C \max_{\substack{t_1, \dots, t_n \in \mathbb{R} \\ (t_1, \dots, t_n) \in K}} \sum_{i=1}^n \sum_{j=1}^n a_{ij} t_i t_j. \quad (55)$$

Such generalized Grothendieck inequalities are investigated in [92], where bounds on C_K are obtained under certain geometric assumptions on K . These assumptions are easy to verify when $K = \{x \in \mathbb{R}^n : \|x\|_p \leq 1\}$, yielding (52). More subtle inequalities of this type for other convex bodies K are discussed in [92], but we will not describe them here. The natural bilinear version of (55) is: if $K \subseteq \mathbb{R}^m$ and $L \subseteq \mathbb{R}^n$ are compact and convex sets that are invariant under reflections with respect to the coordinate hyperplanes then let $C_{K,L}$ denote the smallest constant $C \in (0, \infty)$ such that for every $m \times n$ matrix (a_{ij}) we have

$$\max_{\substack{\{x_i\}_{i=1}^m, \{y_j\}_{j=1}^n \subseteq \mathbb{R}^{n+m} \\ (\|x_1\|_2, \dots, \|x_m\|_2) \in K \\ (\|y_1\|_2, \dots, \|y_n\|_2) \in L}} \sum_{i=1}^m \sum_{j=1}^n a_{ij} \langle x_i, y_j \rangle \leq C \max_{\substack{\{s_i\}_{i=1}^m, \{t_j\}_{j=1}^n \subseteq \mathbb{R} \\ (s_1, \dots, s_m) \in K \\ (t_1, \dots, t_n) \in L}} \sum_{i=1}^m \sum_{j=1}^n a_{ij} s_i t_j. \quad (56)$$

The argument in Remark 5.2 shows that $C_{K,L} \leq K_G$. Under certain geometric assumptions on K, L this bound can be improved [92].

6. HIGHER RANK GROTHENDIECK INEQUALITIES

We have already seen several variants of the classical Grothendieck inequality (1), including the Grothendieck inequality for graphs (36), the variant of the positive semidefinite Grothendieck inequality arising from the Kernel Clustering problem (44), and Grothendieck inequalities for convex bodies other than the cube (52), (54), (55), (56). The literature contains additional variants of the Grothendieck inequality, some of which will be described in this section.

Let $G = (\{1, \dots, n\}, E)$ be a graph and fix $q, r \in \mathbb{N}$. Following [23], let $K(q \rightarrow r, G)$ be the smallest constant $K \in (0, \infty)$ such that for every $n \times n$ matrix $A = (a_{ij})$ we have

$$\max_{x_1, \dots, x_n \in S^{q-1}} \sum_{\substack{i, j \in \{1, \dots, n\} \\ \{i, j\} \in E}} a_{ij} \langle x_i, x_j \rangle \leq K \max_{y_1, \dots, y_n \in S^{r-1}} \sum_{\substack{i, j \in \{1, \dots, n\} \\ \{i, j\} \in E}} a_{ij} \langle y_i, y_j \rangle. \quad (57)$$

Set also $K(r, G) = \sup_{q \in \mathbb{N}} K(q \rightarrow r, G)$. We similarly define $K^+(q \rightarrow r, G)$ to be the smallest constant $K \in (0, \infty)$ satisfying (57) for all positive semidefinite matrices A , and correspondingly $K^+(r, G) = \sup_{q \in \mathbb{N}} K^+(q \rightarrow r, G)$.

To link these definitions to what we have already seen in this article, observe that K_G is the supremum of $K(1, G)$ over all finite bipartite graphs G , and due to the results described in Section 4 we have

$$\sup_{n \in \mathbb{N}} K^+(r, K_n^\circ) = \sup_{n \in \mathbb{N}} \sup_{x_1, \dots, x_n \in S^{r-1}} \frac{1}{C(\langle x_i, x_j \rangle_{i, j=1}^n)}, \quad (58)$$

where K_n° is the complete graph on n -vertices with self loops. Recall that the definition of $C(B)$ for a positive semidefinite matrix B is given in the paragraph following (44).

