

# MAXIMUM CUTS AND FRACTIONAL CUT COVERS: A COMPUTATIONAL STUDY OF A RANDOMIZED SEMIDEFINITE PROGRAMMING APPROACH

NATHAN BENEDETTO PROENÇA<sup>1†</sup>, MARCEL K. CARLI SILVA<sup>2\*</sup>, CRISTIANE M. SATO<sup>3\*</sup>,  
AND LEVENT TUNÇEL<sup>1¶</sup>

ABSTRACT. We present experimental work on a primal-dual framework simultaneously approximating maximum cut and weighted fractional cut-covering instances. In this primal-dual framework, we solve a semidefinite programming (SDP) relaxation to either the maximum cut problem or to the weighted fractional cut-covering problem, and then independently sample a collection of cuts via the random-hyperplane technique. We then simultaneously certify the approximate optimality of a cut and a fractional cut cover. We present several implementations which reliably achieve the celebrated Goemans and Williamson approximation ratio of  $\alpha_{\text{GW}} \approx 0.878$  for both optimization problems simultaneously, after  $\lceil 128 \ln m \rceil$  samples, a number significantly smaller than the best theoretical bounds.

This is the first experimental work approximating the weighted fractional cut-covering problem, and we deliver robust and repeatable results despite the use of randomized algorithms and floating-point arithmetic. Careful pre-processing of instances and post-processing of numeric results allow for good empirical outcomes with both first-order and second-order SDP solvers. Nearly optimal SDP solutions are suitably perturbed to ensure better probabilistic and numerical behavior. Our experiments deviate from theory by using a linear programming (LP) solver to compute fractional cut covers. For most instances studied, LP solving produces certifiably better results than the theoretical algorithm after  $\lceil 128 \ln m \rceil$  samples. All our experiments strictly follow a unified pipeline which explicitly documents all parameters used in each run.

## 1. INTRODUCTION

Let  $G = (V, E)$  be a (simple) graph with nonnegative edge weights  $w \in \mathbb{R}_+^E$ . The *cut* with *shore*  $S \subseteq V$  is the set  $\delta(S)$  of edges with exactly one endpoint in  $S$ . The weight of the cut  $\delta(S)$  is  $w^\top \mathbb{1}_{\delta(S)}$ , where  $\mathbb{1}_F \in \{0, 1\}^E$  denotes the *incidence vector* of  $F$  for every  $F \subseteq E$ . The *maximum cut problem* (or *maxcut problem*) for  $(G, w)$  is to find a cut that achieves the maximum in

$$(1) \quad \text{mc}(G, w) := \max\{w^\top \mathbb{1}_{\delta(S)} : S \subseteq V\}.$$

The maxcut problem is a cornerstone of combinatorial optimization, having attracted continuous research interest from multiple viewpoints for over six decades. Its influence spans a remarkably diverse range of fields, reaching deep into graph theory, computational complexity, statistical physics, circuit design, and quantum computing, among others. One of the earliest results, proved by Erdős [12] in 1965, is that every graph has a cut containing at least half of its edges, from which a simple 2-approximation algorithm follows (see also [11, 38]). The maxcut problem has exceptionally broad applicability, with its relevance growing as new technologies and connections emerge, including

---

<sup>1</sup>DEPARTMENT OF COMBINATORICS AND OPTIMIZATION, UNIVERSITY OF WATERLOO

<sup>2</sup>INSTITUTE OF MATHEMATICS, STATISTICS, AND COMPUTER SCIENCE, UNIVERSITY OF SÃO PAULO

<sup>3</sup>CENTER FOR MATHEMATICS, COMPUTING AND COGNITION, FEDERAL UNIVERSITY OF THE ABC REGION

<sup>¶</sup>Research of this author was supported in part by a Discovery Grant from the Natural Sciences and Engineering Research Council (NSERC) of Canada.

\*This work was partially supported by CNPq (408180/2023-4) and FAPESP (2023/03167-5).

<sup>†</sup>Corresponding Author: Nathan Benedetto Proença. Affiliation: University of Waterloo. E-mail: n2benede@uwaterloo.ca.

Ising spin glass models [4], VLSI circuit layout [1], and quantum computing [2, 15]. Moreover, it holds a prominent place in computational complexity theory [23, 25, 26]. Many computational studies have been conducted using a variety of approaches [3, 8, 10, 19, 27, 31, 36, 37, 41], including heuristics, polyhedral methods, spectral techniques, and semidefinite programming techniques.

Goemans and Williamson's approximation algorithm [16] for the maxcut problem is one of the most celebrated applications of semidefinite programming. They obtained a randomized  $\alpha_{\text{GW}}$ -approximation algorithm, where  $\alpha_{\text{GW}} := \min\{\frac{2}{\pi} \frac{\theta}{1-\cos\theta} : 0 < \theta \leq \pi\} \approx 0.87856$ . The heart of the algorithm uses an optimal solution of a semidefinite programming relaxation in a randomized rounding procedure that samples a cut. The relaxation can be written as follows. Let  $\mathbb{S}^V$  denote the set of  $V \times V$  symmetric matrices and let  $\mathbb{S}_+^V \subseteq \mathbb{S}^V$  denote the set of positive semidefinite matrices. For every  $w \in \mathbb{R}^E$ , we denote by  $\mathcal{L}_G(w) := \sum_{ij \in E} w_{ij} (e_i - e_j)(e_i - e_j)^\top \in \mathbb{S}^V$  the *weighted Laplacian of  $G$* , where the elements of  $\{e_i : i \in V\}$  are the canonical basis vectors of  $\mathbb{R}^V$ . The map  $\text{Diag} : \mathbb{R}^V \rightarrow \mathbb{S}^V$  builds a diagonal matrix from a vector, while  $\text{diag} : \mathbb{S}^V \rightarrow \mathbb{R}^V$  extracts the diagonal of a matrix;  $\mathbb{1}$  denotes the vector of all-ones. For every  $w \in \mathbb{R}_+^E$ , set

$$(2a) \quad \eta(G, w) := \max\{\langle \frac{1}{4}\mathcal{L}_G(w), Y \rangle : Y \in \mathbb{S}_+^V, \text{diag}(Y) = \mathbb{1}\}$$

$$(2b) \quad = \min\{\mathbb{1}^\top x : x \in \mathbb{R}^V, \text{Diag}(x) \succeq \frac{1}{4}\mathcal{L}_G(w)\}.$$

