Submitted to J. ACM.

Improved Approximation Algorithms for Maximum Cut and Satisfiability Problems Using Semidefinite Programming*

Michel X. Goemans† M.I.T.

David P. Williamson‡ Cornell University

Abstract

We present randomized approximation algorithms for the maximum cut (MAX CUT) and maximum 2-satisfiability (MAX 2SAT) problems that always deliver solutions of expected value at least .87856 times the optimal value. These algorithms use a simple and elegant technique that randomly rounds the solution to a nonlinear programming relaxation. This relaxation can be interpreted both as a semidefinite program and as an eigenvalue minimization problem. We then show how to derandomize the algorithm to obtain approximation algorithms with the same performance guarantee of .87856. The previous best-known approximation algorithms for these problems had performance guarantees of $\frac{1}{2}$ for MAX CUT and $\frac{3}{4}$ for MAX 2SAT. Slight extensions of our analysis lead to a .79607-approximation algorithm for the maximum directed cut problem (MAX DICUT) and a .758-approximation algorithm for MAX SAT, where the previous best-known approximation algorithms had performance guarantees of $\frac{1}{4}$ and $\frac{3}{4}$ respectively. Our algorithm gives the first substantial progress in approximating MAX CUT in nearly twenty years, and represents the first use of semidefinite programming in the design of approximation algorithms.

Introduction

Given an undirected graph $G = (V, E)$ and non-negative weights $w_{ij} = w_{ji}$ on the edges $(i, j) \in E$, the maximum cut problem (MAX CUT) is that of finding the set of vertices S that maximizes the weight of the edges in the *cut* $(S, \bar{S})$; that is, the weight of the edges with one endpoint in S and the other in $\bar{S}$. For simplicity, we usually set $w_{ij} = 0$ for $(i, j) \notin E$ and denote the weight of a cut $(S, \bar{S})$ by $w(S, \bar{S}) = \sum_{i \in S, j \notin S} w_{ij}$. The MAX CUT problem is one of the Karp's original NP-complete problems [27], and has long been known to be NP-complete even if the problem is unweighted; that is, if $w_{ij} = 1$ for all $(i, j) \in E$ [13]. The MAX CUT problem is solvable in polynomial time for some special classes of graphs (e.g. if the graph is planar [38, 19]). Besides its theoretical importance, the MAX CUT problem has applications in circuit layout design and statistical physics (Barahona et al. [4]). For a comprehensive survey of the MAX CUT problem, the reader is referred to Poljak and Tuza [48].

*A preliminary version has appeared in the Proceedings of the 26th Annual ACM Symposium on Theory of Computing, Montréal, pages 422–431, 1994.

†Address: Dept. of Mathematics, Room 2-372, M.I.T., Cambridge, MA 02139. Email: goemans@math.mit.edu. Research supported in part by NSF contract 9302476-CCR, Air Force contract F49620-92-J-0125 and DARPA contract N00014-92-J-1799.

‡Address: School of Operations Research and Industrial Engineering, 237 ETC Building, Cornell University, Ithaca, NY 14853. Email: dpw@cs.cornell.edu. Research supported by an NSF Postdoctoral Fellowship. This research was conducted while the author was visiting MIT.

Because it is unlikely that there exist efficient algorithms for NP-hard maximization problems, a typical approach to solving such a problem is to find a *ρ-approximation algorithm*; that is, a polynomial-time algorithm that delivers a solution of value at least ρ times the optimal value. The constant ρ is sometimes called the *performance guarantee* of the algorithm. In 1976, Sahni and Gonzales [52] presented a $\frac{1}{2}$-approximation algorithm for the MAX CUT problem. Their algorithm iterates through the vertices and decides whether or not to assign vertex i to S based on which placement maximizes the weight of the cut of vertices 1 to i. This algorithm is essentially equivalent to the randomized algorithm that flips an unbiased coin for each vertex to decide which vertices are assigned to the set S. Since 1976, a number of researchers have presented approximation algorithms for the unweighted MAX CUT problem with performance guarantees of $\frac{1}{2} + \frac{1}{2m}$ [54], $\frac{1}{2} + \frac{n-1}{4m}$ [47], and $\frac{1}{2} + \frac{1}{2n}$ [22] (where $n = |V|$ and $m = |E|$), but no progress was made in improving the constant in the performance guarantee beyond that of Sahni and Gonzales's straightforward algorithm.

We present a simple, randomized $(\alpha - \epsilon)$-approximation algorithm for the maximum cut problem where

$$\alpha = \min_{0 \leq \theta \leq \pi} \frac{2}{\pi} \frac{\theta}{1 - \cos\theta} > 0.87856,$$

and ϵ is any positive scalar. The algorithm represents the first substantial progress in approximating the MAX CUT problem in nearly twenty years. We also show how to derandomize the algorithm to obtain a deterministic $(\alpha - \epsilon)$-approximation algorithm, for any $\epsilon > 0$. The algorithm for MAX CUT also leads directly to randomized and deterministic $(\alpha - \epsilon)$-approximation algorithms for the maximum 2-satisfiability problem (MAX 2SAT). The previously best known algorithm for this problem has a performance guarantee of $\frac{3}{4}$ and is due to Yannakakis [56] (see also Goemans and Williamson [15]). The improved 2SAT algorithm leads to .7584-approximation algorithm for the overall MAX SAT problem, fractionally better than Yannakakis's $\frac{3}{4}$-approximation algorithm for MAX SAT. Finally, a slight extension of our analysis yields a .79607-approximation algorithm for the maximum directed cut problem (MAX DICUT). The previous best known algorithm for MAX DICUT has a performance guarantee of $\frac{1}{4}$ [41].

Our algorithm depends on a means of randomly rounding a solution to a nonlinear relaxation. This relaxation can either be seen as a *semidefinite program* or as an *eigenvalue minimization problem*. To our knowledge, this is the first time that semidefinite programs have been used in the design and analysis of approximation algorithms. The relaxation plays a crucial role in allowing us to obtain a better performance guarantee: previous approximation algorithms compared the value of the solution obtained to the total sum of the weights $\sum_{i<j} w_{ij}$.

A semidefinite program is the optimization problem of a linear function of a symmetric matrix subject to linear equality constraints and the constraint that the matrix be positive semidefinite. Semidefinite programming is a special case of convex programming and also of the so-called *linear programming over cones* or *cone-LP* since the set of positive semidefinite matrices constitutes a convex cone. To some extent, semidefinite programming is very similar to linear programming; see Alizadeh [1] for a comparison. It inherits the very elegant duality theory of cone-LP (see Wolkowicz [55] and the exposition by Alizadeh [1]). The simplex method and primal-dual algorithms can be generalized to semidefinite programs (Pataki [42]). Given any $\epsilon > 0$, semidefinite programs can be solved within an additive error of ϵ in polynomial time (ϵ is part of the input, so the running time dependence on ϵ is polynomial in $\log \frac{1}{\epsilon}$). This can be done through the ellipsoid algorithm (Grötschel et al. [18]) and other polynomial-time algorithms for convex programming (Vaidya [53]) as well as interior-point methods (Nesterov and Nemirovskii [36, 37] and Alizadeh [1]). To terminate in polynomial time, these algorithms implicitly assume some requirement on the feasible space or on the size of the optimum solution; for details see Grötschel et al. [18] and

Section 3.3 of Alizadeh [1].

The importance of semidefinite programming is that it leads to tighter relaxations than the classical linear programming relaxations for many graph and combinatorial problems. A beautiful application of semidefinite programming is the work of Lovász [30] on the Shannon capacity of a graph. In conjunction with the polynomial-time solvability of semidefinite programs, this leads to the only known polynomial-time algorithm for finding the largest stable set (or the largest clique) in a perfect graph (Grötschel et al. [17]). More recently, there has been increased interest in semidefinite programming [33, 34, 1, 42, 45, 12, 32]. This started with the work of Lovász and Schrijver [33, 34], who developed a machinery to define tighter and tighter relaxations of any integer program based on quadratic and semidefinite programming. Their paper demonstrated the wide applicability and the power of semidefinite programming for combinatorial optimization problems. A consequence of our algorithm is a proof that a semidefinite program is a much tighter relaxation in the worst case than any known linear programming relaxation for the MAX CUT problem.

An eigenvalue minimization problem consists of minimizing a linear combination of the k largest eigenvalues of a matrix subject to equality constraints on the matrix. These problems can be solved in polynomial time by the ellipsoid algorithm [18], since the objective function can be seen to be convex. There is an abundant literature on spectral bounds for combinatorial optimization problems; see the survey paper by Mohar and Poljak [35]. Building on work by Overton and Womersley [40, 39], Alizadeh [1] has shown that eigenvalue minimization problems can be formulated as semidefinite programs.

For MAX CUT, the nonlinear relaxation we consider is equivalent to a spectral bound proposed by Delorme and Poljak [8, 7]. This equivalence to the semidefinite program we consider was established by Poljak and Rendl [45]. As shown by Poljak and Rendl [44, 46] and Delorme and Poljak [9], the spectral bound provides a very good bound on the maximum cut in practice. Delorme and Poljak [8, 7] study the worst-case ratio between the maximum cut and their spectral bound. The worst instance they are aware of is the 5-cycle for which the ratio is $\frac{32}{25+5\sqrt{5}} = 0.88445\ldots$, but they were unable to prove a bound better than 0.5 in the worst-case. Our result implies a worst-case bound of $\alpha = 0.87856\ldots$, very close to the bound for the 5-cycle.

The above discussion on the worst-case behavior indicates that straightforward modifications of our technique will not lead to significant improvements in the MAX CUT result. Furthermore, MAX CUT, MAX 2SAT, and MAX DICUT are MAX SNP-hard [41], and so it is known that there exists a constant $c < 1$ such that a c-approximation algorithm for any of these problems would imply that $P = NP$ [2]. Nevertheless, since the appearance of an abstract of this paper, Feige and Goemans [11] have extended our technique to yield a .931-approximation algorithm for MAX 2SAT and a .859-approximation algorithm for MAX DICUT. In addition, by using semidefinite programming and similar rounding ideas, Karger, Motwani, and Sudan [26] have been able to show how to color a k-colorable graph with $\tilde{O}(n^{1-\frac{3}{k+1}})$ colors in polynomial time. Thus it seems likely that the techniques in this paper will continue to prove useful in designing approximation algorithms.

We expect that in practice the performance of our algorithms will be much better than the worst-case bounds. We have performed some preliminary computational experiments with the MAX CUT algorithm which show that on a number of different types of random graphs the algorithm is usually within .96 of the optimal solution.

The abstract is structured as follows. We present the randomized algorithm for MAX CUT in Section 1, and its analysis in Section 2. We elaborate on our semidefinite programming bound and its relationship with other work on the MAX CUT problem in Section 3. The quality of the relaxation is investigated in Section 4, where we also present computational results. Section

5 presents the derandomization of the algorithm. In Section 6, we show how to extend the algorithm to an algorithm for MAX 2SAT, MAX SAT, MAX DICUT, and other problems. We conclude with a few remarks and open problems in Section 7.

1 The Randomized Approximation Algorithm for MAX CUT

Given a graph with vertex set $V = \{1, \ldots, n\}$ and non-negative weights $w_{ij} = w_{ji}$ for each pair of vertices i and j, the weight of the maximum cut $w(S, \bar{S})$ is given by the following integer quadratic program:

$$\begin{array}{lll} & \text{Maximize} & \frac{1}{2}\sum_{i<j} w_{ij}(1 - y_i y_j) \\ (Q) & \text{subject to:} & y_i \in \{-1, 1\} \qquad \forall i \in V. \end{array}$$

To see this, note that the set $S = \{i | y_i = 1\}$ corresponds to a cut of weight $w(S, \bar{S}) = \frac{1}{2}\sum_{i<j} w_{ij}(1 - y_i y_j)$.