An important special case of (57) is $r = 2$, since the supremum of $K(2, G)$ over all finite bipartite graphs G is at most the complex Grothendieck constant $K_G^{\mathbb{C}}$ (defined analogously to K_G , but over the complex scalar field), a fundamental quantity whose value has been

investigated in [45, 83, 99, 50, 74]. The best known bounds on $K_G^{\mathbb{C}}$ are $1.338 < K_G^{\mathbb{C}} < 1.4049$; see [101, Sec. 4] for more information on this topic. We also refer to [32, 113] for information of the constants $K(2q \rightarrow 2, G)$ where G is a bipartite graph. The supremum of $K(q \rightarrow r, G)$ over all bipartite graphs G was investigated in [78] for $r = 1$ and in [74] for $r = 2$; see also [75] for a unified treatment of these cases. The higher rank constants $K(q \rightarrow r, G)$ when G is bipartite were introduced in [22]. Definition (57) in full generality is due to [23] where several estimates on $K(q \rightarrow r, G)$ are given. One of the motivations of [23] is the case $r = 3$ (and G a subgraph of the grid \mathbb{Z}^3), based on the connection to the polynomial time approximation of ground states of spin glasses as described in Section 3.1.1; the case $r = 1$ was discussed in Section 3.1.1 in connection with the Ising model, but the case $r = 3$ corresponds to the more physically realistic Heisenberg model of vector-valued spins. The parameter $\sup_{n \in \mathbb{N}} K^+(r, K_n^{\circ})$ (recall (58)) was studied in [22] in the context of quantum information theory, and in [24] it was shown that

$$K^+(1, K_n^{\circ}) \leq \frac{\pi}{n} \left(\frac{\Gamma((n+1)/2)}{\Gamma(n/2)} \right)^2 = \frac{\pi}{2} - \frac{\pi}{4n} + O\left(\frac{1}{n^2}\right), \quad (59)$$

and

$$\sup_{n \in \mathbb{N}} K^+(r, K_n^{\circ}) = \frac{r}{2} \left(\frac{\Gamma(r/2)}{\Gamma((r+1)/2)} \right)^2 = 1 + \frac{1}{2r} + O\left(\frac{1}{r^2}\right).$$

We refer to [24] for a corresponding UGC hardness result. Note that (59) improves over (43) for fixed $n \in \mathbb{N}$.

7. HARDNESS OF APPROXIMATION

We have seen examples of how Grothendieck-type inequalities yield upper bounds on the best possible polynomial time approximation ratio of certain optimization problems. From the algorithmic and computational complexity viewpoint it is interesting to prove computational lower bounds as well, i.e., results that rule out the existence of efficient algorithms achieving a certain approximation guarantee. Such results are known as hardness or inapproximability results, and as explained in Section 1.1, at present the state of the art allows one to prove such results while relying on complexity theoretic assumptions such as $P \neq NP$ or the Unique Games Conjecture. A nice feature of the known hardness results for problems in which a Grothendieck-type inequality has been applied is that often the hardness results (lower bounds) exactly match the approximation ratios (upper bounds). In this section we briefly review the known hardness results for optimization problems associated with Grothendieck-type inequalities.

Let $K_{n,n}$ -QP denote the optimization problem associated with the classical Grothendieck inequality (the acronym QP stands for “quadratic programming”). Thus, in the problem $K_{n,n}$ -QP we are given an $n \times n$ real matrix (a_{ij}) and the goal is to determine the quantity

$$\max \left\{ \sum_{i=1}^m \sum_{j=1}^n a_{ij} \varepsilon_i \delta_j : \{\varepsilon_i\}_{i=1}^m, \{\delta_j\}_{j=1}^n \subseteq \{-1, 1\} \right\}.$$

As explained in [8], the MAX DICUT problem can be framed as a special case of the problem $K_{n,n}$ -QP. Hence, as a consequence of [51], we know that for every $\varepsilon \in (0, 1)$, assuming $P \neq NP$ there is no polynomial time algorithm that approximates the $K_{n,n}$ -QP problem within ratio $\frac{13}{12} - \varepsilon$. In [68] it is shown that the lower bound (3) on the Grothendieck constant

can be translated into a hardness result, albeit relying on the Unique Games Conjecture. Namely, letting η_0 be as in (3), for every $\varepsilon \in (0, 1)$ assuming the UGC there is no polynomial time algorithm that approximates the $K_{n,n}$ -QP problem within a ratio $\frac{\pi}{2}e^{\eta_0^2} - \varepsilon$.

We note that all the hardness results cited here rely on the well-known paradigm of *dictatorship testing*. A lower bound on the integrality gap of a semidefinite program, such as the estimate $K_G \geq \frac{\pi}{2}e^{\eta_0^2}$, can be translated into a probabilistic test to check whether a function $f : \{-1, 1\}^n \mapsto \{-1, 1\}$ is a dictatorship, i.e., of the form $f(x) = x_i$ for some fixed $i \in \{1, \dots, n\}$. If f is indeed a dictatorship, then the test passes with probability c and if f is “far from a dictator” (in a formal sense that we do not describe here), the test passes with probability at most s . The ratio c/s corresponds exactly to the UGC-based hardness lower bound. It is well-known how to prove a UGC-based hardness result once we have the appropriate dictatorship test; see the survey [63].