(We denote by  $\langle A, B \rangle := \text{Tr}(AB)$  the *trace inner product* of  $A, B \in \mathbb{S}^V$  and by  $A \succeq B$  the fact that  $A - B \in \mathbb{S}_+^V$ .) That is,  $\eta(G, w)$  is the optimal value of the semidefinite program (SDP) in (2a), which can be approximated to arbitrary precision in polynomial time; note that (2a) and (2b) have Slater points. Goemans and Williamson's randomized hyperplane technique samples a cut with expected weight at least  $\alpha_{\text{GW}}\eta(G, w)$ , which implies that

$$(3) \quad \alpha_{\text{GW}}\eta(G, w) \leq \text{mc}(G, w) \leq \eta(G, w).$$

We are interested in a deeply related problem: the weighted fractional cut-covering problem. Let  $\mathcal{P}(V)$  denote the power set of  $V$ . For any  $z \in \mathbb{R}_+^E$ , the *fractional cut-covering problem for  $(G, z)$*  is to find an optimal solution to

$$(4) \quad \text{fcc}(G, z) := \min\left\{\mathbb{1}^\top y : y \in \mathbb{R}_+^{\mathcal{P}(V)}, \sum_{S \subseteq V} y_S \mathbb{1}_{\delta(S)} \geq z\right\}.$$

A *fractional cut cover for  $(G, z)$*  is a feasible solution to (4). Every fractional cut cover provides a lower bound on the maximum cut value of every  $w \in \mathbb{R}_+^E$ :

$$(5) \quad w^\top z \leq w^\top \left( \sum_{S \subseteq V} y_S \mathbb{1}_{\delta(S)} \right) = \sum_{S \subseteq V} y_S \cdot w^\top \mathbb{1}_{\delta(S)} \leq \left( \sum_{S \subseteq V} y_S \right) \text{mc}(G, w) = (\mathbb{1}^\top y) \text{mc}(G, w).$$

The relationship between the maximum cut value for different weights  $w \in \mathbb{R}_+^E$  and fractional cut covers for  $z \in \mathbb{R}_+^E$  can be precisely captured by antiblocker duality [13, 14]. In particular, linear programming (LP) strong duality implies that

$$(6) \quad \text{fcc}(G, z) = \max\{z^\top w : w \in \mathbb{R}_+^E, \text{mc}(G, w) \leq 1\}.$$

Dually, we have that

$$(7) \quad \text{mc}(G, w) = \max\{w^\top z : z \in \mathbb{R}_+^E, \text{fcc}(G, z) \leq 1\}.$$

Indeed, using that  $\text{fcc}(G, \mathbb{1}_{\delta(S)}) \leq 1$  for every  $S \subseteq V$  and (5),

$$\max\{w^\top \mathbb{1}_{\delta(S)} : S \subseteq V\} \leq \max\{w^\top z : z \in \mathbb{R}_+^E, \text{fcc}(G, z) \leq 1\} \leq \text{mc}(G, w).$$

Antiblocker duality is yet another manifestation of the same duality that underlies LP and SDP strong duality. When solving an LP or SDP, we typically obtain an optimal dual solution either as a byproduct or as a crucial part of solving the primal problem. Therefore, it should come as no

surprise that solving an instance of the fractional cut-covering problem implies solving a maximum cut instance, which can be established via standard results on the ellipsoid method [17] and (6).

The rich literature on the maximum cut problem contrasts with the scarcity of papers on the fractional cut-covering problem. To the best of our knowledge, we introduced the weighted fractional cut-covering problem (4) in our recent work [5]. Šámal [42] introduced the unweighted version  $\text{fcc}(G) := \text{fcc}(G, \mathbb{1})$  to study the existence of cut-continuous maps between graphs. Neto and Ben-Ameur [33] surveyed lower and upper bounds on  $\text{fcc}(G)$ . There is more literature on the unweighted *cut-covering problem*, where one wishes to find a smallest collection of cuts containing every edge in the graph [18, 20, 28, 30]. Studying fractional versions of combinatorial problems is a staple of combinatorial optimization, with [17] already proving that computing the fractional chromatic number is NP-hard. There is recent interest from the combinatorics community on fractional versions of their invariants [9, 24, 45].

In [5], we presented a randomized SDP-based primal-dual algorithm that receives as input an instance  $(G, w)$  of the maxcut problem, computes a paired instance  $(G, z)$  for the weighted fractional cut-covering problem (or vice-versa), and simultaneously approximately solves both while providing a certificate for the quality of the approximation. This paper serves as the computational counterpart to [5], conducting an extensive computational study of that framework with some algorithmic refinements. We thus arrive at the central question this paper explores:

In practice, with reasonable resource usage,  
(Practicality) can we simultaneously approximate mc and fcc?

**Our framework.** We now provide an overview of the framework implemented in this study (for a detailed description see Section 2.1). Our algorithm in [5] builds upon the algorithm by Goemans and Williamson [16]. Assume an instance  $(G, w)$  of the maximum cut problem is given as input. For each  $Y \in \mathbb{S}_+^V$  we denote by  $\text{GW}(Y)$  a random shore produced by the random hyperplane technique of [16] applied to the matrix  $Y$ . To approximately compute  $\text{mc}(G, w)$ , as well as to compute an approximately optimal  $z \in \mathbb{R}_+^E$  in (7), we:

STEP 1: compute an optimal solution  $Y \in \mathbb{S}_+^V$  to (2a), and set  $z := \frac{1}{4}\mathcal{L}_G^*(Y)$ ;

STEP 2: let  $\mathcal{F} \subseteq \mathcal{P}(V)$  be a finite set of shores independently sampled from  $\text{GW}(Y)$ ;

STEP 3: output a shore  $S \in \mathcal{F}$  that maximizes  $w^\top \mathbb{1}_{\delta(S)}$  over  $\mathcal{F}$  and a fractional cut cover  $y \in \mathbb{R}_+^{\mathcal{F}}$  for  $(G, z)$  that minimizes  $\mathbb{1}^\top y$  over fractional cut covers in  $\mathbb{R}_+^{\mathcal{F}}$ .

STEP 3 is an algorithmic refinement over [5]: here we optimize  $y$  by solving the LP in (4) only with the shores in  $\mathcal{F}$ , whereas in [5],  $y$  is obtained from  $\mathcal{F}$  by scaling the average of the sampled incidence vectors. One of the main features of the algorithm in [5] is the construction of a simple, nearly combinatorial certificate for the approximation ratio of both problems (see Section 2.2), which our pipeline also produces.