Since solving this integer quadratic program is NP-complete, we consider relaxations. We can interpret (Q) as restricting y_i to be a 1-dimensional vector of unit norm. Some very interesting relaxations can be defined by allowing y_i to be a multi-dimensional vector v_i of unit norm. Since the linear space spanned by the vectors v_i has dimension at most n, we can assume that these vectors belong to $\mathbb{R}^n$ (or $\mathbb{R}^m$ for some $m \leq n$), or more precisely to the n-dimensional unit sphere S_n (or S_m for $m \leq n$). To ensure that the resulting optimization problem is indeed a relaxation, we need to define the objective function in such a way that it reduces to $\frac{1}{2}\sum_{i<j} w_{ij}(1 - y_i y_j)$ in the case of vectors lying in a 1-dimensional space. There are two very natural ways of guaranteeing this property. One can either replace $(1 - y_i y_j)$ by $(1 - v_i \cdot v_j)$ where $v_i \cdot v_j$ represents the inner product (or dot product) of v_i and v_j, or by $\frac{2\arccos(v_i \cdot v_j)}{\pi}$ which corresponds to the angle between v_i and v_j scaled to range between 0 and 2. The resulting relaxations are denoted by (P) and (R):

$$\begin{array}{lll} & \text{Maximize} & \frac{1}{2}\sum_{i<j} w_{ij}(1 - v_i \cdot v_j) \\ (P) & \text{subject to:} & v_i \in S_n \qquad \forall i \in V \end{array}$$

and

$$\begin{array}{lll} & \text{Maximize} & \sum_{i<j} w_{ij} \frac{\arccos(v_i \cdot v_j)}{\pi} \\ (R) & \text{subject to:} & v_i \in S_n \qquad \forall i \in V. \end{array}$$

Relaxation (P) plays a crucial role in the design of our approximation algorithm, while the analysis of the algorithm will show that (R) is in fact equivalent to the MAX CUT problem. We will show in Section 3 we can solve the relaxation (P) using semidefinite programming. We can now present our simple randomized algorithm for the MAX CUT problem.

1. Solve (P), obtaining an optimal set of vectors v_i.
2. Let r be a vector uniformly distributed on the unit sphere S_n.
3. Set $S = \{i | v_i \cdot r \geq 0\}$.

In other words, we choose a random hyperplane through the origin (with r as its normal) and partition the vertices into those vectors that lie "above" the plane (i.e. have a non-negative inner product with r) and those that lie "below" it (i.e. have a negative inner product with r). The motivation for this randomized step comes from the fact that (P) (and (R) too) is independent of the coordinate system: applying any orthonormal transformation to a set of vectors results in a solution with the same objective value.

Let W denote the value of the cut produced in this way, and $E[W]$ its expectation. We will show in Theorem 2.1 in Section 2 that, given any set of vectors $v_i \in S_n$, the expected weight of the cut defined by a random hyperplane is

$$E[W] = \sum_{i<j} w_{ij} \frac{\arccos(v_i \cdot v_j)}{\pi}.$$

Therefore there must exist a cut of value at least $E[W]$, implying (somewhat surprisingly) that the relaxation (R) is in fact a reformulation of the MAX CUT problem. Moreover, we will also show that

$$E[W] \geq \alpha \cdot \frac{1}{2} \sum_{i<j} w_{ij}(1 - v_i \cdot v_j),$$

where $\alpha = \min_{0 \leq \theta \leq \pi} \frac{2}{\pi} \frac{\theta}{1-\cos\theta} > .878$. If Z^*_{MC} is the optimal value of the maximum cut and Z^*_P is the optimal value of the relaxation (P), then since the expected weight of the cut generated by the algorithm is equal to $E[W] \geq \alpha Z^*_P \geq \alpha Z^*_{MC}$, the algorithm has a performance guarantee of α for the MAX CUT problem.

We must argue that the algorithm can be implemented in polynomial time. For simplicity, we assume that the weights are integral and we denote throughout the paper the sum $\sum_{i<j} w_{ij}$ of the weights by W_{tot}. In Section 3, we describe two equivalent formulations to (P). One is based on semidefinite programming, while the other is defined in terms of eigenvalues. Strictly speaking, we cannot solve (P) or any of its reformulations to optimality in polynomial time; both the vectors and the optimal value Z^*_P might in fact be irrational. However, using an algorithm for semidefinite programming (see introduction), one can obtain, for any $\epsilon > 0$, a set of vectors v_i's of value greater than $Z^*_P - \epsilon$ in time polynomial in the input size and $\log \frac{1}{\epsilon}$. For example, Alizadeh's adaptation of Ye's interior-point algorithm to semidefinite programming [1] performs $O(\sqrt{n}(\log W_{tot} + \log \frac{1}{\epsilon}))$ iterations, each requiring $O(n^3)$ time. On these approximately optimal vectors, the randomized algorithm will produce a cut of expected value greater than or equal to $\alpha(Z^*_P - \epsilon) \geq (\alpha - \epsilon)Z^*_{MC}$. The point on the unit sphere S_n can be generated by drawing n values $x_1, x_2, \ldots, x_n$ independently from the standard normal distribution, and normalizing the vector obtained (see Knuth [28, p. 130]); for our purposes, there is no need to normalize the resulting vector x. The standard normal distribution can be simulated using the uniform distribution between 0 and 1 (see Knuth [28, p. 117]). The algorithm is thus a randomized $(\alpha - \epsilon)$-approximation algorithm for MAX CUT, whose running time is $O(n^{3.5}(\log W_{tot} + \log \frac{1}{\epsilon}))$.

2 Analysis of the Algorithm

In this section, we analyze the performance of the algorithm. Let $\{v_1, \ldots, v_n\}$ be any vectors belonging to S_n, and let $E[W]$ be the expected value of the cut $w(S, \bar{S})$ produced by the randomized algorithm given in the previous section. We start by characterizing the expected value of the cut.

Theorem 2.1

$$E[W] = \frac{1}{\pi} \sum_{i<j} w_{ij} \arccos(v_i \cdot v_j).$$

Given a vector r drawn uniformly from the unit sphere S_n, we know by the linearity of expectations that

$$E[W] = \sum_{i<j} w_{ij} \cdot \Pr[\mathrm{sgn}(v_i \cdot r) \neq \mathrm{sgn}(v_j \cdot r)],$$

where $\mathrm{sgn}(x) = 1$ if $x \geq 0$, and -1 otherwise. Theorem 2.1 is thus implied by the following lemma.

Lemma 2.2

$$\Pr[sgn(v_i \cdot r) \neq sgn(v_j \cdot r)] = \frac{1}{\pi} \arccos(v_i \cdot v_j).$$

Proof: A restatement of the lemma is that the probability the random hyperplane separates the two vectors is directly proportional to the angle between the two vectors; that is, it is proportional to the angle $\theta = \arccos(v_i \cdot v_j)$. By symmetry, $\Pr[\mathrm{sgn}(v_i \cdot r) \neq \mathrm{sgn}(v_j \cdot r)] = 2 \Pr[v_i \cdot r \geq 0, v_j \cdot r < 0]$. The set $\{r : v_i \cdot r \geq 0, v_j \cdot r < 0\}$ corresponds to the intersection of two half-spaces whose dihedral angle is precisely θ; its intersection with the sphere is a spherical digon of angle θ and, by symmetry of the sphere, has thus measure equal to $\frac{\theta}{2\pi}$ times the measure of the full sphere. In other words, $\Pr[v_i \cdot r \geq 0, v_j \cdot r < 0] = \frac{\theta}{2\pi}$, and the lemma follows. ■

Our main theorem is the following. As we have argued above, this theorem applied to vectors v_i's of value greater than $Z_P^* - \epsilon$ implies that the algorithm has a performance guarantee of $\alpha - \epsilon$.

Theorem 2.3

$$E[W] \geq \alpha \frac{1}{2} \sum_{i<j} w_{ij}(1 - v_i \cdot v_j),$$

where $\alpha = \min_{0 \leq \theta \leq \pi} \frac{2}{\pi} \frac{\theta}{1-\cos\theta} > 0.87856$.

By using Theorem 2.1 and the nonnegativity of w_{ij}, the result follows from the following lemma applied to $y = v_i \cdot v_j$.

Lemma 2.4 *For* $-1 \leq y \leq 1$, $\frac{1}{\pi} \arccos(y) \geq \alpha \cdot \frac{1}{2}(1 - y)$ *where* $\alpha = \min_{0 < \theta \leq \pi} \frac{2}{\pi} \frac{\theta}{1-\cos\theta} > 0.87856$.

Proof: The expression for α follows straightforwardly by using the change of variables $\cos\theta = y$. See Figure 1, part (i). Using simple calculus, one can see that α achieves its value for $\theta = 2.331122\ldots$, the non-zero root of $\cos\theta + \theta \sin\theta = 1$. To formally prove the bound of 0.87856, we first observe that $\frac{2}{\pi} \frac{\theta}{1-\cos\theta} \geq 1$ for $0 < \theta \leq \pi/2$. The concavity of $f(\theta) = 1 - \cos\theta$ in the range $\pi/2 \leq \theta \leq \pi$ implies that, for any θ_0, we have $f(\theta) \leq f(\theta_0) + (\theta - \theta_0) f'(\theta)$, or $1 - \cos\theta \leq 1 - \cos\theta_0 + (\theta - \theta_0)\sin\theta_0 = \theta \sin\theta_0 + (1 - \cos\theta_0 - \theta_0 \sin\theta_0)$. Choosing $\theta_0 = 2.331122$ for which $1 - \cos\theta_0 - \theta_0 \sin\theta_0 < 0$, we derive that $1 - \cos\theta < \theta \sin\theta_0$, implying that $\alpha > \frac{2}{\pi \sin\theta_0} > 0.87856$. ■

We can refine the analysis and prove a better performance guarantee whenever the value of the relaxation (P) constitutes a large fraction of the total weight.

Theorem 2.5 *Let* $R = (\sum_{i<j} w_{ij} \frac{1 - v_i \cdot v_j}{2})/W_{tot}$. *Let* $h(t) = \arccos(1 - 2t)/\pi$. *If* $R \geq \gamma = \arg\min_{0 < t \leq 1} \frac{h(t)}{t} = 0.84458\ldots$ *then*

$$E[W] \geq \frac{h(R)}{R} \sum_{i<j} w_{ij} \frac{1 - v_i \cdot v_j}{2}.$$

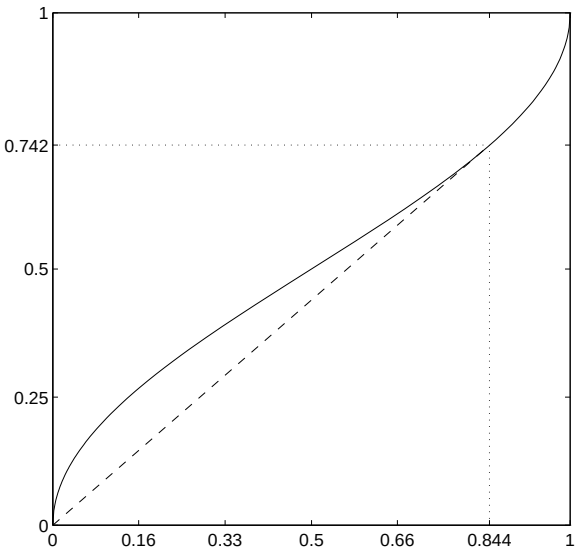

Figure 1: (i) Plot of $z = \arccos(y)/\pi$ as a function of $t = \frac{1}{2}(1-y)$. The ratio z/t is thus the slope of the line between $(0,0)$ and (z,t). The minimum slope is precisely $\alpha = 0.742/0.844$ and corresponds to the dashed line. (ii) The plot also represents $h(t) = \arccos(1-2t)/\pi$ as a function of t. As t approaches 1, $h(t)/t$ also approaches 1. (iii) The dashed line corresponds to $\tilde{h}$ between 0 and $\gamma > 0.844$.