The above quoted result of [68] relied on explicitly knowing the lower bound construction [105] leading to the estimate $K_G \geq \frac{\pi}{2}e^{\eta_0^2}$. On the other hand, in [104], building on the earlier work [103], it is shown that *any* lower bound on the Grothendieck constant can be translated into a UGC-based hardness result, even without explicitly knowing the construction! Thus, modulo the UGC, the best polynomial time algorithm to approximate the $K_{n,n}$ -QP problem is via the Grothendieck inequality, even though we do not know the precise value of K_G . Formally, for every $\varepsilon \in (0, 1)$, assuming the UGC there is no polynomial time algorithm that approximates the $K_{n,n}$ -QP problem within a factor $K_G - \varepsilon$.

Let $K_{n,n}$ -QP_{PSD} be the special case of the $K_{n,n}$ -QP problem where the input matrix (a_{ij}) is assumed to be positive semidefinite. By considering matrices that are Laplacians of graphs one sees that the MAX CUT problem is a special case of the problem $K_{n,n}$ -QP_{PSD} (see [66]). Hence, due to [51], we know that for every $\varepsilon \in (0, 1)$, assuming $P \neq NP$ there is no polynomial time algorithm that approximates the $K_{n,n}$ -QP_{PSD} problem within ratio $\frac{17}{16} - \varepsilon$. Moreover, it is proved in [66] that for every $\varepsilon \in (0, 1)$, assuming the UGC there is no polynomial time algorithm that approximates the $K_{n,n}$ -QP_{PSD} problem within ratio $\frac{\pi}{2} - \varepsilon$, an optimal hardness result due to the positive semidefinite Grothendieck inequality (43). This follows from the more general results for the Kernel Clustering problem described later.

Let (a_{ij}) be an $n \times n$ real matrix with zeroes on the diagonal. The K_n -QP problem seeks to determine the quantity

$$\max \left\{ \sum_{i=1}^m \sum_{j=1}^n a_{ij} \varepsilon_i \varepsilon_j : \{\varepsilon_i\}_{i=1}^m \subseteq \{-1, 1\} \right\}.$$

In [69] it is proved that for every $\gamma \in (0, 1/6)$, assuming that NP does not have a $2^{(\log n)^{O(1)}}$ time deterministic algorithm, there is no polynomial time algorithm that approximates the K_n -QP problem within ratio $(\log n)^\gamma$. This improves over [12] where a hardness factor of $(\log n)^c$ was proved, under the same complexity assumption, for an unspecified universal constant $c > 0$. Recall that, as explained in Section 3, there is an algorithm for K_n -QP that achieves a ratio of $O(\log n)$, so there remains an asymptotic gap in our understanding of the complexity of the K_n -QP problem. For the maximum acyclic subgraph problem, as discussed in Section 2.1.3, the gap between the upper and lower bounds is even larger. We have already seen that an approximation factor of $O(\log n)$ is achievable, but from the hardness perspective we know due to [97] that there exists $\varepsilon_0 > 0$ such that assuming $P \neq NP$ there is no polynomial time algorithm for the maximum acyclic subgraph problem

that achieves an approximation ratio less than $1 + \varepsilon_0$. In [47] it was shown that assuming the UGC there is no polynomial time algorithm for the maximum acyclic subgraph problem that achieves any constant approximation ratio.

Fix $p \in (0, \infty)$. As discussed in Section 5, the L_p Grothendieck problem is as follows. Given an $n \times n$ real matrix $A = (a_{ij})$ with zeros on the diagonal, the goal is to determine the quantity $M_p(A)$ defined in (49). For $p \in (2, \infty)$ it was shown in [48] that for every $\varepsilon \in (0, 1)$, assuming $P \neq NP$ there is no polynomial time algorithm that approximates the L_p Grothendieck problem within a ratio $\gamma_p^2 - \varepsilon$. Here γ_p is defined as in (51). This result (nontrivially) builds on the previous result of [70] that obtained the same conclusion while assuming the UGC rather than $P \neq NP$.