Naturally, the framework performs well for the maxcut problem since each sample of  $\text{GW}(Y)$  yields a cut with expected weight at least  $\alpha_{\text{GW}}\eta(G, w)$ , implying that a reasonable number of independent samples should provide good concentration. For the fractional cut cover, the situation is more challenging since every single edge must be covered. In fact, one of the main obstacles to obtaining good results is the existence of *thin edges*, i.e., edges with relative weight much smaller than others. Since  $\text{GW}(Y)$  favors large cuts, thin edges are naturally harder to cover. [5] proposes multiple solutions for this issue, which we explore in this paper. Among these solutions, and central to this work, is the idea of *perturbation*: instead of sampling from the nearly optimal solution  $Y$  to (2a), we sample from  $(1 - \varepsilon)Y + \varepsilon I$  for a small  $\varepsilon \in (0, 1)$ .

**Our contributions.** This paper provides a comprehensive computational validation of our primal-dual SDP framework in [5] with the linear optimization refinement described in Step 3, demonstrating that it delivers practical algorithms for simultaneously solving maximum cut and fractional cut-covering problems with strong empirical performance. Through computational experiments, we

explore the ?? of our framework, addressing metrics such as: the number of samples needed in practice, the approximation ratio that can be achieved, the usage of time and memory, the quality of the certificates for our primal-dual problems, and performance when starting from either a maximum cut instance or a fractional cut covering instance. To the best of our knowledge, this work is the first one to provide computational experiments for the weighted version of the fractional cut-covering problem, with [33] having some explorations on bounds for  $\text{fcc}(G, \mathbb{1})$ , but not producing actual cut covers. On the maximum cut side, as the primal-dual framework from [5] is new, we have in our experiments a novel context to measure and evaluate maximum cut algorithms.

We gather experimental evidence supporting the ?? of our framework using the dataset from [31], a recent experimental study of maximum cut algorithms. To explore performance under resource constraints, we run experiments on three machines: a budget personal computer, a higher-performance personal computer, and a shared university server.

We answer ?? in the affirmative: the max-cut problem and the fractional cut-covering problem can indeed be approximated simultaneously with reasonable resource usage. Below we outline our key experimental findings:

- (1) Sample sizes significantly smaller than theoretical bounds suffice in practice, with perturbation playing a critical role in ensuring edge coverage. Obtaining fractional cut covers from modest sample sizes is a significant empirical finding.
- (2) Approximation ratios exceeded the theoretical guarantee of  $\alpha_{\text{GW}}$  in nearly all instances, even with sample sizes below theoretical predictions, considering the use of perturbation and the linear optimization refinement.
- (3) The algorithms deliver repeatable, consistent results across runs despite their probabilistic nature.
- (4) The framework produces approximate optimality certificates for both maximum cut and fractional cut-covering, validating the theoretical primal-dual guarantees. Crucial to this guarantee is the use of *sanitization* to obtain, from SDP solver outputs, optimality certificates that can be easily checked using floating-point arithmetic.
- (5) The framework works well with first-order and second-order SDP solvers, in part due to pre-processing of instances and post-processing of SDP outputs.
- (6) Floating-point implementations preserve theoretical guarantees satisfactorily. This is achieved by building on simple primitives (such as square-root computation and normal distribution sampling) and using robust solver technology.

**1.1. Notation and Organization of the Text.** We denote by  $\mathbb{S}^V$  the set of symmetric matrices indexed by the finite set  $V$ , and by  $\mathbb{S}_+^V \subseteq \mathbb{S}^V$  the set of positive semidefinite matrices. We use the *Löwner order*, writing  $A \succeq B$  or  $B \preceq A$  as a shorthand for  $A - B \in \mathbb{S}_+^V$ . We equip the space of symmetric matrices with the *trace inner product*, writing  $\langle A, B \rangle = \text{Tr}(AB) = \sum_{i \in V} (AB)_{ii}$ . We denote by  $\|X\|_{\text{Fr}} := \langle X, X \rangle^{1/2}$  the *Frobenius norm* of a symmetric matrix. For vector inner products, we usually write  $w^\top z := \sum_{e \in E} w_e z_e$  for every  $w, z \in \mathbb{R}^E$ , as we have done in the introduction. We denote by  $\mathbb{1}$  the all-ones vector, and by  $\mathbb{1}_F \in \{0, 1\}^E$  the incidence vector of  $F \subseteq E$ .

Let  $G = (V, E)$  be a graph. Throughout the text, we denote by  $n := |V|$  the number of vertices of the graph, and by  $m := |E|$  the number of edges of a graph.

Section 2 presents the theoretical framework which our algorithm builds upon. Section 3 presents our first set of experiments, taking as input maximum cut instances. Section 4 presents experiments which take as input fractional cut-covering instances. Section 5 studies different approaches to ensure feasibility with  $O(\ln m)$  samples. Section 6 studies different choices of solvers. Section 7 compares the LP-solving approach with the simpler, averaging-and-scaling approach presented in [5].

## 2. THEORETICAL AND COMPUTATIONAL FOUNDATIONS

The theoretical framework we present here is from [5]; we highlight the concepts crucial for this manuscript. First, we present our primal-dual framework in exact arithmetic. This allows us to then discuss how we handle numerical errors in our algorithms. We then present our *pipeline*, a single abstraction which captures all the experiments performed in this manuscript.

Our first task is to define the meaning of *simultaneously solving* the fractional cut-covering and maximum cut problems. Equations (6) and (7) lead us into the following definition:

**Definition 1.** Let  $G = (V, E)$  be a graph and let  $\beta \in (0, 1]$ . A  $\beta$ -pairing on  $G$  is a pair  $(w, z) \in \mathbb{R}_+^E \times \mathbb{R}_+^E$  such that there exist  $\rho, \mu \in \mathbb{R}_+$ , such that  $\rho = 0 = \mu$  if and only if  $w = 0 = z$ , and

$$(9) \quad w^\top z \stackrel{(9a)}{=} \rho\mu \quad \text{and} \quad \beta\rho\mu \stackrel{(9b)}{\leq} \text{mc}(G, w) \stackrel{(9c)}{\leq} \rho\mu \stackrel{(9d)}{\leq} \rho \text{fcc}(G, z) \stackrel{(9e)}{\leq} \frac{1}{\beta}\rho\mu.$$

We define an *exact pairing* on  $G$  to be a 1-pairing on  $G$ .

When  $\rho > 0$  and  $\mu > 0$ , which we regard as the ‘‘typical’’ case, we may restate (9) as

$$w^\top z \stackrel{(9a)}{=} \rho\mu, \quad \beta\rho \stackrel{(9b)}{\leq} \text{mc}(G, w) \stackrel{(9c)}{\leq} \rho, \quad \text{and} \quad \mu \stackrel{(9d)}{\leq} \text{fcc}(G, z) \stackrel{(9e)}{\leq} \frac{1}{\beta}\mu.$$

The definition is made to accommodate the case  $0 \in \{\rho, \mu\}$ .