The theorem implies a performance guarantee of $h(R)/R - \epsilon$ when $R \geq \gamma$. As R varies between γ and 1, one easily verifies that $h(R)/R$ varies between α and 1 (see Figure 1, part (ii)). *Proof*: Letting $\lambda_e = w_{ij}/W_{tot}$ and $x_e = \frac{1-v_i \cdot v_j}{2}$ for $e = (i,j)$, we can rewrite R as $R = \sum_e \lambda_e x_e$. The expected weight of the cut produced by the algorithm is equal to

$$E[W] = \sum_{i<j} w_{ij} \frac{\arccos(v_i \cdot v_j)}{\pi} = W_{tot} \sum_e \lambda_e \frac{\arccos(1-2x_e)}{\pi} = W_{tot} \sum_e \lambda_e h(x_e).$$

To bound $E[W]$, we evaluate

$$\begin{aligned} \text{Min} \quad & \sum_e \lambda_e h(x_e) \\ \text{subject to:} & \sum_e \lambda_e x_e = R \\ & 0 \leq x_e \leq 1. \end{aligned}$$

Consider the relaxation obtained by replacing $h(t)$ by the largest (pointwise) convex function $\tilde{h}(t)$ smaller or equal to $h(t)$. It is easy to see that $\tilde{h}(t)$ is linear with slope $\alpha = 0.87856\ldots$ between 0 and γ, and then equal to $h(t)$ for any t greater than γ. See Figure 1, part (iii). But

$$\sum_e \lambda_e h(x_e) \geq \sum_e \lambda_e \tilde{h}(x_e) \geq \tilde{h}(\sum_e \lambda_e x_e) \geq \tilde{h}(R) = h(R),$$

where we have used the fact that $\sum_e \lambda_e = 1$, $\lambda_e \geq 0$ and that $\tilde{h}$ is a convex function. This shows that

$$E[W] \geq W_{tot} h(R) = \frac{h(R)}{R} \sum_{i<j} w_{ij} \frac{1 - v_i \cdot v_j}{2},$$

proving the result. ■

So far, we have assumed that the weights are non-negative. If some of the weights are negative, the definition of the performance guarantee has to be modified since the optimum value could be positive or negative. We now give a corresponding generalization of Theorem 2.3 to arbitrary weights.

Theorem 2.6 *Let* $W_- = \sum_{i<j} w_{ij}^-$, *where* $x^- = \min(0, x)$. *Then*

$$\{E[W] - W_-\} \geq \alpha \left\{ \frac{1}{2} \sum_{i<j} w_{ij}(1 - v_i \cdot v_j) - W_- \right\}.$$

Proof: The quantity $E[W] - W_-$ can be written as

$$\sum_{i<j:w_{ij}>0} w_{ij} \frac{\arccos(v_i \cdot v_j)}{\pi} + \sum_{i<j:w_{ij}<0} |w_{ij}| \left(1 - \frac{\arccos(v_i \cdot v_j)}{\pi}\right).$$

Similarly, $\left\{\frac{1}{2}\sum_{i<j} w_{ij}(1 - v_i \cdot v_j) - W_-\right\}$ is equal to

$$\sum_{i<j:w_{ij}>0} w_{ij} \frac{1 - v_i \cdot v_j}{2} + \sum_{i<j:w_{ij}<0} |w_{ij}| \frac{1 + v_i \cdot v_j}{2}.$$

The result therefore follows from Lemma 2.4 and the following variation of it. ■

Lemma 2.7 *For* $-1 \leq z \leq 1, 1 - \frac{1}{\pi}\arccos(z) \geq \alpha \cdot \frac{1}{2}(1 + z)$.

Proof: The lemma follows from Lemma 2.4 by using a change of variables $z = -y$ and noting that $\pi - \arccos(z) = \arccos(-z)$. ■

3 Relaxations and Duality

The relaxation (P) can be reformulated in several ways.

We begin by defining some terms and notation. All matrices under consideration are defined over the reals. An $n \times n$ matrix A is said to be positive semidefinite if for every vector $x \in \mathbb{R}^n$, $x^T A x \geq 0$. The following statements are equivalent for a symmetric matrix A (see e.g. [29]): (i) A is positive semidefinite, (ii) all eigenvalues of A are non-negative, and (iii) there exists a matrix B such that $A = B^T B$. In (iii), B can either be a (possibly singular) $n \times n$ matrix, or an $m \times n$ matrix for some $m \leq n$. Given a symmetric positive semidefinite matrix A, an $m \times n$ matrix B of full row-rank satisfying (iii) can be obtained in $O(n^3)$ time using an incomplete Cholesky decomposition [16, p. 90, P5.2-3].

Using the decomposition $Y = B^T B$, one can see that a positive semidefinite Y with $y_{ii} = 1$ corresponds precisely to a set of unit vectors $v_1, \ldots, v_n \in S_m$: simply correspond the vector v_i to the ith column of B. Then $y_{ij} = v_i \cdot v_j$. The matrix Y is known as the Gram matrix of $\{v_1, \ldots, v_n\}$ [29, p. 110]. Using this equivalence, we can reformulate (P) as a semidefinite program:

$$\begin{array}{lllll} & Z_P^* & = & \text{Max} & \frac{1}{2}\sum_{i<j} w_{ij}(1 - y_{ij}) \\ (SD) & & \text{subject to:} & y_{ii} = 1 & \forall i \in V \\ & & & Y \text{ symmetric positive semidefinite} & \end{array}$$

where $Y = (y_{ij})$. As previously mentioned, (SD) can be solved to within an additive error of ϵ in time $O(n^{3.5}(\log W_{tot} + \log \frac{1}{\epsilon}))$. Once an almost optimal solution to (SD) is found, one can use an incomplete Cholesky decomposition to obtain vectors $v_1, \ldots, v_n \in S_m$ for some $m \leq n$ such that $\frac{1}{2}\sum_{i<j} w_{ij}(1 - v_i \cdot v_j) \geq Z_P^* - \epsilon$.

Among all optimum solutions to (SD), one can show the existence of a solution of low rank. This follows from a general statement about semidefinite programs which is implicit in Barvinok [5] and Pataki [43]: any extreme solution (i.e. which cannot be expressed as the strict convex combination of other feasible solutions) of a semidefinite program with k linear equalities has rank at most l where $\frac{l(l+1)}{2} \leq k$, i.e. $l \leq \frac{\sqrt{8k+1}-1}{2} < \sqrt{2k}$. This therefore implies that there exists a primal optimum solution Y^* to (SD) of rank less than $\sqrt{2n}$, and that the optimum vectors v_i of (P) can be embedded in $\mathbb{R}^m$ with $m < \sqrt{2n}$.

As mentioned in the introduction, there is an elegant duality theory for semidefinite programming. We now turn to discussing the dual of the program (SD). It is typical to assume that the objective function of a semidefinite program is symmetric. For this purpose, we can rewrite the objective function of (SD) as $\frac{1}{4}\sum_{i=1}^{n}\sum_{j=1}^{n} w_{ij}(1 - y_{ij})$, or even as $\frac{1}{2}W_{tot} - \frac{1}{4}\sum_i \sum_j w_{ij}y_{ij}$. In matrix form, the objective function can be conveniently written as $\frac{1}{2}W_{tot} - \frac{1}{4}Tr(WY)$, where $W = (w_{ij})$ and Tr denotes the trace.

The dual of (SD) has a very simple description:

$$
\begin{array}{lrll}
Z_D^* & = & \frac{1}{2}W_{tot} + \frac{1}{4}\text{Min} & \sum_i \gamma_i \\
(D) & & \text{subject to:} & W + diag(\gamma) \text{ positive semidefinite,}
\end{array}
$$

where $diag(\gamma)$ denotes the diagonal matrix whose ith diagonal entry is γ_i. The dual has a simple interpretation. Since $W + diag(\gamma)$ is positive semidefinite, it can be expressed as C^TC; in other words, the weight w_{ij} can be viewed as $c_i \cdot c_j$ for some vectors c_i's and $\gamma_i = c_i \cdot c_i = \|c_i\|^2$. The weight of any cut is thus $w(S, \bar{S}) = (\sum_{i \in S} c_i) \cdot \left(\sum_{j \notin S} c_j\right)$, which is never greater than

$$\left\| \frac{\sum_{i \in V} c_i}{2} \right\|^2 = \frac{1}{2}W_{tot} + \frac{1}{4}\sum_{i \in V} \gamma_i.$$

Showing weak duality between (P) and (D), namely that $Z_P^* \leq Z_D^*$, is easy. Consider any primal feasible solution Y and any dual vector γ. Since both Y and $W + diag(\gamma)$ are positive semidefinite, we derive that $Tr((diag(\gamma)+W)Y) \geq 0$ (see [29, p. 218, ex. 14]). But $Tr((diag(\gamma)+W)Y) = Tr(diag(\gamma)Y) + Tr(WY) = \sum_i \gamma_i + Tr(WY)$, implying that the difference of the dual objective function value and the primal objective function value is non-negative.

For semidefinite programs in their full generality, there is no guarantee that the primal optimum value is equal to the dual optimum value. Also, the maximum (resp. minimum) is in fact a supremum (resp. infimum) and there is no guarantee that the supremum (resp. infimum) is attained. These, however, are pathological cases. Our programs (SD) and (D) behave nicely; both programs attain their optimum values, and these values are equal (i.e. $Z_P^* = Z_D^*$). This can be shown in a variety of ways (see [45, 1, 5]); the key idea is to show that certain sets are closed in order to apply the separating hyperplane theorem.

Given that strong duality holds in our case, the argument showing weak duality implies that, for the optimum primal solution Y^* and the optimum dual solution γ^*, we have $Tr((diag(\gamma^*)+W)Y^*) = 0$. Since both $diag(\gamma^*)+W$ and Y^* are positive semidefinite, we derive that $(diag(\gamma^*)+W)Y^* = 0$ (see [29, p. 218, ex. 14]). This is the strong form of complementary slackness for semidefinite programs (see Alizadeh [1]); the component-wise product expressed by the trace is

replaced by matrix multiplication. This implies, for example, that Y^* and $diag(\gamma^*) + W$ share a system of eigenvectors and that $rank(Y^*) + rank(diag(\gamma^*) + W) \leq n$.

The relaxation (D) (and thus (P)) is also equivalent to an eigenvalue upper bound on the value of the maximum cut Z^*_{MC} introduced by Delorme and Poljak [8, 7]. To describe the bound, we first introduce some notation. The Laplacian matrix $L = (l_{ij})$ is defined by $l_{ij} = -w_{ij}$ for $i \neq j$ and $l_{ii} = \sum_{k=1}^{n} w_{ik}$. The maximum eigenvalue of a matrix A is denoted by $\lambda_{max}(A)$.

Lemma 3.1 *[[8]] Let $u \in \mathbb{R}^n$ satisfy $u_1 + \ldots + u_n = 0$. Then*

$$g(u) = \frac{n}{4}\lambda_{max}(L + diag(u))$$

is an upper bound on Z^*_{MC}.

The proof is simple. Let y be an optimal solution to the integer quadratic program (Q). Notice that $y^T L y = 4Z^*_{MC}$. By using the Rayleigh principle ($\lambda_{max}(M) = \max_{\|x\|=1} x^T M x$), we obtain

$$\begin{aligned} \lambda_{max}(L + diag(u)) &\geq \frac{y^T(L + diag(u))y}{y^T y} \\ &= \frac{1}{n}\left(y^T L y + \sum_{i=1}^{n} y_i^2 u_i\right) \\ &= \frac{4Z^*_{MC}}{n}, \end{aligned}$$

proving the lemma. A vector u satisfying $\sum_{i=1}^{n} u_i = 0$ is called a *correcting vector*. The bound proposed by Delorme and Poljak [8] is to optimize $g(u)$ over all correcting vectors:

$$\begin{aligned} Z^*_{EIG} &= \text{Inf} \quad g(u) \\ (EIG) \qquad &\text{subject to: } \sum_{i=1}^{n} u_i = 0. \end{aligned}$$

As mentioned in the introduction, eigenvalue minimization problems can be formulated as semidefinite programs. For MAX CUT, the equivalence between (SD) and (EIG) was established by Poljak and Rendl [45].