For the Kernel Clustering problem with a $k \times k$ hypothesis matrix B , an optimal hardness result is obtained in [67] in terms of the parameters $R(B)$ and $C(B)$ described in Section 4. Specifically for a fixed $k \times k$ symmetric positive semidefinite matrix B and for every $\varepsilon \in (0, 1)$, assuming the UGC there is no polynomial time algorithm that, given an $n \times n$ matrix A approximates the quantity $\mathbf{Clust}(A|B)$ within ratio $\frac{R(B)^2}{C(B)} - \varepsilon$. When $B = I_k$ is the $k \times k$ identity matrix, the following hardness result is obtained in [66]. Let $\varepsilon > 0$ be an arbitrarily small constant. Assuming the UGC, there is no polynomial time algorithm that approximates $\mathbf{Clust}(A|I_2)$ within ratio $\frac{\pi}{2} - \varepsilon$. Similarly, assuming the UGC there is no polynomial time algorithm that approximates $\mathbf{Clust}(A|I_3)$ within ratio $\frac{16\pi}{27} - \varepsilon$, and, using also the solution of the propeller conjecture in \mathbb{R}^3 given in [53], there is no polynomial time algorithm that approximates $\mathbf{Clust}(A|I_4)$ within ratio $\frac{2\pi}{3} - \varepsilon$. Furthermore, for $k \geq 5$, assuming the propeller conjecture and the UGC, there is no polynomial time algorithm that approximates $\mathbf{Clust}(A|I_k)$ within ratio $\frac{8\pi}{9} \left(1 - \frac{1}{k}\right) - \varepsilon$.

Acknowledgements. We are grateful to Oded Regev for many helpful suggestions.

REFERENCES

- [1] A. Acín, N. Gisin, and B. Toner. Grothendieck’s constant and local models for noisy entangled quantum states. *Phys. Rev. A* (3), 73(6, part A):062105, 5, 2006.
- [2] N. Alon. Maximizing a quadratic form on the ℓ_1^n unit ball. Unpublished manuscript, 2006.
- [3] N. Alon and E. Berger. The Grothendieck constant of random and pseudo-random graphs. *Discrete Optim.*, 5(2):323–327, 2008.
- [4] N. Alon, A. Coja-Oghlan, H. Hàn, M. Kang, V. Rödl, and M. Schacht. Quasi-randomness and algorithmic regularity for graphs with general degree distributions. *SIAM J. Comput.*, 39(6):2336–2362, 2010.
- [5] N. Alon, R. A. Duke, H. Lefmann, V. Rödl, and R. Yuster. The algorithmic aspects of the regularity lemma. *J. Algorithms*, 16(1):80–109, 1994.
- [6] N. Alon, W. Fernandez de la Vega, R. Kannan, and M. Karpinski. Random sampling and approximation of MAX-CSPs. *J. Comput. System Sci.*, 67(2):212–243, 2003. Special issue on STOC2002 (Montreal, QC).
- [7] N. Alon, K. Makarychev, Y. Makarychev, and A. Naor. Quadratic forms on graphs. *Invent. Math.*, 163(3):499–522, 2006.
- [8] N. Alon and A. Naor. Approximating the cut-norm via Grothendieck’s inequality. *SIAM J. Comput.*, 35(4):787–803 (electronic), 2006.
- [9] N. Alon and A. Orlitsky. Repeated communication and Ramsey graphs. *IEEE Trans. Inform. Theory*, 41(5):1276–1289, 1995.