The goal of our algorithms is to obtain  $\beta$ -pairings, either given  $G = (V, E)$  and  $w \in \mathbb{R}_+^E$  as input, or  $G$  and  $z \in \mathbb{R}_+^E$  as input. Yet, note that the definition of  $\beta$ -pairing does not specify how one can prove either of the inequalities involved. The inequalities (9c) and (9d) are particularly challenging, as they require either limiting the value attainable by any cut on the weighted graph  $(G, w)$ , or the value of any fractional cut cover for  $z$ . This is a common situation when working with approximation algorithms, and the SDP relaxations effectively address this issue. We have that

$$(10) \quad \text{mc}(G, w) \leq \mathbb{1}^\top x \text{ if } \text{Diag}(x) \succeq \frac{1}{4}\mathcal{L}_G(w),$$

since, for every  $S \subseteq V$ , we have

$$\mathbb{1}_{\delta(S)}^\top w = \langle (2\mathbb{1}_S - \mathbb{1})(2\mathbb{1}_S - \mathbb{1})^\top, \frac{1}{4}\mathcal{L}_G(w) \rangle \leq \langle (2\mathbb{1}_S - \mathbb{1})(2\mathbb{1}_S - \mathbb{1})^\top, \text{Diag}(x) \rangle = \mathbb{1}^\top x.$$

Hence  $\text{Diag}(x) \succeq \frac{1}{4}\mathcal{L}_G(w)$  certifies (9c) and we can choose any  $\rho \geq \mathbb{1}^\top x$  as an upper bound for  $\text{mc}(G, w)$ . We also get a certificate for (9d), since, by (6),

$$\text{fcc}(G, z) = \max\{z^\top w : w \in \mathbb{R}_+^E, \text{mc}(G, w) \leq 1\} \geq z^\top \left( \frac{1}{\mathbb{1}^\top x} w \right),$$

and we can set  $\mu := w^\top z / \rho$  so that (9a) holds. To certify (9b) and (9e), we simply need a cut of weight at least  $\beta\rho$  and a fractional cut cover with value at most  $\frac{1}{\beta}\mu$ . Hence, we arrive at an strategy to certify  $(w, z)$  to be a  $\beta$ -pairing.

**Definition 2.** Let  $G = (V, E)$  be a graph and let  $\beta \in (0, 1]$ . Let  $(w, z) \in \mathbb{R}_+^E \times \mathbb{R}_+^E$ . A  $\beta$ -certificate for  $(w, z)$  is a tuple  $(\rho, \mu, S, y, x, B)$  such that  $\rho = 0 = \mu$  if and only if  $w = 0 = z$ , and

$$(11.i) \quad \rho, \mu \in \mathbb{R}_+ \text{ are such that } \rho\mu = w^\top z,$$

$$(11.ii) \quad S \subseteq V \text{ is such that } w^\top \mathbb{1}_{\delta(S)} \geq \beta\rho,$$

$$(11.iii) \quad y \in \mathbb{R}_+^{\mathcal{P}(V)} \text{ is such that } \sum_{U \subseteq V} y_U \mathbb{1}_{\delta(U)} \geq z \text{ and } \mathbb{1}^\top y \leq \frac{1}{\beta}\mu, \text{ and}$$

$$(11.iv) \quad x \in \mathbb{R}^V \text{ and } B \in \mathbb{R}^{V \times V} \text{ are such that } \rho \geq \mathbb{1}^\top x \text{ and } \text{Diag}(x) - \frac{1}{4}\mathcal{L}_G(w) = B^\top B.$$

Given an instance  $(G, w)$  for the max-cut problem or an instance instance  $(G, z)$  for the fractional cut-covering problem, our goal is obtaining a  $\beta$ -pairing together with a  $\beta$ -certificate. The value of  $\beta$  is mostly determined by the value of the solution obtained for the SDP relaxation, but we aim at  $\beta \geq \alpha_{\text{GW}}$ , the approximation rate guarantee in the algorithm by Goemans and Williamson [16].

**2.1. Algorithm in the Real-Number Machine Model.** This section presents the algorithm assuming exact arithmetic and that optimal SDP solutions are obtained.

**First step: solving an SDP relaxation.** Our framework takes as input either a maximum cut instance  $(G, w)$  or a fractional cut-covering instance  $(G, z)$ , and computes the corresponding paired instance. The two cases differ only in the first step: computing an optimal SDP solution  $Y \in \mathbb{S}_+^V$ , which is a relaxation for the maximum cut problem in the first case and for the fractional cut-covering problem in the second. When starting from a maximum cut instance, we described the first step as

STEP 1: (mc) compute an optimal solution  $Y \in \mathbb{S}_+^V$  to (2a), and set  $z := \frac{1}{4}\mathcal{L}_G^*(Y)$ .

Note how we obtain the instance for the fractional cut covering problem as  $(G, z)$  with  $z := \frac{1}{4}\mathcal{L}_G^*(Y)$ . When starting from a fractional cut covering instance, we use the following SDP relaxation. Let  $G = (V, E)$  be a graph, and let  $z \in \mathbb{R}_+^E$ . We define

$$(12a) \quad \eta^\circ(G, z) := \min\{\mu : \mu \in \mathbb{R}_+, Y \in \mathbb{S}_+^V, \text{diag}(Y) = \mu \mathbb{1}, \frac{1}{4}\mathcal{L}_G^*(Y) \geq z\}$$

$$(12b) \quad = \max\{z^\top w : w \in \mathbb{R}_+^E, x \in \mathbb{R}^V, \text{Diag}(x) \succeq \frac{1}{4}\mathcal{L}_G(w), \mathbb{1}^\top x \leq 1\}$$

Both problems admit Slater points and so, by Strong Duality, equality holds and both problems attain their optimal value. The first step then is

STEP 1: (fcc) compute an optimal solution  $Y \in \mathbb{S}_+^V$ ,  $w \in \mathbb{R}_+^E$ , and  $x \in \mathbb{R}^V$  to (12).

In practice, we only approximately solve either (2) or (12). Our framework includes a `Solver` for this task and a `Scaler` to normalize the input before solving and rescale the output afterwards.

The *perturbation* mentioned in Section 1 is implemented in this step and differs between the two starting points. For the maximum cut problem, we perturb the nearly optimal solution  $Y$  to (2a) to  $(1 - \varepsilon)Y + \varepsilon I$  for a small  $\varepsilon \in (0, 1)$ , so that  $z$  becomes  $\frac{1}{4}\mathcal{L}_G^*((1 - \varepsilon)Y + \varepsilon I)$ . For the fractional cut covering problem, the perturbation is applied to the SDP relaxation itself as in (33) in Section 4.