For completeness, we derive the equivalence between (EIG) and the dual (D). For the optimum dual vector γ^*, the smallest eigenvalue of $diag(\gamma^*) + W$ must be 0, since otherwise we could decrease all γ^*_i by some $\epsilon > 0$ and thus reduce the dual objective function. This can be rewritten as $\lambda_{max}(-W - diag(\gamma^*)) = 0$. Define λ as $(\sum_i \gamma^*_i + 2W_{tot})/n$. By definition, $Z^*_D = n\lambda/4$. Moreover, defining $u_i = \lambda - \gamma^*_i - \sum_j w_{ij}$, one easily verifies that $\sum_i u_i = 0$ and that $-W - diag(\gamma^*) = L + diag(u) - \lambda I$, implying that $\lambda_{max}(L + diag(u)) = \lambda$. This shows that $Z^*_{EIG} \leq Z^*_D$. The converse inequality follows by reversing the argument.

4 Quality of the Relaxation

Theorem 2.3 implies the following corollary:

Corollary 4.1 *For any instance of MAX CUT,*

$$\frac{Z^*_{MC}}{Z^*_P} \geq \alpha > 0.87856.$$

For the 5-cycle, Delorme and Poljak [8] have shown that $Z^*_{MC}/Z^*_{EIG} = \frac{32}{25+5\sqrt{5}} = 0.88445\ldots$, implying that our worst-case analysis is almost tight. One can obtain this bound from the relaxation (P) by observing that for the 5-cycle $1-2-3-4-5-1$, the optimal vectors lie in a 2-dimensional subspace and can be expressed as $v_i = (\cos(\frac{4i\pi}{5}), \sin(\frac{4i\pi}{5}))$ for $i = 1, \ldots, 5$ corresponding to $Z^*_P = \frac{5}{2}(1+\cos\frac{\pi}{5}) = \frac{25+5\sqrt{5}}{8}$. Since $Z^*_{MC} = 4$ for the 5-cycle, this yields the bound of Delorme and Poljak. Delorme and Poljak have shown that $Z^*_{MC}/Z^*_{EIG} \geq \frac{32}{25+5\sqrt{5}}$ holds for special subclasses of graphs, such as planar graphs or line graphs. However, they were unable to prove a bound better than 0.5 in the absolute worst-case.

Although the worst-case value of Z^*_{MC}/Z^*_P is not completely settled, we have constructed instances for which $E[W]/Z^*_P < 0.8786$, showing that the analysis of our algorithm is practically tight. These instances are based on strongly self-dual polytopes due to Lovász [31]. A polytope P in $\mathbb{R}^n$ is said to be strongly self-dual [31] if (i) P is inscribed in the unit sphere, (ii) P is circumscribed around the sphere with origin as center and with radius r for some $0 < r < 1$, and (iii) there is a bijection σ between vertices and facets of P such that, for every vertex v of P, the facet $\sigma(v)$ is orthogonal to the vector v. For example, in $\mathbb{R}^2$, the strongly self-dual polytopes are precisely the regular odd polygons. One can associate a graph $G = (V, E)$ to a self-dual polytope P: the vertex set V corresponds to the vertices of P and there is an edge (v, w) if w belongs to the facet $\sigma(v)$ (or, equivalently, v belongs to $\sigma(w)$). For the regular odd polygons, these graphs are simply the odd cycles. Because of conditions (ii) and (iii), the inner product $v \cdot w$ for any pair of adjacent vertices is equal to $-r$. As a result, a strongly self-dual polytope leads to a feasible solution of (P) of value $\frac{1+r}{2}W_{tot}$. Lovász [31] gives a recursive construction of a class of strongly self-dual polytopes. One can show that, by choosing the dimension n large enough, his construction leads to strongly self-dual polytopes for which r is arbitrarily close to the critical value giving a bound of α. However, it is unclear whether, in general, for such polytopes, non-negative weights can be selected such that the vectors given by the polytope constitute an *optimum* solution to (P). Even if this could be shown, this would only establish that the ratio $E[W]/Z^*_P$ can be arbitrarily close to α; this would not imply anything for the ratio Z^*_{MC}/Z^*_P. In our investigation of self-dual polytopes, we have found the following instances. For the graph of Figure 2 on 11 vertices, the ratio $E[W]/Z^*_P$ is less than .8796 in the unweighted case; moreover, weights can be selected such that the ratio drops to .8788. A more complicated weighted instance on 103 vertices can be shown to lead to a bound $E[W]/Z^*_P < 0.8786$.

In practice, however, we expect that the algorithm will perform much better than this worst-case bound. Poljak and Rendl [44, 46] (see also Delorme and Poljak [9]) report computational results showing that the bound Z^*_{EIG} is typically less than 2-5% and never worse than 8% away from Z^*_{MC}. We also performed our own computational experiments, in which the cuts computed by the algorithm were usually within 4% of the semidefinite bound Z^*_P, and never less than 9% from the bound. To implement the algorithm, we used code supplied by R. Vanderbei [50] for a special class of semidefinite programs. We use Vanderbei's code to solve the semidefinite program, then we generate 50 random vectors r. We output the best of the 50 cuts induced. We applied our code to a small subset of the instances considered by Poljak and Rendl [44]. In particular, we considered several different types of random graphs, as well as complete geometric graphs defined by Traveling Salesman Problem (TSP) instances from the TSPLIB (see Reinelt [49]).

For four different types of random graphs, we ran 50 instances on graphs of 50 vertices, 20 on graphs of size 100, and 5 on graphs of size 200. In the Type A random graph, each edge (i, j) is included with probability 1/2. In the Type B random graph, each edge is given a weight drawn uniformly from the interval [-50,50]; the ratio of Theorem 2.6 is used in reporting nearness to

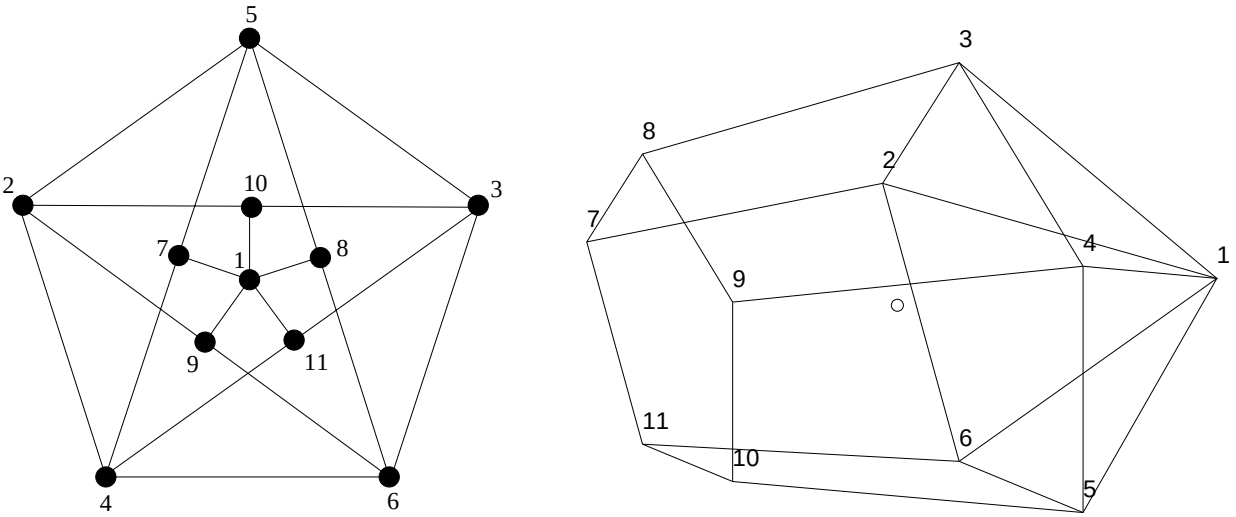

Figure 2: Graph on 11 vertices for which the ratio $E[W]/Z_P^*$ is less than .8796 in the unweighted case. The convex hull of the optimum vectors is depicted on the right; the circle represents the center of the sphere.

the semidefinite bound. In the Type C random graph of size $n \geq 10$, an edge (i, j) is included with probability $10/n$, leading to constant expected degree. Finally, in the Type D random graphs, an edge (i, j) is included with probability .1 if $i <= n/2$ and $j > n/2$ and probability .05 otherwise, leading to a large cut between the vertices in $[1, \ldots, n/2]$ and those in $[n/2+1, \ldots, n]$. We summarize the results of our experiments in Table 1 below. CPU Times are given in CPU seconds on a Sun SPARCstation 1.

In the case of the TSP instances, we used Euclidean instances from the TSPLIB, and set the edge weight w_{ij} to the Euclidean distance between the points i and j. We summarize our results in Table 2. In all 10 instances, we compute the optimal solution; for 5 instances, the value of the cut produced is equal to Z_P^*, and for the others, we have been able to exploit additional information from the dual semidefinite program to prove optimality (for the problem gr48, Poljak and Rendl [44] also show that our solution is optimal). For all TSPLIB instances, the maximum cut value is within .995 of the semidefinite bound. Given these computational results, it is tempting to speculate that a much stronger bound can be proven for these Euclidean instances. However, the instance defined by a unit length equilateral triangle has a maximum cut value of 2, but $Z_P^* = \frac{9}{4}$, for a ratio of $\frac{8}{9} = 0.8889$.

Homer and Peinado [24] have implemented our algorithm on a CM-5, and have shown that it produces optimal or very nearly optimal solutions to a number of MAX CUT instances derived from via minimization problems. These instances were provided by Michael Jünger [25] and have between 828 and 1366 vertices.

5 Derandomization

The algorithm can be derandomized, resulting in a deterministic $(\alpha - \epsilon)$-approximation algorithm for any $\epsilon > 0$. The method we use is the classical method of conditional expectations.

Assume we are given a set of unit vectors $v_1, \ldots, v_n$ spanning a linear space of dimension m.

Type of Graph	Size	Num Trials	Ave Int Gap	Ave CPU Time
Type A	50	50	.96988	36.28
	100	20	.96783	323.08
	200	5	.97209	4629.62
Type B	50	50	.97202	23.06
	100	20	.97097	217.42
	200	5	.97237	2989.00
Type C	50	50	.95746	23.53
	100	20	.94214	306.84
	200	5	.92362	2546.42
Type D	50	50	.95855	27.35
	100	20	.93984	355.32
	200	5	.93635	10709.42

Table 1: Summary of results of algorithm on random instances. Ave Int Gap is the average ratio of the value of the cut generated to the semidefinite bound, except for Type B graphs, where it is the ratio of the value of the cut generated minus the negative edge weights to the semidefinite bound minus the negative edge weights.

Instance	Size	SD Val	Cut Val	Time
dantzig42	42	42638	42638	43.35
gr120	120	2156775	2156667	754.87
gr48	48	321815	320277	26.17
gr96	96	105470	105295	531.50
hk48	48	771712	771712	66.52
kroA100	100	5897392	5897392	420.83
kroB100	100	5763047	5763047	917.47
kroC100	100	5890760	5890760	398.78
kroD100	100	5463946	5463250	469.48
kroE100	100	5986675	5986591	375.68

Table 2: Summary of results of algorithm on TSPLIB instances. SD Val is the value produced by the semidefinite relaxation. Cut Val is the value of the best cut output by the algorithm.

We can thus assume that the vectors v_i belong to S_m. Define $f(v_1, \ldots, v_n) = \frac{1}{\pi} \sum_{i<j} w_{ij} \arccos(v_i \cdot v_j) = E[W]$. In the randomized algorithm, we directly obtain from the v_i's a set of y_i's belonging to S_1 (i.e. scalars taking values in $\{-1, +1\}$), proving the existence of $y_1, \ldots, y_n$ such that $f(y_1, \ldots, y_n) \geq f(v_1, \ldots, v_n)$. In the deterministic algorithm, we shall decrease the dimension m one unit at a time.