- [10] N. Alon and J. H. Spencer. *The probabilistic method*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience [John Wiley & Sons], New York, second edition, 2000. With an appendix on the life and work of Paul Erdős.
- [11] S. Arora and B. Barak. *Computational complexity*. Cambridge University Press, Cambridge, 2009. A modern approach.
- [12] S. Arora, E. Berger, G. Kindler, E. Hazan, and S. Safra. On non-approximability for quadratic programs. In *46th Annual Symposium on Foundations of Computer Science*, pages 206–215. IEEE Computer Society, 2005.
- [13] C. P. Bachas. Computer-intractability of the frustration model of a spin glass. *J. Phys. A*, 17(13):L709–L712, 1984.
- [14] N. Bansal, A. Blum, and S. Chawla. Correlation clustering. In *43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 238–247, 2002.
- [15] N. Bansal and R. Williams. Regularity lemmas and combinatorial algorithms. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2009)*, pages 745–754. IEEE Computer Soc., Los Alamitos, CA, 2009.
- [16] F. Barahona. On the computational complexity of Ising spin glass models. *J. Phys. A*, 15(10):3241–3253, 1982.
- [17] A. Bhaskara and A. Vijayaraghavan. Approximating matrix p -norms. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 497–511, 2011.
- [18] I. Bieche, R. Maynard, R. Rammal, and J.-P. Uhry. On the ground states of the frustration model of a spin glass by a matching method of graph theory. *J. Phys. A*, 13(8):2553–2576, 1980.
- [19] R. Blei. *Analysis in integer and fractional dimensions*, volume 71 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2001.
- [20] R. C. Blei. Multidimensional extensions of the Grothendieck inequality and applications. *Ark. Mat.*, 17(1):51–68, 1979.
- [21] M. Braverman, K. Makarychev, Y. Makarychev, and A. Naor. The Grothendieck constant is strictly smaller than Krivine’s bound. An extended abstract will appear in 52nd Annual IEEE Symposium on Foundations of Computer Science. Preprint available at <http://arxiv.org/abs/1103.6161>, 2011.
- [22] J. Briet, H. Buhrman, and B. Toner. A generalized Grothendieck inequality and entanglement in XOR games. Preprint available at <http://arxiv.org/abs/0901.2009>, 2009.
- [23] J. Briet, F. M. de Oliveira Filho, and F. Vallentin. Grothendieck inequalities for semidefinite programs with rank constraint. Preprint available at <http://arxiv.org/abs/1011.1754>, 2010.
- [24] J. Briët, F. M. de Oliveira Filho, and F. Vallentin. The positive semidefinite Grothendieck problem with rank constraint. In *Automata, Languages and Programming, 37th International Colloquium, Part I*, pages 31–42, 2010.
- [25] M. Charikar, V. Guruswami, and A. Wirth. Clustering with qualitative information. *J. Comput. System Sci.*, 71(3):360–383, 2005.
- [26] M. Charikar, K. Makarychev, and Y. Makarychev. On the advantage over random for maximum acyclic subgraph. In *48th Annual IEEE Symposium on Foundations of Computer Science*, pages 625–633, 2007.
- [27] M. Charikar and A. Wirth. Maximizing quadratic programs: extending Grothendieck’s inequality. In *45th Annual Symposium on Foundations of Computer Science*, pages 54–60. IEEE Computer Society, 2004.
- [28] R. Cleve, P. Høyer, B. Toner, and J. Watrous. Consequences and limits of nonlocal strategies. In *19th Annual IEEE Conference on Computational Complexity*, pages 236–249, 2004.
- [29] A. Coja-Oghlan, C. Cooper, and A. Frieze. An efficient sparse regularity concept. *SIAM J. Discrete Math.*, 23(4):2000–2034, 2009/10.
- [30] D. Conlon and J. Fox. Bounds for graph regularity and removal lemmas. Preprint available at <http://arxiv.org/abs/1107.4829>, 2011.
- [31] S. Cook. The P versus NP problem. In *The millennium prize problems*, pages 87–104. Clay Math. Inst., Cambridge, MA, 2006.
- [32] A. M. Davie. Matrix norms related to Grothendieck’s inequality. In *Banach spaces (Columbia, Mo., 1984)*, volume 1166 of *Lecture Notes in Math.*, pages 22–26. Springer, Berlin, 1985.