**Second step: sampling cuts.** For any positive semidefinite matrix  $Y \in \mathbb{S}_+^V$  and  $g \in \mathbb{R}^V$ , set  $\text{GW}(Y, g) := \{i \in V : e_i^\top Y^{1/2} g \geq 0\}$ , where  $Y^{1/2} \in \mathbb{S}_+^V$  is the unique positive semidefinite square root of  $Y$ . Let  $(\Omega, \Sigma, \mathbb{P})$  be a probability space, and let  $h: \Omega \rightarrow \mathbb{R}^V$  be a uniformly distributed unit vector. We denote by  $\text{GW}(Y): \Omega \rightarrow \mathcal{P}(V)$  the random shore defined by  $\omega \in \Omega \mapsto \text{GW}(Y, h(\omega))$ . In particular,  $\text{GW}(Y)$  is a possible implementation of the randomized rounding algorithm in [16]. The second step is then

STEP 2: let  $\mathcal{F} := \{S_1, \dots, S_T\}$  be a finite set of shores independently sampled from  $\text{GW}(Y)$ ;

The main idea is to choose  $T \in \mathbb{N}$  large enough so (with high probability) some sampled shore induces a  $\beta$ -approximate maximum cut, and the collection  $\mathcal{F}$  is rich enough to yield a good fractional cut covering. For maximum cut this is straightforward: each sample is independent and  $\text{GW}(Y)$  yields a cut with expected weight at least  $\alpha_{\text{GW}}\eta(G, w)$ , so high concentration follows directly from Chernoff's inequality. But why should  $\mathcal{F}$  contain cuts that cover  $z \in \mathbb{R}_+^E$ ? The strategy is motivated by the observation that

$$(13) \quad \mathbb{E}[\mathbb{1}_{\delta(\text{GW}(Y))}] \geq \alpha_{\text{GW}} \frac{1}{4} \mathcal{L}_G^*(Y).$$

Since  $\frac{1}{4}\mathcal{L}_G^*(Y) \geq z$  holds whether we solved (2) or (12), this implies  $\frac{1}{\alpha_{\text{GW}}}\mathbb{E}[\mathbb{1}_{\delta(\text{GW}(Y))}] \geq z$ . Inequality (13) holds since

$$\mathbb{P}(ij \in \text{GW}(Y)) = \frac{\arccos Y_{ij}}{\pi} \geq \alpha_{\text{GW}} \frac{1 - Y_{ij}}{2}$$

for every edge  $ij \in E$ . The formula for the probability of an edge being covered comes from [16], and the inequality follows from the definition of  $\alpha_{\text{GW}}$ . Using perturbation, we guarantee that each  $z_{ij}$  is bounded away from zero, so every edge has a non-negligible probability of being covered. Chernoff's inequality combined with a union bound over the edges then gives high probability of covering all edges.

Our framework includes a **Sampler** component for this step. We also explore mixing in a small number of cuts generated by including each vertex in the shore independently with probability  $\frac{1}{2}$  (see Section 5.3), which are sampled by a component we call **Presampler**.

**Third step: computing a cut and a fractional cut cover.** At this point we have a rich collection of shores  $\mathcal{F} \subseteq \mathcal{P}(V)$ . We define

$$(14) \quad \text{mc}(\mathcal{F}, w) := \max\{w^\top \mathbb{1}_{\delta(S)} : S \in \mathcal{F}\} \quad \text{for every } w \in \mathbb{R}_+^E$$

and

$$(15) \quad \text{fcc}(\mathcal{F}, z) := \min\left\{\mathbb{1}^\top y : y \in \mathbb{R}_+^{\mathcal{F}}, \sum_{S \in \mathcal{F}} y_S \mathbb{1}_{\delta(S)} \geq z\right\} \quad \text{for every } z \in \mathbb{R}_+^E.$$

The utility of these definitions is immediate: if  $|\mathcal{F}|$  is bounded by a polynomial on the size of the graph  $G$ , then both  $\text{mc}(\mathcal{F}, w)$  and  $\text{fcc}(\mathcal{F}, z)$ , for any  $w \in \mathbb{R}_+^E$  and  $z \in \mathbb{R}_+^E$ , are computable in polynomial time. Moreover,

$$(16) \quad \text{mc}(\mathcal{F}, \cdot) \leq \text{mc}(G, \cdot) \text{ and } \text{fcc}(\mathcal{F}, \cdot) \geq \text{fcc}(G, \cdot).$$

Our third step computes  $\text{mc}(\mathcal{F}, w)$  and  $\text{fcc}(\mathcal{F}, z)$ :

STEP 3: output a shore  $S \in \mathcal{F}$  maximizing  $w^\top \mathbb{1}_{\delta(S)}$  over  $\mathcal{F}$ , and a fractional cut cover  $y \in \mathbb{R}_+^{\mathcal{F}}$  for  $(G, z)$  minimizing  $\mathbb{1}^\top y$  over fractional cut covers in  $\mathbb{R}_+^{\mathcal{F}}$ .

Computing  $\text{mc}(\mathcal{F}, w)$  amounts to selecting the shore inducing the largest-weight cut with respect to  $w$ . Computing  $\text{fcc}(\mathcal{F}, z)$  requires solving a linear program restricted to the shores in  $\mathcal{F}$ ; our framework includes a **Cover Producer** component using Gurobi for this task.

Given the probabilistic nature of our algorithms, it is possible that the LP defining  $\text{fcc}(\mathcal{F}, z)$  is infeasible, which is an outcome reflecting insufficient samples taken. We call this outcome *Restricted LP Infeasible* (RLI):

$$(RLI) \quad \text{there exists } ij \in E, \text{ s.t. } z_{ij} > 0 \text{ and } ij \notin \bigcup_{S \in \mathcal{F}} \delta(S).$$

**Certificates.** Let  $G = (V, E)$  be a graph. Whether we are given a maximum cut instance or a fractional cut-covering instance as input, we always end up with a pair  $z \in \mathbb{R}_+^E$  and  $w \in \mathbb{R}_+^E$  and feasible solutions to the following problems:

$$(2b') \quad \eta(G, w) = \min\{\rho \in \mathbb{R}_+ : x \in \mathbb{R}^V, \text{Diag}(x) \succeq \frac{1}{4}\mathcal{L}_G(w), \mathbb{1}^\top x \leq \rho\},$$

$$(12a) \quad \eta^\circ(G, z) = \min\{\mu \in \mathbb{R}_+ : Y \in \mathbb{S}_+^V, \text{diag}(Y) = \mu \mathbb{1}, \frac{1}{4}\mathcal{L}_G^*(Y) \geq z\}.$$

Indeed, given a maximum cut instance  $(G, w)$  as input, if  $Y \in \mathbb{S}_+^V$  is feasible in (2a), then  $(1, Y)$  is feasible in (12a) for  $z := \frac{1}{4}\mathcal{L}_G^*(Y)$ . Dually, given a fractional cut-covering instance  $(G, z)$ , if  $(w, x)$  is feasible in (12b), then  $(1, x)$  is feasible in (17a) for  $w$ .