Consider a vector r drawn uniformly from S_m. For simplicity of exposition, we shall assume that r is obtained by generating independent standard normal variables $x_1, \ldots, x_m$. We express the last two coordinates of x in polar coordinates: $(x_{m-1}, x_m) = (s \cos\theta, s \sin\theta)$. We restrict θ to lie in $[-\pi/2, \pi/2)$, implying that s is allowed to be negative. Observe that θ is undefined on a set of measure 0 (when $s = 0$). Although we don't really need this fact, it is easy to see that θ is uniformly distributed in the range $[-\frac{\pi}{2}, \frac{\pi}{2}]$. (Indeed, (x_{m-1}, x_m) is a bivariate normal variable and its density depends only on $s^2 = x_{m-1}^2 + x_m^2$.) We shall condition on the value of θ. For a given value of θ, let $f_\theta(v_1, \ldots, v_n) = E[W|\theta]$. We know that

$$f(v_1, \ldots, v_n) = E[W] = E_\theta[E[W|\theta]] = \int_{-\pi/2}^{\pi/2} f_\theta(v_1, \ldots, v_n) g(\theta) d\theta \leq \max_{-\frac{\pi}{2} \leq \theta \leq \frac{\pi}{2}} f_\theta(v_1, v_2, \ldots, v_n),$$

where $g(\theta)$ is the probability density function. As argued above, $g(\theta) = \frac{1}{\pi}$, but the fact that θ is uniformly distributed is not needed to derive the inequality.

We first describe a set of vectors $u_1(\theta), \ldots, u_n(\theta) \in S_{m-1}$ such that $f_\theta(v_1, \ldots, v_n) = f(u_1(\theta), \ldots, u_n(\theta))$ and, thus, for a fixed value of θ, f_θ can easily be evaluated. We will then describe a discrete set Θ of size polynomial in n and $1/\delta$ such that

$$\max_{-\frac{\pi}{2} \leq \theta \leq \frac{\pi}{2}} f_\theta(v_1, v_2, \ldots, v_n) - \delta W_{tot} \leq \max_{\theta \in \Theta} f_\theta(v_1, v_2, \ldots, v_n).$$

Selecting for θ the value maximizing f_θ over Θ, we obtain vectors $u_1, \ldots, u_n \in S_{m-1}$ such that $f(u_1, \ldots, u_n) \geq f(v_1, \ldots, v_n) - \delta W_{tot}$. We can then continue the process until the dimension is equal to 1. At that point, we have scalars $y_i \in \{-1, +1\}$ such that

$$f(y_1, \ldots, y_n) \geq f(v_1, \ldots, v_n) - n\delta W_{tot}.$$

We can assume without loss of generality that $f(v_1, \ldots, v_n) \geq W_{tot}/2$ since this is true for an orthonormal basis of n vectors. Thus

$$f(y_1, \ldots, y_n) \geq (1 - 2n\delta) f(v_1, \ldots, v_n).$$

Recalling that $f(v_1, \ldots, v_n) \geq \alpha(Z^*_{MC} - \epsilon) \geq \alpha(1-\epsilon) Z^*_{MC}$, we see that the derandomized algorithm achieves a performance guarantee of $(1 - 2n\delta)\alpha(1 - \epsilon)$, which can be made to be $(\alpha - \epsilon')$ for any $\epsilon' > 0$, by setting $\epsilon = \frac{\epsilon'}{2}$ and $\delta = \frac{\epsilon'}{4n}$. The running time of the resulting deterministic $(\alpha - \epsilon')$-approximation algorithm is thus polynomial in $\frac{1}{\epsilon'}$. Our derandomized algorithm heavily involves square roots and trigonometric functions. As usual, we implicitly assume that all these operations are computed with precision polynomial in the input size and $\log \frac{1}{\epsilon'}$, in order to guarantee polynomiality and to avoid losing in the performance guarantee.

We first show how to define $u_i(\theta)$ from v_i and θ. Given a vector $z \in \mathbb{R}^m$, we define, for any scalar a, $z|a$ to be the vector in $\mathbb{R}^{m-1}$ having coordinates $z_1, z_2, \ldots, z_{m-2}$ and a. Recall that $(x_{m-1}, x_m) = (s \cos\theta, s \sin\theta)$. We also express the last two coordinates of v_i in polar form: $(v_{i,m-1}, v_{im}) = (t_i \cos\gamma_i, t_i \sin\gamma_i)$, where $t_i = \sqrt{v_{i,m-1}^2 + v_{im}^2} \geq 0$ and $\gamma_i \in [0, 2\pi)$. We observe that

$$x_{m-1} v_{i,m-1} + x_m v_{im} = s t_i (\cos\gamma_i \cos\theta + \sin\gamma_i \sin\theta) = s t_i \cos(\gamma_i - \theta).$$

Setting $x' = x|s$ and

$$a_i(\theta) = v_i|[t_i \cos(\gamma_i - \theta)],$$

we derive that $x \cdot v_i = x' \cdot a_i(\theta)$. Let $u_i(\theta) = a_i(\theta)/\|a_i(\theta)\|$. To ensure that $u_i(\theta)$ is well-defined, we implicitly assume that the basis vectors are in general position, implying that $v_i|0$ is not the zero vector; at the end of the section, we will discuss how to remove this assumption. The above discussion shows that $\mathrm{sgn}(x \cdot v_i) = \mathrm{sgn}(x' \cdot u_i(\theta))$. Moreover, once we condition on θ, s is normally distributed (and independent of $x_1, \ldots, x_{m-2}$), implying that $x'/\|x'\|$ is uniformly distributed on S_{m-1}. Therefore, $f_\theta(v_1, \ldots, v_n)$ can be evaluated as $f(u_1(\theta), \ldots, u_n(\theta))$.

In order to compute approximately the maximum of $f(u_1(\theta), \ldots, u_n(\theta))$ over $\theta \in [-\pi/2, \pi/2]$, we discretize the interval in the following way. For every i, let $\alpha_i(\theta)$ denote the angle between the vectors $v_i|0$ and $a_i(\theta)$. Since $v_i|0$ and $a_i(\theta) - (v_i|0)$ are orthogonal, we derive that

$$\tan \alpha_i(\theta) = \frac{\|a_i(\theta) - (v_i|0)\|}{\|v_i|0\|} = \frac{t_i|\cos(\gamma_i - \theta)|}{\sqrt{\sum_{j=1}^{m-2} v_{ij}^2}},$$

and from this that $\alpha_i(\theta)$ ranges between 0 and

$$m_i = \arctan \frac{t_i}{\sqrt{\sum_{j=1}^{m-2} v_{ij}^2}} < \frac{\pi}{2}.$$

As θ varies, $\alpha_i(\theta)$ covers the interval $[0, m_i]$ twice, because the dependence of $\alpha_i(\theta)$ on θ can be stated in terms of $\cos(\gamma_i - \theta)$. Therefore, we can choose $2m_i/\nu \leq \pi/\nu$ values of θ, say $\theta_j^{(i)}$ for $j \leq 2m_i/\nu$, such that for all $\theta_j^{(i)} \leq \theta \leq \theta_{j+1}^{(i)}$ we have $|\alpha_i(\theta) - \alpha_i(\theta_j^{(i)})| \leq \nu$. Let Θ denote the union of these angles over $i = 1, \ldots, n$. Notice that $|\Theta| \leq \pi n/\nu$.

Lemma 5.1 *There exists Θ with $|\Theta| \leq \frac{2n}{\delta}$ such that*

$$\max_{-\frac{\pi}{2} \leq \theta \leq \frac{\pi}{2}} f_\theta(v_1, v_2, \ldots, v_n) - \delta W_{tot} \leq \max_{\theta \in \Theta} f_\theta(v_1, v_2, \ldots, v_n),$$

where $W_{tot} = \sum_{i<j} w_{ij}$.

Proof of Lemma 5.1: Let $\Theta = \{\theta_1 < \ldots < \theta_p\}$ be the set corresponding to $\nu = \delta\pi/2$. We claim that for any $i < j$ and any $\theta_k \leq \theta \leq \theta_{k+1}$, we have

$$|\arccos(u_i(\theta) \cdot u_j(\theta)) - \arccos(u_i(\theta_k) \cdot u_j(\theta_k))| \leq 2\nu = \delta\pi.$$

Summing the above inequality over all $i < j$ and using the definition of f, one derives the lemma.

To prove the claim, notice that the angle between $u_i(\theta)$ and $u_j(\theta)$ differs from the angle between $u_i(\theta_k)$ and $u_j(\theta_k)$ by at most $|\alpha_i(\theta) - \alpha_i(\theta_k)| + |\alpha_j(\theta) - \alpha_j(\theta_k)|$, which by construction of Θ is at most 2ν. ■

To complete the discussion, we show how to remove our assumption that $v_i|0$ is nonzero for all i. If $v_i|0$ is the zero vector, then $u_i(\theta)$ is not defined for only one value of θ: either $\gamma_i - \pi/2$ or $\gamma_i - \frac{3}{2}\pi$. This value of θ is a possible point of discontinuity of $f_\theta(v_1, \ldots, v_n)$. Moreover, below or above this value, $u_i(\theta)$ is independent of θ. For the index i, we only include in the set Θ the two values of θ just before and after this threshold value. The proof of Lemma 5.1 must be modified slightly. We first restrict our attention to each interval where f_θ is continuous. We then observe that the increase in the angle between $u_i(\theta)$ and $u_j(\theta)$ is even easier to handle if $v_i|0$ is the zero vector (or $v_j|0$ or both) since $u_i(\theta)$ is independent of θ in any such interval.

6 Generalizations

We can use the same technique as in Section 1 to approximate several other problems.

MAX RES CUT For the variation of MAX CUT in which pairs of vertices are forced to be on either the same side of the cut or on different sides of the cut, we have an $(\alpha - \epsilon)$-approximation algorithm.

MAX 2SAT For the maximum 2-satisfiability problem, we also have an $(\alpha - \epsilon)$-approximation algorithm.

MAX DICUT For the maximum directed cut problem, we have a $(\beta - \epsilon)$-approximation algorithm where $\beta > 0.79607$.

The extension to the MAX RES CUT problem is trivial. We merely need to add the following constraints to (P): $v_i \cdot v_j = 1$ for $(i,j) \in E^+$ and $v_i \cdot v_j = -1$ for $(i,j) \in E^-$, where E^+ (resp. E^-) corresponds to the pair of vertices forced to be on the same side (resp. different sides) of the cut. Using the randomized algorithm of Section 1 and setting $y_i = 1$ if $r \cdot v_i \geq 0$ and $y_i = -1$ otherwise gives a feasible solution to MAX RES CUT, assuming that a feasible solution exists. Indeed, it is easy to see that if $v_i \cdot v_j = 1$, then the algorithm will produce a solution such that $y_i y_j = 1$. If $v_i \cdot v_j = -1$ then the only case in which the algorithm produces a solution such that $y_i y_j \neq -1$ is when $v_i \cdot r = v_j \cdot r = 0$, an event that happens with probability 0. The analysis of the expected value of the cut is unchanged and, therefore, the resulting algorithm is a randomized $(\alpha - \epsilon)$-approximation algorithm.

In the next section, we show that the algorithm for MAX CUT can also be used to approximate a more general integer quadratic program, and that MAX 2SAT can be modelled as such. We also show how the algorithm for MAX 2SAT leads to a slightly improved algorithm for MAX SAT. In Section 6.2, we show how to approximate another general integer quadratic program, which can be used to model the MAX DICUT problem. Before describing these extensions, we should point out that nothing prevents the combination of these three generalizations.