- [33] J. Diestel, J. H. Fourie, and J. Swart. *The metric theory of tensor products*. American Mathematical Society, Providence, RI, 2008. Grothendieck’s résumé revisited.
- [34] J. Diestel, H. Jarchow, and A. Tonge. *Absolutely summing operators*, volume 43 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1995.
- [35] R. G. Downey and M. R. Fellows. *Parameterized complexity*. Monographs in Computer Science. Springer-Verlag, New York, 1999.
- [36] P. Erdős and A. Rényi. On the evolution of random graphs. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 5:17–61, 1960.
- [37] P. C. Fishburn and J. A. Reeds. Bell inequalities, Grothendieck’s constant, and root two. *SIAM J. Discrete Math.*, 7(1):48–56, 1994.
- [38] A. Frieze and R. Kannan. Quick approximation to matrices and applications. *Combinatorica*, 19(2):175–220, 1999.
- [39] M. R. Garey and D. S. Johnson. *Computers and intractability*. W. H. Freeman and Co., San Francisco, Calif., 1979. A guide to the theory of NP-completeness, A Series of Books in the Mathematical Sciences.
- [40] D. J. H. Garling. *Inequalities: a journey into linear analysis*. Cambridge University Press, Cambridge, 2007.
- [41] S. Gerke and A. Steger. A characterization for sparse ϵ -regular pairs. *Electron. J. Combin.*, 14(1):Research Paper 4, 12 pp. (electronic), 2007.
- [42] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. Assoc. Comput. Mach.*, 42(6):1115–1145, 1995.
- [43] G. H. Golub and C. F. Van Loan. *Matrix computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, third edition, 1996.
- [44] W. T. Gowers. Lower bounds of tower type for Szemerédi’s uniformity lemma. *Geom. Funct. Anal.*, 7(2):322–337, 1997.
- [45] A. Grothendieck. Résumé de la théorie métrique des produits tensoriels topologiques. *Bol. Soc. Mat. São Paulo*, 8:1–79, 1953.
- [46] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, second edition, 1993.
- [47] V. Guruswami, R. Manokaran, and P. Raghavendra. Beating the random ordering is hard: Inapproximability of maximum acyclic subgraph. In *49th Annual IEEE Symposium on Foundations of Computer Science*, pages 573–582, 2008.
- [48] V. Guruswami, P. Raghavendra, R. Saket, and Y. Wu. Bypassing UGC from some optimal geometric inapproximability results. *Electronic Colloquium on Computational Complexity (ECCC)*, 17:177, 2010.
- [49] U. Haagerup. The Grothendieck inequality for bilinear forms on C^* -algebras. *Adv. in Math.*, 56(2):93–116, 1985.
- [50] U. Haagerup. A new upper bound for the complex Grothendieck constant. *Israel J. Math.*, 60(2):199–224, 1987.
- [51] J. Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859 (electronic), 2001.
- [52] J. Håstad and S. Venkatesh. On the advantage over a random assignment. *Random Structures Algorithms*, 25(2):117–149, 2004.
- [53] S. Heilman, A. Jagannath, and A. Naor. Solution of the propeller conjecture in \mathbb{R}^3 . Manuscript, 2011.
- [54] H. Heydari. Quantum correlation and Grothendieck’s constant. *J. Phys. A*, 39(38):11869–11875, 2006.
- [55] G. J. O. Jameson. *Summing and nuclear norms in Banach space theory*, volume 8 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1987.
- [56] W. B. Johnson and J. Lindenstrauss. Basic concepts in the geometry of Banach spaces. In *Handbook of the geometry of Banach spaces, Vol. I*, pages 1–84. North-Holland, Amsterdam, 2001.
- [57] D. B. Judin and A. S. Nemirovskii. Informational complexity and effective methods for the solution of convex extremal problems. *Ėkonom. i Mat. Metody*, 12(2):357–369, 1976.
- [58] F. Juhász. The asymptotic behaviour of Lovász’ θ function for random graphs. *Combinatorica*, 2(2):153–155, 1982.
- [59] D. Karger, R. Motwani, and M. Sudan. Approximate graph coloring by semidefinite programming. *J. ACM*, 45(2):246–265, 1998.