Now let  $z \in \mathbb{R}_+^E$  and  $w \in \mathbb{R}_+^E$ , and let  $(\rho, x)$  and  $(\mu, Y)$  be feasible in (17a) and (12a). STEP 3 obtains a shore  $S \subseteq V$  such that  $w^\top \mathbb{1}_{\delta(S)} = \text{mc}(\mathcal{F}, w)$  and a fractional cut cover  $y \in \mathbb{R}_+^{\mathcal{P}(V)}$  with support in  $\mathcal{F}$  with value  $\text{fcc}(\mathcal{F}, z)$ . Set

$$(17) \quad \sigma = 1 - \frac{z^\top w}{\rho \mu}, \quad \beta_{\text{fcc}} := \frac{(1 - \sigma)\mu}{\text{fcc}(\mathcal{F}, z)}, \quad \beta_{\text{mc}} := \frac{\text{mc}(\mathcal{F}, w)}{\rho}.$$

The value of  $\sigma$  measures how close to optimality were the feasible solutions used to obtain the pairing between  $w$  and  $z$ . Indeed,  $\sigma \geq 0$ , since

$$z^\top w \leq \langle \frac{1}{4}\mathcal{L}_G^*(Y), w \rangle = \langle Y, \frac{1}{4}\mathcal{L}_G(w) \rangle \leq \langle Y, \text{Diag}(x) \rangle \leq \mu \rho.$$

In particular, “nearly solving the relevant SDP” means computing feasible solutions making  $\sigma$  small. We claim that

$$(18) \quad \beta_{\text{fcc}} \leq \beta_{\text{mc}}.$$

Indeed, for every  $y \in \mathbb{R}_+^{\mathcal{F}}$  such that  $\sum_{S \in \mathcal{F}} y_S \mathbb{1}_{\delta(S)} \geq z$ ,

$$(1 - \sigma)\rho\mu \leq z^\top w \leq \sum_{S \in \mathcal{F}} y_S \mathbb{1}_{\delta(S)}^\top w \leq y^\top \mathbb{1}_{\text{mc}}(\mathcal{F}, w).$$

Hence  $(1 - \sigma)\rho\mu \leq \text{fcc}(\mathcal{F}, z) \text{mc}(\mathcal{F}, w)$ , and thus

$$\beta_{\text{fcc}} = \frac{(1 - \sigma)\mu}{\text{fcc}(\mathcal{F}, z)} \leq \frac{\text{mc}(\mathcal{F}, w)}{\rho} = \beta_{\text{mc}}.$$

The definitions of  $\beta_{\text{mc}}$  and  $\beta_{\text{fcc}}$  are made so as to ensure that both values are always at most one. We do not use the definitions in (18) when the Restricted LP (15) is Infeasible (RLI).

**2.2. Floating-Point Arithmetic and Certificate Computation.** Let  $G = (V, E)$  be a graph, and let  $w, z \in \mathbb{R}_+^E$ . Let  $\rho, \mu \in \mathbb{R}_+$  be such that

$$(19) \quad \eta(G, w) \leq \rho, \eta^\circ(G, z) \leq \mu.$$

As we just saw, these inequalities play a crucial role in constructing  $\beta$ -certificates. To store a numerical proof of ??, we store Cholesky factorizations of the positive semidefinite matrices appearing in (17a) and (12a). Namely, we store, for each  $w \in \mathbb{R}_+^E$  and  $z \in \mathbb{R}_+^E$ :

$$(20) \quad \begin{aligned} &\rho, \mu \in \mathbb{R}_+, R \in \mathbb{R}^{k \times V}, B \in \mathbb{R}^{V \times V}, x \in \mathbb{R}^V \text{ such that} \\ &\rho \geq \mathbb{1}^\top x \text{ and } \frac{1}{4} \|Re_i - Re_j\|^2 \geq z_{ij}, \forall ij \in E, \\ &\| \text{diag}(R^\top R) - \mu \mathbb{1} \|_2 \approx 0 \\ &\| \text{Diag}(x) - \frac{1}{4} \mathcal{L}_G(w) - B^\top B \|_{\text{Fr}} \approx 0. \end{aligned}$$

By prioritizing the certification of the positive semidefinite inequalities, we must handle numerical error in the constraints  $\text{diag}(R^\top R) = \mu \mathbb{1}$  and  $S^\top S = \text{Diag}(x) - \frac{1}{4} \mathcal{L}_G(w)$ .

A certificate of the type (20) is different from the direct output of a solver in one important way: a solver may output solutions which are not feasible, but only close to an actual feasible solution. On the other hand, the point of (20) is to have a *numerical proof* of feasibility to (17a) and (12a). The act of computing (20) from the solver output is called *sanitize* in our implementation, and it must bridge this distinction. Intuitively, we accomplish this by “moving numerical error into the objective value”, where we weaken the objective values produced in order to ensure feasibility. More precisely, we present in Algorithms 1 and 2 how we compute (20) from the information produced by a solver. Both algorithms use LAPACK’s routines to perform numerical linear algebra.

An important invariant of both algorithms is that they do not modify neither  $w$  nor  $z$ . This is necessary to provide uniform treatment for both the case in which a maximum cut instance is given as input, and the case in which a fractional cut-covering instance is given as input. Furthermore, they are by design simple procedures that either fail, or provide the feasibility guarantee needed by (20). For the sake of simplicity of exposition, the pseudocode represents all failure cases with  $\perp$ ; in practice, the failure cases are properly discriminated.

Only one type of failure was observed when running Algorithms 1 and 2. Algorithm 2 returned  $\perp$  for some inputs, which occurs when

$$(\perp_Y) \quad \exists ij \in E, \text{ s.t. } z_{ij} > 0 \text{ and } \|Re_i - Re_j\|_2^2 = 0$$

for  $R \in \mathbb{R}^{V \times k}$  as computed in step (6) of Algorithm 2. In exact arithmetic, this is impossible, as it contradicts the assumption that the input  $Y$  satisfies  $\frac{1}{4} \mathcal{L}_G^*(Y) \geq z$ .