6.1 MAX 2SAT

We consider the integer quadratic program

$$\begin{array}{lll} & \text{Maximize} & \displaystyle\sum_{i<j} [a_{ij}(1 - y_i y_j) + b_{ij}(1 + y_i y_j)] \\ (Q') & \text{subject to:} & y_i \in \{-1, 1\} \qquad \forall i \in V, \end{array}$$

where a_{ij} and b_{ij} are non-negative. The objective function of (Q') is thus a non-negative linear form in $1 \pm y_i y_j$. We relax (Q') to:

$$\begin{array}{lll} & \text{Maximize} & \displaystyle\sum_{i<j} [a_{ij}(1 - v_i \cdot v_j) + b_{ij}(1 + v_i \cdot v_j)] \\ (P') & \text{subject to:} & v_i \in S_n \qquad \forall i \in V. \end{array}$$

Let $E[V]$ be the expected value of the solution produced by the randomized algorithm. By the linearity of expectations,

$$E[V] = 2 \sum_{i<j} a_{ij} \Pr[\text{sgn}(v_i \cdot r) \neq \text{sgn}(v_j \cdot r)] + 2 \sum_{i<j} b_{ij} \Pr[\text{sgn}(v_i \cdot r) = \text{sgn}(v_j \cdot r)].$$

Using the analysis of the max cut algorithm, we note that $\Pr[\mathrm{sgn}(v_i \cdot r) = \mathrm{sgn}(v_j \cdot r)] = 1 - \frac{1}{\pi}\arccos(v_i \cdot v_j)$, and thus the approximation ratio for the more general program follows from Lemmas 2.4 and 2.7.

Hence we can show the following theorem, which implies that the algorithm is an $(\alpha - \epsilon)$-approximation algorithm for (Q').

Theorem 6.1

$$E[V] \geq \alpha \sum_{i<j} \left[a_{ij}(1 - v_i \cdot v_j) + b_{ij}(1 + v_i \cdot v_j)\right].$$

We now show that the maximum 2-satisfiability problem can be modelled by (Q'). An instance of the maximum satisfiability problem (MAX SAT) is defined by a collection $\mathcal{C}$ of boolean clauses, where each clause is a disjunction of literals drawn from a set of variables $\{x_1, x_2, \ldots, x_n\}$. A *literal* is either a variable x or its negation $\bar{x}$. The *length* $l(C_j)$ of a clause C_j is the number of distinct literals in the clause. In addition, for each clause $C_j \in \mathcal{C}$ there is an associated non-negative weight w_j. An optimal solution to a MAX SAT instance is an assignment of truth values to the variables $x_1, \ldots, x_n$ that maximizes the sum of the weight of the satisfied clauses. MAX 2SAT consists of MAX SAT instances in which each clause has length at most two. MAX 2SAT is NP-complete [13]; the best approximation algorithm known previously has a performance guarantee of $\frac{3}{4}$ and is due to Yannakakis [56] (see also Goemans and Williamson [15]). As with the MAX CUT problem, MAX 2SAT is known to be MAX SNP-hard [41]; thus there exists some constant $c < 1$ such that the existence of a c-approximation algorithm implies that $P = NP$ [2]. Haglin [20, 21] has shown that any ρ-approximation algorithm for MAX RES CUT can be translated into an ρ-approximation algorithm for MAX 2SAT, but we will show a direct algorithm here.

We will model MAX 2SAT using the integer quadratic program (Q'). We introduce a variable y_i in the quadratic program for each boolean variable x_i in the 2SAT instance; we also introduce an additional variable y_0. The value of y_0 will determine whether -1 or 1 will correspond to "true" in the MAX 2SAT instance. More precisely, x_i is true if $y_i = y_0$ and false otherwise. Given a boolean formula C, we define its value $v(C)$ to be 1 if the formula is true and 0 if the formula is false. Thus, $v(x_i) = \frac{1+y_0y_i}{2}$ and $v(\bar{x}_i) = 1 - v(x_i) = \frac{1-y_0y_i}{2}$. Observe that

$$\begin{aligned} v(x_i \vee x_j) &= 1 - v(\bar{x}_i \wedge \bar{x}_j) = 1 - v(\bar{x}_i)v(\bar{x}_j) = 1 - \frac{1 - y_0y_i}{2}\frac{1 - y_0y_j}{2} \\ &= \frac{1}{4}\left(3 + y_0y_i + y_0y_j - y_0^2y_iy_j\right) \\ &= \frac{1 + y_0y_i}{4} + \frac{1 + y_0y_j}{4} + \frac{1 - y_iy_j}{4}. \end{aligned}$$

The value of other clauses with 2 literals can be similarly expressed; for instance, if x_i is negated one only needs to replace y_i by $-y_i$. Therefore, the value $v(C)$ of any clause with at most two literals per clause can be expressed in the form required in (Q'). As a result, the MAX 2SAT problem can be modelled as

$$\begin{aligned} &\text{Maximize} && \sum_{C_j \in \mathcal{C}} w_j v(C_j) \\ (SAT) \qquad &\text{subject to:} && y_i \in \{-1, 1\} \qquad\qquad \forall i \in \{0, 1, \ldots, n\}, \end{aligned}$$

where the $v(C_j)$ are non-negative linear combinations of $1+y_iy_j$ and $1-y_iy_j$. The (SAT) program is in the same form as (Q'), so by the theorem above, our algorithm is an $(\alpha - \epsilon)$-approximation algorithm for the MAX 2SAT problem.

6.1.1 MAX SAT

The improved MAX 2SAT algorithm leads to a slightly improved approximation algorithm for MAX SAT. In [15], we developed several randomized $\frac{3}{4}$-approximation algorithms for MAX SAT. We considered the following linear programming relaxation of the MAX SAT problem, where I_j^+ denotes the set of non-negative variables in clause C_j and I_j^- is the set of negative variables in C_j:

$$\begin{aligned}
\text{Max} \quad & \sum_{C_j \in \mathcal{C}} w_j z_j \\
\text{subject to:} & \\
& \sum_{i \in I_j^+} y_i + \sum_{i \in I_j^-} (1 - y_i) \geq z_j && \forall C_j \in \mathcal{C} \\
& 0 \leq y_i \leq 1 && 1 \leq i \leq n \\
& 0 \leq z_j \leq 1 && \forall C_j \in \mathcal{C}.
\end{aligned}$$

By associating $y_i = 1$ with x_i set true, $y_i = 0$ with x_i set false, $z_j = 1$ with clause C_j satisfied, and $z_j = 0$ with clause C_j not satisfied, the program exactly corresponds to the MAX SAT problem. We showed that for any feasible solution (y, z), if x_i is set to be true with probability y_i, then the probability that clause j will be satisfied is at least $(1 - (1 - \frac{1}{k})^k) z_j$, for $k = l(C_j)$. We then considered choosing randomly between the following two algorithms: (1) set x_i true independently with probability y_i; (2) set x_i true independently with probability $\frac{1}{2}$. Given this combined algorithm, the probability that a length k clause is satisfied is at least $\frac{1}{2}(1-2^{-k}) + \frac{1}{2}(1-(1-\frac{1}{k})^k) z_j$. This expression can be shown to be at least $\frac{3}{4} z_j$ for all lengths k. Thus if an optimal solution to the linear program is used, the algorithm results in a $\frac{3}{4}$-approximation algorithm for MAX SAT, since the expected value of the algorithm is at least $\frac{3}{4} \sum_j w_j z_j$.

We formulate a slightly different relaxation of the MAX SAT problem. Let $u(C)$ denote a relaxed version of the expression v used in the previous section in which the products $y_i y_j$ are replaced by inner products of vectors $v_i \cdot v_j$. Thus $u(x_i) = \frac{1}{2}(1 + v_i \cdot v_0)$, $u(\bar{x}_i) = \frac{1}{2}(1 - v_i \cdot v_0)$, and $u(x_i \vee x_j) = \frac{1}{4}(1 + v_i \cdot v_0) + \frac{1}{4}(1 + v_j \cdot v_0) + \frac{1}{4}(1 - v_i \cdot v_j)$. We then consider the following relaxation,

$$\begin{aligned}
\text{Max} \quad & \sum_{C_j \in \mathcal{C}} w_j z_j \\
\text{subject to:} & \\
& \sum_{i \in I_j^+} u(x_i) + \sum_{i \in I_j^-} u(\bar{x}_i) \geq z_j && \forall C_j \in \mathcal{C} \\
& u(C_j) \geq z_j && \forall C_j \in \mathcal{C}, l(C_j) = 2 \\
& v_i \cdot v_i = 1 && 0 \leq i \leq n \\
& 0 \leq z_j \leq 1 && \forall C_j \in \mathcal{C}.
\end{aligned}$$

Thus if we set x_i to be true with probability $u(x_i)$ for the optimal solution to the corresponding semidefinite program, then by the arguments of [15], we satisfy a clause C_j of length k with probability at least $(1 - (1 - \frac{1}{k})^k) z_j$.

To obtain the improved bound, we consider the three algorithms: (1) set x_i true independently with probability $\frac{1}{2}$; (2) set x_i true independently with probability $u(x_i)$ (given the optimal solution to the program); (3) pick a random unit vector r and set x_i true iff $\text{sgn}(v_i \cdot r) = \text{sgn}(v_0 \cdot r)$. Suppose we use algorithm i with probability p_i, where $p_1 + p_2 + p_3 = 1$. From the previous

section, for algorithm (3) the probability that a clause C_j of length 1 or 2 is satisfied is at least $.87856u(C_j) \geq .87856z_j$. Thus the expected value of the solution is at least

$$\sum_{j:l(C_j)=1} w_j(.5p_1 + (p_2 + .87856p_3)z_j) + \sum_{j:l(C_j)=2} w_j(.75p_1 + (.75p_2 + .87856p_3)z_j)$$
$$+ \sum_{j:l(C_j)\geq 3} w_j((1-2^{-l(C_j)})p_1 + (1-(1-\frac{1}{l(C_j)})^{l(C_j)})p_2 z_j).$$

If we set $p_1 = p_2 = .4785$ and $p_3 = .0430$, then the expected value is at least $.7554\sum_j w_j z_j$, yielding a .7554-approximation algorithm. To see this, we check the value of the expression for lengths 1 through 4, and notice that $\min_k(1-(1-\frac{1}{k})^k) \geq 1-\frac{1}{e}$ and $.4785((1-2^{-5})+1-\frac{1}{e}) \geq .7554$.

We can obtain even a slightly better approximation algorithm for the MAX SAT problem. The bottleneck in the analysis above is that algorithm (3) contributes no expected weight for clauses of length 3 or greater. For a given clause C_j of length 3 or more, let P_j be a set of length 2 clauses formed by taking the literals of C_j two at a time; thus P_j will contain $\binom{l(C_j)}{2}$ clauses. If at least one of the literals in C_j is set true, then at least $l(C_j)-1$ of the clauses in P_j will be satisfied. Thus the following program is a relaxation of the MAX SAT problem:

$$\text{Max} \quad \sum_{C_j \in \mathcal{C}} w_j z_j$$

subject to:

$$\sum_{i \in I_j^+} u(x_i) + \sum_{i \in I_j^-} u(\bar{x}_i) \geq z_j \qquad \forall C_j \in \mathcal{C}$$
$$u(C_j) \geq z_j \qquad \forall C_j \in \mathcal{C}, l(C_j) = 2$$
$$\frac{1}{l(C_j)-1} \sum_{C \in P_j} u(C) \geq z_j \qquad \forall C_j \in \mathcal{C}, l(C_j) \geq 3$$
$$v_i \cdot v_i = 1 \qquad 0 \leq i \leq n$$
$$0 \leq z_j \leq 1 \qquad \forall C_j \in \mathcal{C}.$$

Algorithm (3) has expected value of $.87856u(C)$ for each $C \in P_j$ for any j, so that its expected value for any clause of length 3 or more becomes at least

$$.87856 \cdot \frac{1}{\binom{l(C_j)}{2}} \sum_{C \in P_j} u(C) = .87856 \cdot \frac{2}{l(C_j)} \frac{1}{l(C_j)-1} \sum_{C \in P_j} u(C)$$
$$\geq .87856 \cdot \frac{2}{l(C_j)} z_j,$$

so that the overall expectation of the algorithm will be at least

$$\sum_{j:l(C_j)=1} w_j(.5p_1 + (p_2 + .87856p_3)z_j) + \sum_{j:l(C_j)=2} w_j(.75p_1 + (.75p_2 + .87856p_3)z_j)$$
$$+ \sum_{j:l(C_j)\geq 3} w_j((1-2^{-l(C_j)})p_1 + ((1-(1-\frac{1}{l(C_j)})^{l(C_j)})p_2 + .87856 \cdot \frac{2}{l(C_j)}p_3)z_j).$$

By setting $p_1 = p_2 = .467$ and $p_3 = .066$, we obtain a .7584-approximation algorithm, which can be verified by checking the expression for lengths 1 through 6, and noticing that $.467((1-\frac{1}{e}) + (1-2^{-7})) \geq .7584$. Other small improvements are possible by tightening the analysis.