- [60] B. S. Kashin and S. Ī. Sharek. On the Gram matrices of systems of uniformly bounded functions. *Tr. Mat. Inst. Steklova*, 243(Funkts. Prostran., Priblizh., Differ. Uravn.):237–243, 2003.
- [61] J. Kempe, H. Kobayashi, K. Matsumoto, B. Toner, and T. Vidick. Entangled games are hard to approximate. In *49th Annual IEEE Symposium on Foundations of Computer Science*, pages 447–456, 2008.
- [62] S. Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, pages 767–775 (electronic), New York, 2002. ACM.
- [63] S. Khot. On the unique games conjecture (invited survey). In *Proceedings of the 25th Annual IEEE Conference on Computational Complexity*, pages 99–121, 2010.
- [64] S. Khot, G. Kindler, E. Mossel, and R. O’Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM J. Comput.*, 37(1):319–357 (electronic), 2007.
- [65] S. Khot and A. Naor. Linear equations modulo 2 and the L_1 diameter of convex bodies. *SIAM J. Comput.*, 38(4):1448–1463, 2008.
- [66] S. Khot and A. Naor. Approximate kernel clustering. *Mathematika*, 55(1-2):129–165, 2009.
- [67] S. Khot and A. Naor. Sharp kernel clustering algorithms and their associated Grothendieck inequalities. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 664–683, 2010. Full version to appear in *Random Structures and Algorithms*.
- [68] S. Khot and R. O’Donnell. SDP gaps and UGC-hardness for max-cut-gain. *Theory Comput.*, 5:83–117, 2009.
- [69] S. Khot and S. Safra. A two prover one round game with strong soundness. To appear in 52nd Annual IEEE Symposium on Foundations of Computer Science, 2011.
- [70] G. Kindler, A. Naor, and G. Schechtman. The UGC hardness threshold of the L_p Grothendieck problem. *Math. Oper. Res.*, 35(2):267–283, 2010.
- [71] Y. Kohayakawa. Szemerédi’s regularity lemma for sparse graphs. In *Foundations of computational mathematics (Rio de Janeiro, 1997)*, pages 216–230. Springer, Berlin, 1997.
- [72] Y. Kohayakawa and V. Rödl. Szemerédi’s regularity lemma and quasi-randomness. In *Recent advances in algorithms and combinatorics*, volume 11 of *CMS Books Math./Ouvrages Math. SMC*, pages 289–351. Springer, New York, 2003.
- [73] Y. Kohayakawa, V. Rödl, and L. Thoma. An optimal algorithm for checking regularity. *SIAM J. Comput.*, 32(5):1210–1235 (electronic), 2003.
- [74] H. König. On the complex Grothendieck constant in the n -dimensional case. In *Geometry of Banach spaces (Strobl, 1989)*, volume 158 of *London Math. Soc. Lecture Note Ser.*, pages 181–198. Cambridge Univ. Press, Cambridge, 1990.
- [75] H. König. Some remarks on the Grothendieck inequality. In *General inequalities, 6 (Oberwolfach, 1990)*, volume 103 of *Internat. Ser. Numer. Math.*, pages 201–206. Birkhäuser, Basel, 1992.
- [76] H. König. On an extremal problem originating in questions of unconditional convergence. In *Recent progress in multivariate approximation (Witten-Bommerholz, 2000)*, volume 137 of *Internat. Ser. Numer. Math.*, pages 185–192. Birkhäuser, Basel, 2001.
- [77] J.-L. Krivine. Sur la constante de Grothendieck. *C. R. Acad. Sci. Paris Sér. A-B*, 284(8):A445–A446, 1977.
- [78] J.-L. Krivine. Constantes de Grothendieck et fonctions de type positif sur les sphères. *Adv. in Math.*, 31(1):16–30, 1979.
- [79] M. Laurent and A. Varvitsiotis. Computing the Grothendieck constant of some graph classes. Preprint available at <http://arxiv.org/abs/1106.2735>, 2011.
- [80] T. Lee, G. Schechtman, and A. Shraibman. Lower bounds on quantum multiparty communication complexity. In *Proceedings of the 24th Annual IEEE Conference on Computational Complexity*, pages 254–262, 2009.
- [81] T. Lee and A. Shraibman. Lower bounds in communication complexity. *Found. Trends Theor. Comput. Sci.*, 3(4):front matter, 263–399 (2009), 2007.
- [82] A. S. Lewis and M. L. Overton. Eigenvalue optimization. In *Acta numerica, 1996*, volume 5 of *Acta Numer.*, pages 149–190. Cambridge Univ. Press, Cambridge, 1996.
- [83] J. Lindenstrauss and A. Pełczyński. Absolutely summing operators in L_p -spaces and their applications. *Studia Math.*, 29:275–326, 1968.