---

**Algorithm 1** Sanitize slack inequality “ $\text{Diag}(x) \succeq \frac{1}{4}\mathcal{L}_G(w)$ ”
 

---

**Parameters:**  $\gamma \in (0, 1)$ .

 $\triangleright$  We have used  $\gamma := 10^{-8}$  throughout.

**Input:**  $\tilde{x} \in \mathbb{R}^V$  and  $w \in \mathbb{R}_+^E$ 
**Output:**  $\text{SLACKSANITIZE}_\gamma(\tilde{x}, w)$  can either fail, returning  $\perp$ , or return  $(\rho, x, S)$  such that

$$\frac{1}{n} \|\text{Diag}(x) - \frac{1}{4}\mathcal{L}_G(w) - S^\top S\|_{\text{Fr}} \approx 0 \text{ and } \rho = \mathbb{1}^\top x.$$

```

1 procedure SLACKSANITIZE $_\gamma(\tilde{x}, w)$ 
2    $\tau \leftarrow \lambda_{\min}(\text{Diag}(\tilde{x}) - \frac{1}{4}\mathcal{L}_G(w))$   $\triangleright$  Computed with LAPACK's dsyevd
3   if  $\tau < 0$  then  $\tilde{x} \leftarrow \tilde{x} + (\gamma - \tau)\mathbb{1}$ 
4    $B \leftarrow \text{Cholesky}(\text{Diag}(\tilde{x}) - \frac{1}{4}\mathcal{L}_G(w))$   $\triangleright$  Computed with LAPACK's dpotrf2
5   if Cholesky computation failed then return  $\perp$ 
6    $\rho \leftarrow \mathbb{1}^\top \tilde{x}$ 
7   return  $(\rho, \tilde{x}, B)$ 
    
```

---



---

**Algorithm 2** Sanitize “ $Y \in \mathbb{S}_+^V$  such that  $\frac{1}{4}\mathcal{L}_G^*(Y) \geq z$ ”
 

---

**Parameters:**  $\gamma \in [0, 1)$ .

 $\triangleright$  We have always used  $\gamma \in \{0, 10^{-8}\}$ .

**Input:**  $\tilde{Y} \in \mathbb{S}^V$  and  $z \in \mathbb{R}_+^E$  with  $\|z\|_\infty = 1$ .

**Output:**  $\text{REPRESENTATIONSANITIZE}_\gamma(\tilde{Y}, z)$  can either fail, returning  $\perp$ , or return  $(\mu, R)$  such that

$$\frac{1}{4} \|Re_i - Re_j\|_2^2 \geq z_{ij}, \forall ij \in E \text{ and } \|\text{diag}(R^\top R) - \mu\mathbb{1}\|_2 \approx 0,$$

```

1 procedure REPRESENTATIONSANITIZE $_\gamma(\tilde{Y}, z)$ 
2   Set  $h_i \leftarrow \tilde{Y}_{ii}^{-\frac{1}{2}}$  for every  $i \in V$ 
3   Set  $Y_0 \leftarrow \text{Diag}(h)\tilde{Y}\text{Diag}(h)$ 
4   Let  $R_0 \in \mathbb{R}^{V \times V}$ ,  $\lambda \in \mathbb{R}^V$  be such that  $Y_0 = R_0 \text{Diag}(\lambda)R_0^\top$ .  $\triangleright$  Computed with dsyevd
5   Let  $R_1 \in \mathbb{R}^{V \times k}$  and  $\nu \in \mathbb{R}^k$  contain all eigenpairs satisfying  $\lambda_i > \gamma$ 
6    $R \leftarrow \text{Diag}(\nu)^{\frac{1}{2}}R_1^\top$ 
7    $\mu \leftarrow 0$ 
8   for each  $ij \in E$  do  $\mu \leftarrow \max\left(\mu, \frac{z_{ij}}{\frac{1}{4}\|Re_i - Re_j\|_2^2}\right)$ 
9   if  $\mu \in \{+\infty, 0\}$  then return  $\perp$ 
10  return  $(\mu, \sqrt{\mu}R)$ 
    
```

---

Let  $G = (V, E)$  be a graph, and let  $w, z \in \mathbb{R}_+^E \setminus \{0\}$ . To account for floating-point error, we enhance Corollary 2 by saying that a  $\beta$ -certificate is a tuple  $(\rho, \mu, S, y, x, B, \tau)$  satisfying (11.i) to (11.iii) and

(11.iv')  $x \in \mathbb{R}^V$  and  $B \in \mathbb{R}^{V \times V}$  are such that  $\rho \geq \mathbb{1}^\top x$  and  $\|\text{Diag}(x) - \frac{1}{4}\mathcal{L}_G(w) - B^\top B\|_{\text{Fr}} \leq \tau\rho$ .

We claim that if  $(\rho, \mu, S, y, x, B, \tau)$  is a  $\beta$ -certificate, then

$$(21) \quad (w, z) \text{ is a } \tilde{\beta}\text{-pairing, with } \tilde{\beta} := \frac{1}{1 + \tau n}\beta.$$

Since  $\|\text{Diag}(x) - \frac{1}{4}\mathcal{L}_G(w) - \tilde{B}^\top \tilde{B}\|_{\text{Fr}} \leq \tau\rho$ , it follows that  $\text{Diag}(x) - \frac{1}{4}\mathcal{L}_G(w) - B^\top B \succeq -\tau\rho I$ . Thus  $\text{Diag}(x + \tau\rho\mathbb{1}) - \frac{1}{4}\mathcal{L}_G(w) \succeq B^\top B \succeq 0$ , so  $\text{mc}(G, w) \leq (1 + \tau n)\rho$  by (10). Using (11.ii) we have

$$(1 + \tau n)\rho \geq \text{mc}(G, w) \geq w^\top \mathbb{1}_{\delta(S)} \geq \beta\rho = \tilde{\beta}(1 + \tau n)\rho,$$

and using (11.i), (6), and (11.iii), we have

$$\frac{\mu}{1 + \tau n} = \frac{w^\top z}{(1 + \tau n)\rho} \leq \text{fcc}(G, z) \leq \mathbb{1}^\top y \leq \beta^{-1}\mu = \tilde{\beta}^{-1} \frac{\mu}{1 + \tau n}.$$

Hence (21) holds.