6.2 MAX DICUT

We consider the integer quadratic program

$$\begin{array}{lll} & \text{Max} & \sum_{i,j,k}[c_{ijk}(1 - y_iy_j - y_iy_k + y_jy_k) + d_{ijk}(1 + y_iy_j + y_iy_k + y_jy_k)] \\ (Q'') & \text{subject to:} & y_i \in \{-1, 1\} \qquad \forall i \in V, \end{array}$$

where c_{ijk} and d_{ijk} are non-negative. Observe that $1 - y_iy_j - y_iy_k + y_jy_k$ can also be written as $(1 - y_iy_j)(1 - y_iy_k)$ (or as $(1 - y_iy_j)(1 + y_jy_k)$), and, thus, the objective function of (Q'') can be interpreted as a non-negative restricted quadratic form in $1 \pm y_iy_j$. Moreover, $1 - y_iy_j - y_iy_k + y_jy_k$ is equal to 4 if $y_i = -y_j = -y_k$ and 0 otherwise, while $1 + y_iy_j + y_iy_k + y_jy_k$ is 4 if $y_i = y_j = y_k$ and is 0 otherwise. We relax (Q'') to:

$$\begin{array}{lll} & \text{Max} & \sum_{i,j,k}[c_{ijk}(1 - v_i \cdot v_j - v_i \cdot v_k + v_j \cdot v_k) + d_{ijk}(1 + v_i \cdot v_j + v_i \cdot v_k + v_j \cdot v_k)] \\ (P'') & \text{subject to:} & v_i \in S_n \qquad \forall i \in V. \end{array}$$

We approximate (Q'') by using exactly the same algorithm as before. The analysis is somewhat more complicated, however, and the performance guarantee β is slightly weaker, namely

$$\beta = \min_{0 \le \theta < \arccos(-1/3)} \frac{2}{\pi} \frac{2\pi - 3\theta}{1 + 3\cos\theta} > 0.79607.$$

Given a vector r drawn uniformly from the unit sphere S_n, we know by the linearity of expectations that the expected value $E[U]$ of the solution output is

$$4 \sum_{i,j,k} [c_{ijk} \cdot \Pr[\text{sgn}(v_i \cdot r) \neq \text{sgn}(v_j \cdot r) = \text{sgn}(v_k \cdot r)] + d_{ijk} \cdot \Pr[\text{sgn}(v_i \cdot r) = \text{sgn}(v_j \cdot r) = \text{sgn}(v_k \cdot r)]] .$$

Consider any term in the sum, say $d_{ijk} \cdot \Pr[\text{sgn}(v_i \cdot r) = \text{sgn}(v_j \cdot r) = \text{sgn}(v_k \cdot r)]$. The c_{ijk} terms can be dealt with similarly by simply replacing v_i by $-v_i$. The performance guarantee follows from the proof of the following two lemmas.

Lemma 6.2

$$\Pr[sgn(v_i \cdot r) = sgn(v_j \cdot r) = sgn(v_k \cdot r)] = 1 - \frac{1}{2\pi}(\arccos(v_i \cdot v_j) + \arccos(v_i \cdot v_k) + \arccos(v_j \cdot v_k)).$$

Lemma 6.3 *For any* $v_i, v_j, v_k \in S_n$,

$$1 - \frac{1}{2\pi}(\arccos(v_i \cdot v_j) + \arccos(v_i \cdot v_k) + \arccos(v_j \cdot v_k)) \ge \frac{\beta}{4}[1 + v_i \cdot v_j + v_i \cdot v_k + v_j \cdot v_k],$$

where $\beta = \min_{0 \le \theta < \arccos(-1/3)} \frac{2}{\pi} \frac{2\pi - 3\theta}{1 + 3\cos\theta} > 0.79607$.

Proof of Lemma 6.2: The desired probability can be seen to be equal to twice the area of the spherical triangle polar to the spherical triangle defined by v_i, v_j and v_k. Stated this way, the result is a corollary to Girard's formula (1629 [14], see [51]) expressing the area of a spherical triangle with angles α, β and γ as its excess $\alpha + \beta + \gamma - \pi$.

We present a proof of the lemma from first principles. In fact, our proof parallels Euler's proof (1781 [10], see [51]) of Girard's formula. We define the following events:

$$\begin{array}{ll} A: & \operatorname{sgn}(v_i \cdot r) = \operatorname{sgn}(v_j \cdot r) = \operatorname{sgn}(v_k \cdot r) \\ B_i: & \operatorname{sgn}(v_i \cdot r) \neq \operatorname{sgn}(v_j \cdot r) = \operatorname{sgn}(v_k \cdot r) \\ C_i: & \operatorname{sgn}(v_j \cdot r) = \operatorname{sgn}(v_k \cdot r) \\ C_j: & \operatorname{sgn}(v_i \cdot r) = \operatorname{sgn}(v_k \cdot r) \\ C_k: & \operatorname{sgn}(v_i \cdot r) = \operatorname{sgn}(v_j \cdot r). \end{array}$$

Note that $B_i = C_i - A$. We define B_j and B_k similarly, so that $B_j = C_j - A$ and $B_k = C_k - A$. Clearly,

$$\Pr[A] + \Pr[B_i] + \Pr[B_j] + \Pr[B_k] = 1. \tag{1}$$

Also, $\Pr[C_i] = \Pr[A] + \Pr[B_i]$ and similarly for j and k. Adding up these equalities and subtracting (1), we obtain

$$\Pr[C_i] + \Pr[C_j] + \Pr[C_k] = 1 + 2\Pr[A]. \tag{2}$$

By Lemma 2.2, $\Pr[C_i] = 1 - \frac{1}{\pi}\arccos(v_j \cdot v_k)$ and similarly for j and k. Together with (2), we derive

$$\Pr[A] = 1 - \frac{1}{2\pi}(\arccos(v_i \cdot v_j) + \arccos(v_i \cdot v_k) + \arccos(v_j \cdot v_k)),$$

proving the lemma. ■

Proof of Lemma 6.3: Let $a = \arccos(v_i \cdot v_j)$, $b = \arccos(v_i \cdot v_k)$ and $c = \arccos(v_j \cdot v_k)$. From the theory of spherical triangles, it follows that the possible values for (a, b, c) over all possible vectors v_i, v_j and v_k define the set

$$S = \{(a,b,c) : 0 \leq a \leq \pi, 0 \leq b \leq \pi, 0 \leq c \leq \pi, c \leq a+b, b \leq a+c, a \leq b+c, a+b+c \leq 2\pi\}.$$

(see [6, Corollary 18.6.12.3]). The claim can thus be restated as $1 - \frac{1}{2\pi}(a+b+c) \geq \frac{\beta}{4}(1+\cos(a)+\cos(b)+\cos(c))$ for all $(a,b,c) \in S$.

Let (a,b,c) minimize $h(a,b,c) = 1 - \frac{1}{2\pi}(a+b+c) - \frac{\beta}{4}(1+\cos(a)+\cos(b)+\cos(c))$ over all $(a,b,c) \in S$. We consider several cases:

1. $a + b + c = 2\pi$. We have $1 - \frac{1}{2\pi}(a+b+c) = 0$. On the other hand,

$$\begin{aligned} & 1 + \cos(a) + \cos(b) + \cos(c) \\ &= 1 + \cos(a) + \cos(b) + \cos(a+b) \\ &= 1 + \cos(a+b) + 2\cos(\frac{a+b}{2})\cos(\frac{a-b}{2}) \\ &= 2\cos^2(\frac{a+b}{2}) + 2\cos(\frac{a+b}{2})\cos(\frac{a-b}{2}) \\ &= 2\cos(\frac{a+b}{2})\left[\cos(\frac{a+b}{2}) + \cos(\frac{a-b}{2})\right]. \end{aligned} \tag{3}$$

We now derive that $h(a,b,c) \geq -\frac{\beta}{2}\cos(\frac{a+b}{2})[\cos(\frac{a+b}{2}) + \cos(\frac{a-b}{2})] \geq 0$, the last inequality following from the fact that $\frac{\pi}{2} \leq \frac{a+b}{2} \leq \pi - |\frac{a-b}{2}|$ and thus $\cos(\frac{a+b}{2}) \leq 0$ and $[\cos(\frac{a+b}{2}) + \cos(\frac{a-b}{2})] \geq 0$.

2. $a = b + c$ or $b = a + c$ or $c = a + b$. By symmetry, assume that $c = a + b$. Observe that $1 - \frac{1}{2\pi}(a + b + c) = 1 - \frac{a+b}{\pi}$. On the other hand, by (3) we have that

$$\begin{aligned} & 1 + \cos(a) + \cos(b) + \cos(c) \\ = \quad & 2\cos(\frac{a+b}{2})(\cos(\frac{a-b}{2}) + \cos(\frac{a+b}{2})) \\ \leq \quad & 2\cos(\frac{a+b}{2})(1 + \cos(\frac{a+b}{2})). \end{aligned}$$

   Letting $x = \frac{a+b}{2}$, we observe that the claim is equivalent to $1 - \frac{2x}{\pi} \geq \frac{\beta}{2}\cos(x)(1 + \cos(x))$ for any $0 \leq x \leq \pi/2$. One can in fact verify that $1 - \frac{2x}{\pi} \geq \frac{0.81989}{2}\cos(x)(1 + \cos(x))$ for any $0 \leq x \leq \pi/2$, implying the claim.

3. $a = 0$ or $b = 0$ or $c = 0$. Without loss of generality, let $a = 0$. The definition of S implies that $b = c$, and thus $b = a + c$. This case therefore reduces to the previous one.

4. $a = \pi$ or $b = \pi$ or $c = \pi$. Assume $a = \pi$. This implies that $b + c = \pi$ and, thus, $a + b + c = 2\pi$. We have thus reduced the problem to case 1.

5. In the last case, (a, b, c) belongs to the interior of S. This implies that the gradient of h must vanish and the hessian of h must be positive semidefinite at (a, b, c). In other words, $\sin a = \sin b = \sin c = \frac{2}{\beta\pi}$, and $\cos a \geq 0, \cos b \geq 0$ and $\cos c \geq 0$. From this, we derive that $a = b = c$. But $h(a, a, a) = 1 - \frac{3a}{2\pi} - \frac{\beta}{4}(1 + 3\cos(a))$. The lemma now follows from the fact that $a \leq \frac{2\pi}{3}$, the definition of β and the fact that $1 + 3\cos a \leq 0$ for $a \geq \arccos\frac{-1}{3}$.

■

Suppose we are given a directed graph $G = (V, A)$ and weights w_{ij} on each directed arc $(i, j) \in A$, where i is the *tail* of the arc and j is the *head*. The maximum directed cut problem is that of finding the set of vertices S that maximizes the weight of the edges with their tails in S and their heads in $\bar{S}$. The problem is NP-hard via a straightforward reduction from MAX CUT. The previous best known approximation algorithm for MAX DICUT has a performance guarantee of $\frac{1}{4}$ [41].