- [84] N. Linial, S. Mendelson, G. Schechtman, and A. Shraibman. Complexity measures of sign matrices. *Combinatorica*, 27(4):439–463, 2007.
- [85] N. Linial and A. Shraibman. Learning complexity vs. communication complexity. *Combin. Probab. Comput.*, 18(1-2):227–245, 2009.
- [86] N. Linial and A. Shraibman. Lower bounds in communication complexity based on factorization norms. *Random Structures Algorithms*, 34(3):368–394, 2009.
- [87] L. Lovász. On the Shannon capacity of a graph. *IEEE Trans. Inform. Theory*, 25(1):1–7, 1979.
- [88] L. Lovász and M. D. Plummer. *Matching theory*, volume 121 of *North-Holland Mathematics Studies*. North-Holland Publishing Co., Amsterdam, 1986. Annals of Discrete Mathematics, 29.
- [89] L. Lovász and B. Szegedy. Szemerédi’s lemma for the analyst. *Geom. Funct. Anal.*, 17(1):252–270, 2007.
- [90] D. W. Matula. On the complete subgraphs of a random graph. In *Proc. Second Chapel Hill Conf. on Combinatorial Mathematics and its Applications (Univ. North Carolina, Chapel Hill, N.C., 1970)*, pages 356–369. Univ. North Carolina, Chapel Hill, N.C., 1970.
- [91] A. Megretski. Relaxations of quadratic programs in operator theory and system analysis. In *Systems, approximation, singular integral operators, and related topics (Bordeaux, 2000)*, volume 129 of *Oper. Theory Adv. Appl.*, pages 365–392. Birkhäuser, Basel, 2001.
- [92] A. Naor and G. Schechtman. An approximation scheme for quadratic form maximization on convex bodies. Manuscript, 2009.
- [93] A. Nemirovski. Advances in convex optimization: conic programming. In *International Congress of Mathematicians. Vol. I*, pages 413–444. Eur. Math. Soc., Zürich, 2007.
- [94] A. Nemirovski, C. Roos, and T. Terlaky. On maximization of quadratic form over intersection of ellipsoids with common center. *Math. Program.*, 86(3, Ser. A):463–473, 1999.
- [95] Y. Nesterov. Semidefinite relaxation and nonconvex quadratic optimization. *Optim. Methods Softw.*, 9(1-3):141–160, 1998.
- [96] Y. Nesterov, H. Wolkowicz, and Y. Ye. Semidefinite programming relaxations of nonconvex quadratic optimization. In *Handbook of semidefinite programming*, volume 27 of *Internat. Ser. Oper. Res. Management Sci.*, pages 361–419. Kluwer Acad. Publ., Boston, MA, 2000.
- [97] C. H. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *J. Comput. System Sci.*, 43(3):425–440, 1991.
- [98] D. Pérez-García, M. M. Wolf, C. Palazuelos, I. Villanueva, and M. Junge. Unbounded violation of tripartite Bell inequalities. *Comm. Math. Phys.*, 279(2):455–486, 2008.
- [99] G. Pisier. Grothendieck’s theorem for noncommutative C^* -algebras, with an appendix on Grothendieck’s constants. *J. Funct. Anal.*, 29(3):397–415, 1978.
- [100] G. Pisier. *Factorization of linear operators and geometry of Banach spaces*, volume 60 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC, 1986.
- [101] G. Pisier. Grothendieck’s theorem, past and present. Preprint available at <http://arxiv.org/abs/1101.4195>, 2011.
- [102] I. Pitowsky. New Bell inequalities for the singlet state: going beyond the Grothendieck bound. *J. Math. Phys.*, 49(1):012101, 11, 2008.
- [103] P. Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 245–254, 2008.
- [104] P. Raghavendra and D. Steurer. Towards computing the Grothendieck constant. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 525–534, 2009.
- [105] J. A. Reeds. A new lower bound on the real Grothendieck constant. Unpublished manuscript, available at <http://www.dtc.umn.edu/reedsj/bound2.dvi>, 1991.
- [106] O. Regev and B. Toner. Simulating quantum correlations with finite communication. *SIAM J. Comput.*, 39(4):1562–1580, 2009/10.
- [107] R. E. Rietz. A proof of the Grothendieck inequality. *Israel J. Math.*, 19:271–276, 1974.
- [108] M. Sipser. *Introduction to the theory of computation*. PWS Publishing Company, 1997.
- [109] R. R. Smith. Completely bounded multilinear maps and Grothendieck’s inequality. *Bull. London Math. Soc.*, 20(6):606–612, 1988.

- [110] L. Song, A. Smola, A. Gretton, and K. A. Borgwardt. A dependence maximization view of clustering. In *Proceedings of the 24th international conference on Machine learning*, pages 815 – 822, 2007.
- [111] E. Szemerédi. Regular partitions of graphs. In *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976)*, volume 260 of *Colloq. Internat. CNRS*, pages 399–401. CNRS, Paris, 1978.
- [112] A. Tonge. The von Neumann inequality for polynomials in several Hilbert-Schmidt operators. *J. London Math. Soc. (2)*, 18(3):519–526, 1978.
- [113] A. Tonge. The complex Grothendieck inequality for 2×2 matrices. *Bull. Soc. Math. Grèce (N.S.)*, 27:133–136, 1986.
- [114] B. S. Tsirelson. Quantum analogues of Bell’s inequalities. The case of two spatially divided domains. *Zap. Nauchn. Sem. Leningrad. Otdel. Mat. Inst. Steklov. (LOMI)*, 142:174–194, 200, 1985. Problems of the theory of probability distributions, IX.
- [115] N. T. Varopoulos. On an inequality of von Neumann and an application of the metric theory of tensor products to operators theory. *J. Functional Analysis*, 16:83–100, 1974.

COURANT INSTITUTE, NEW YORK UNIVERSITY, 251 MERCER STREET, NEW YORK NY 10012, USA
E-mail address: khot@cims.nyu.edu

COURANT INSTITUTE, NEW YORK UNIVERSITY, 251 MERCER STREET, NEW YORK NY 10012, USA
E-mail address: naor@cims.nyu.edu