We can model the MAX DICUT problem using the program (Q''). We introduce a variable y_i for each $i \in V$, and, as with the MAX 2SAT program, we introduce a variable y_0 that will denote the S side of the cut. Thus $i \in S$ iff $y_i = y_0$. Then arc (i, j) contributes weight $\frac{1}{4}w_{ij}(1 + y_i y_0)(1 - y_j y_0)$ to the cut. Summing over all arcs $(i, j) \in A$ gives a program of the same form as (Q''). Hence we obtain a .79607-approximation algorithm for MAX DICUT. Finally, if the directed graph has weighted indegree of every vertex equal to weighted outdegree, the program (Q'') reduces to one of the form (Q'), and therefore our approximation algorithm has a performance guarantee of $(\alpha - \epsilon)$.

7 Concluding Remarks

Our motivation for studying semidefinite programming relaxations came from a realization that the standard tool of using linear programming relaxations for approximation algorithms had limits which might not be easily surpassed (see the conclusion of Goemans and Williamson [15]). In fact, a classical linear programming relaxation for the maximum cut problem can be shown to be arbitrarily close to twice the value of the maximum cut in the worst case. Given the work of Lovász and Schrijver [33, 34], which showed that tighter and tighter relaxations could

be obtained through semidefinite programming, it seemed worthwhile to investigate the power of such relaxations from a worst-case perspective. The results of this paper constitute a first step in this direction. As we mentioned in the introduction, further steps have already been made, with improved results for MAX 2SAT and MAX DICUT by Feige and Goemans, and for coloring by Karger, Motwani, and Sudan. We think that the continued investigation of these methods is promising.

While this paper leaves many open questions, we think there are two especially interesting problems. The first question is whether a .878-approximation algorithm for MAX CUT can be obtained without explicitly solving the semidefinite program. For example, the first 2-approximation algorithms for weighted vertex cover involved solving a linear program [23], but later Bar-Yehuda and Even [3] devised a primal-dual algorithm in which linear programming was used only in the analysis of the algorithm. Perhaps a semidefinite analog is possible for MAX CUT. The second question is whether adding additional constraints to the semidefinite program leads to a better worst-case bound. There is some reason to think this might be true. Linear constraints are known for which the program would find an optimal solution on planar graphs, whereas there is a gap of $\frac{32}{25+5\sqrt{5}}$ for the current semidefinite program.

One consequence of this paper is that the situation with several MAX SNP problems is no longer clear-cut. When the best-known approximation results for MAX CUT and MAX SAT had such long-standing and well-defined bounds as $\frac{1}{2}$ and $\frac{3}{4}$, it was tempting to believe that perhaps no further work could be done in approximating these problems, and that it was only a matter of time before matching hardness results would be found. The improved results in this paper should rescue algorithm designers from such fatalism. Although MAX SNP problems cannot be approximated arbitrarily closely, there still is work to do in designing improved approximation algorithms.

Acknowledgements

Since the appearence of an extended abstract of this paper, we have benefited from discussions with a very large number of colleagues. In particular, we would like to thank Don Coppersmith and David Shmoys for pointing out to us that our MAX 2SAT algorithm could lead to an improved MAX SAT algorithm, Jon Kleinberg for bringing reference [31] to our attention, Joseph Cheriyan for pointing out a mistake in the derandomization section, Bob Vanderbei for kindly giving us his semidefinite code, Francisco Barahona and Michael Jünger for providing problem instances, Joel Spencer for motivating Theorem 2.5, Farid Alizadeh, Gabor Pataki and Rob Freund for results on semidefinite programming, and Shang-Hua Teng for bringing reference [28] to our attention.

References

[1] F. Alizadeh. Interior point methods in semidefinite programming with applications to combinatorial optimization. To appear in *SIAM Journal on Optimization*. A preliminary version appeared in the *Proceedings of the 2nd Mathematical Programming Society Conference on Integer Programming and Combinatorial Optimization*, 1992.

[2] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and hardness of approximation problems. In *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science*, pages 14–23, 1992.

[3] R. Bar-Yehuda and S. Even. A linear time approximation algorithm for the weighted vertex cover problem. *Journal of Algorithms*, 2:198–203, 1981.

[4] F. Barahona, M. Grötschel, M. Jünger, and G. Reinelt. An application of combinatorial optimization to statistical physics and circuit layout design. *Operations Research*, 36:493–513, 1988.

[5] A. I. Barvinok. Problems of distance geometry and convex properties of quadratic maps. Submitted to *Discrete and Computational Geometry*. A preliminary version appeared as "Topology and convex geometry of quadratic equations", Mathematical Sciences Institute, Cornell University, Technical Report 94-2, 1994.

[6] M. Berger. *Geometry II*. Springer-Verlag, Berlin, 1987.

[7] C. Delorme and S. Poljak. Combinatorial properties and the complexity of a max-cut approximation. *European Journal of Combinatorics*, 14:313–333, 1993.

[8] C. Delorme and S. Poljak. Laplacian eigenvalues and the maximum cut problem. *Mathematical Programming*, 62:557–574, 1993.

[9] C. Delorme and S. Poljak. The performance of an eigenvalue bound on the max-cut problem in some classes of graphs. *Discrete Mathematics*, 111:145–156, 1993.

[10] L. Euler. De mensura angulorum solidorum. Petersburg, 1781.

[11] U. Feige and M. X. Goemans. Personal communication, 1994.

[12] U. Feige and L. Lovász. Two-prover one-round proof systems: Their power and their problems. In *Proceedings of the 24th Annual ACM Symposium on the Theory of Computing*, pages 733–744, 1992.

[13] M. Garey, D. Johnson, and L. Stockmeyer. Some simplified NP-complete graph problems. *Theoretical Computer Science*, 1:237–267, 1976.

[14] A. Girard. De la mesure de la superficie des triangles et polygones sphériques. Appendix to "Invention nouvelle en l'algèbre", Amsterdam, 1629.

[15] M. X. Goemans and D. P. Williamson. New 3/4-approximation algorithms for MAX SAT. To appear in *SIAM J. Disc. Math.* Earlier version appeared in the *Proceedings of the 3rd Mathematical Programming Society Conference on Integer Programming and Combinatorial Optimization*, pages 313–321, 1993.

[16] G. H. Golub and C. F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, MD, 1983.

[17] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1:169–197, 1981.

[18] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, Berlin, 1988.

[19] F. Hadlock. Finding a maximum cut of a planar graph in polynomial time. *SIAM Journal on Computing*, 4:221–225, 1975.

[20] D. J. Haglin. Approximating maximum 2-CNF satisfiability. *Parallel Processing Letters*, 2:181–187, 1992.

[21] D. J. Haglin. Personal communication, 1994.

[22] D. J. Haglin and S. M. Venkatesan. Approximation and intractability results for the maximum cut problem and its variants. *IEEE Transactions on Computers*, 40:110–113, 1991.

[23] D. S. Hochbaum. Approximation algorithms for the set covering and vertex cover problems. *SIAM Journal on Computing*, 11:555–556, 1982.

[24] S. Homer and M. Peinado. A highly parallel implementation of the Goemans/Williamson algorithm to approximate MaxCut. Manuscript, 1994.

[25] M. Jünger. Personal communication, 1994.

[26] D. Karger, R. Motwani, and M. Sudan. Approximate graph coloring by semidefinite programming. In preparation, 1994.

[27] R. M. Karp. Reducibility among combinatorial problems. In R. Miller and J. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, New York, NY, 1972.

[28] D. E. Knuth. *Seminumerical Algorithms*, volume 2 of *The Art of Computer Programming*. Addison-Wesley, Reading, MA, second edition, 1981.

[29] P. Lancaster and M. Tismenetsky. *The Theory of Matrices*. Academic Press, Orlando, FL, 1985.

[30] L. Lovász. On the Shannon capacity of a graph. *IEEE Transactions on Information Theory*, IT-25:1–7, 1979.

[31] L. Lovász. Self-dual polytopes and the chromatic number of distance graphs on the sphere. *Acta Scientiarum Mathematicarum*, 45:317–323, 1983.

[32] L. Lovász. Combinatorial optimization: Some problems and trends. DIMACS Technical Report 92-53, 1992.

[33] L. Lovász and A. Schrijver. Matrix cones, projection representations, and stable set polyhedra. In *Polyhedral Combinatorics*, volume 1 of *DIMACS series in Discrete Mathematics and Theoretical Computer Science*, pages 1–17. American Mathematical Society, 1989.

[34] L. Lovász and A. Schrijver. Cones of matrices and setfunctions, and 0-1 optimization. *SIAM Journal on Optimization*, 1:166–190, 1990.

[35] B. Mohar and S. Poljak. Eigenvalues in combinatorial optimization. Technical report, University of Ljubljana, 1992.

[36] Y. Nesterov and A. Nemirovskii. *Self-Concordant Functions and Polynomial Time Methods in Convex Programming*. Central Economic and Mathematical Institute, USSR Academy of Science, Moscow, 1989.

[37] Y. Nesterov and A. Nemirovskii. *Interior Point Polynomial Methods in Convex Programming*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1994.

[38] G. I. Orlova and Y. G. Dorfman. Finding the maximal cut in a graph. *Engineering Cybernetics*, pages 502–506, 1972.

[39] M. L. Overton and R. S. Womersley. On the sum of the largest eigenvalues of a symmetric matrix. *SIAM Journal on Matrix Analysis and Applications*, 13:41–45, 1992.

[40] M. L. Overton and R. S. Womersley. Optimality conditions and duality theory for minimizing sums of the largest eigenvalues of symmetric matrices. *Mathematical Programming*, 62:321–357, 1993.

[41] C. H. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *Journal of Computer and System Sciences*, 43:425–440, 1991.

[42] G. Pataki. Algorithms for cone-optimization problems and semi-definite programming. Manuscript, 1993.

[43] G. Pataki. On the multiplicity of optimal eigenvalues. Manuscript, 1994.

[44] S. Poljak and F. Rendl. Solving the max-cut problem using eigenvalues. Report 91735-OR, Institute für Diskrete Mathematik, Universität Bonn, 1991.

[45] S. Poljak and F. Rendl. Nonpolyhedral relaxations of graph-bisection problems. DIMACS Technical Report 92-55, 1992.

[46] S. Poljak and F. Rendl. Computational experiments with node and edge relaxations of the max-cut problem. Report 266, Technische Universität Graz, 1993.

[47] S. Poljak and D. Turzík. A polynomial algorithm for constructing a large bipartite subgraph, with an application to a satisfiability problem. *Canadian Journal of Mathematics*, 34:519–524, 1982.

[48] S. Poljak and Z. Tuza. The max-cut problem — a survey. In L. Lovász and P. Seymour, editors, *Special Year on Combinatorial Optimization*, DIMACS series in Discrete Mathematics and Theoretical Computer Science. American Mathematical Society, 1993. To be published.

[49] G. Reinelt. TSPLIB – A traveling salesman problem library. *ORSA Journal on Computing*, 3:376–384, 1991.

[50] F. Rendl, R. Vanderbei, and H. Wolkowicz. Interior point methods for max-min eigenvalue problems. Report 264, Technische Universität Graz, 1993.

[51] B. Rosenfeld. *A history of non-Euclidean geometry.* Springer-Verlag, Berlin, 1988.

[52] S. Sahni and T. Gonzalez. P-complete approximation problems. *Journal of the ACM*, 23:555–565, 1976.

[53] P. M. Vaidya. A new algorithm for minimizing convex functions over convex sets. In *Proceedings of the 30th Annual Symposium on Foundations of Computer Science*, pages 338–343, 1989.

[54] P. M. Vitányi. How well can a graph be n-colored? *Discrete Mathematics*, 34:69–80, 1981.

[55] H. Wolkowicz. Some applications of optimization in matrix theory. *Linear Algebra and Its Applications*, 40:101–118, 1981.

[56] M. Yannakakis. On the approximation of maximum satisfiability. In *Proceedings of the 3rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1–9, 1992.