Let  $G = (V, E)$  be a graph, and let  $w, z \in \mathbb{Q}_+^E$ . Let  $(\rho, \mu, S, y, x, B)$  be a  $\beta$ -certificate as first defined, with  $\rho, \mu \in \mathbb{Q}$ ,  $y \in \mathbb{Q}^{\mathcal{F}}$ ,  $x \in \mathbb{Q}^V$ , and  $B \in \mathbb{Q}^{V \times V}$ . If we compute

$$t := \frac{1}{\rho^2} \|\text{Diag}(x) - \frac{1}{4}\mathcal{L}_G(w) - B^\top B\|_{\text{Fr}}^2 \in \mathbb{Q},$$

we get a formal proof, verifiable in a Turing machine, that  $(w, z)$  is a  $\tilde{\beta}$ -pairing for  $\tilde{\beta} := (1 + \sqrt{tn})^{-1}\beta$ . We did not compute such values of  $t$ , and hence our experiments do not produce formal proofs. In floating-point arithmetic, the largest value of  $\tau$  we have observed was  $1.231 \times 10^{-16}$ , and the largest value of  $\tau n$  was  $2.068 \times 10^{-14}$ .

**Proposition 3.** Let  $G = (V, E)$  be a graph. Let  $w \in \mathbb{R}_+^E$  and  $z \in \mathbb{R}_+^E$  be such that  $(\rho, \tilde{\mu}, R, B, x)$  satisfying (20) exist. Let  $\mathcal{F} \subseteq \mathcal{P}(V)$  be such that  $\text{fcc}(\mathcal{F}, z) < +\infty$ . Set

$$\begin{aligned} \tau &:= \rho^{-1} \|\text{Diag}(x) - \frac{1}{4}\mathcal{L}_G(w) - B^\top B\|_{\text{Fr}}, \quad \sigma := 1 - \frac{w^\top z}{\rho\tilde{\mu}}, \\ \mu &:= (1 - \sigma)\tilde{\mu}, \quad \beta_{\text{mc}} := \frac{\text{mc}(\mathcal{F}, w)}{\rho}, \quad \beta_{\text{fcc}} := \frac{\mu}{\text{fcc}(\mathcal{F}, z)} = \frac{(1 - \sigma)\tilde{\mu}}{\text{fcc}(\mathcal{F}, z)}, \end{aligned}$$

and let  $S \subseteq V$  and  $y \in \mathbb{R}_+^{\mathcal{F}}$  be optimal solutions to  $\text{mc}(\mathcal{F}, w)$  and  $\text{fcc}(\mathcal{F}, z)$ , respectively. Then  $(\rho, \mu, S, y, x, B, \tau)$  is a  $\beta_{\text{fcc}}$ -certificate.

*Proof.* We have that (11.i) holds since  $w^\top z = \rho\tilde{\mu}(1 - \sigma) = \rho\mu$  by definition of  $\sigma$ . Using the definitions of  $S$  and  $\beta_{\text{mc}}$ , as well as  $\beta_{\text{mc}} \geq \beta_{\text{fcc}}$ , we have that  $w^\top \mathbb{1}_{\delta(S)} = \text{mc}(\mathcal{F}, w) = \beta_{\text{mc}}\rho \geq \beta_{\text{fcc}}\rho$ , so (11.ii) holds. Moreover, using the definitions of  $y$  and  $\beta_{\text{fcc}}$ , we have that  $\mathbb{1}^\top y = \text{fcc}(\mathcal{F}, z) = \frac{1}{\beta_{\text{fcc}}}\mu$ , so (11.iii) holds. Finally, (11.iv') holds by definition of  $\tau$ .  $\square$

**2.3. Pipeline.** To improve the scientific value of our data, the execution of every experiment in this work fits within a unified set of steps. We refer to this sequence of steps as the *pipeline*. This abstraction allows for blanket statements about all the data collected, and helps streamline the work involved in testing our implementations. A design priority was to produce verifiable certificates and to explicitly record parameters used in each execution — see Section A. Figure 1 provides an overview of our pipeline. We start by describing the main set of data structures used in our experiments:

- **Weighted Graph** is a graph with weights on each edge; depending on the context (i.e., what Solver does with it) this may describe a maximum cut or a fractional cut-covering instance.
- **Raw Output** is the nearly optimal solutions produced by the solver, represented by  $Y \in \mathbb{S}^V$ ,  $x \in \mathbb{R}^V$ , and  $\{w, z\} \subseteq \mathbb{R}_+^E$ .
- **SDP certificate** are the objects described in (20) produced by Algorithms 1 and 2.
- **Cuts** is a matrix in  $\{\pm 1\}^{V \times T}$ , where each column represents the shore of a cut.
- **$\beta$ -certificate** is the final output, as in Corollary 2.

Although there is a single implementation of **Sanitizer**, as described in Algorithms 1 and 2, there are many implementations of **Solver**, **Scaler**, **Presampler**, **Sampler**, and **Cover Producer**. These

provide the customization points, as selecting among the available implementations is how one produces certificates for the problem one is interested in.

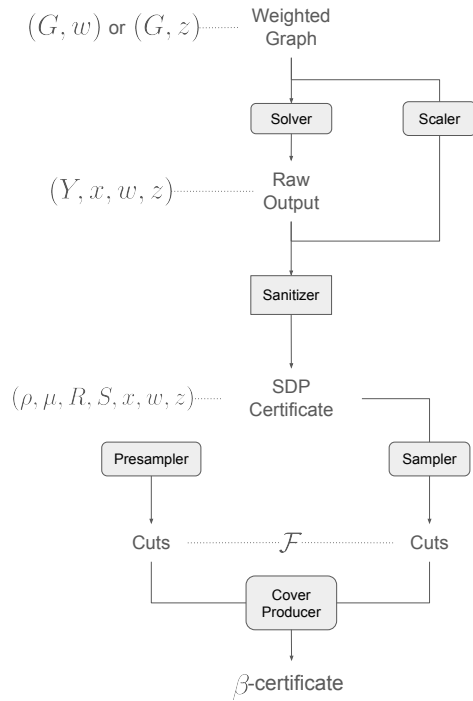


FIGURE 1. Pipeline

The customization points are:

- **Solver** receives as input an weighted graph, which it uses to approximately solve either (2) or (12). Different implementations correspond to different solvers and optimization problems.
- **Scaler** ensures that the weighted graph sent to **Solver** is normalized appropriately. After the SDP relaxation is solved, it scales back both the input (i.e., **Weighted Graph**) and the output of the solver (i.e. **Raw Output**).
- **Sampler** and **Presampler** are the same abstract class, which receives as input a **SDP certificate** and produces **Cuts**. They allow for **Cover Producer** to work with cuts generated in different ways.
- **Cover Producer** receives two sequence of **Cuts**, one arising from the **Presampler**, and one arising from the **Sampler**, and produces the final  $\beta$ -certificate.

### 3. MAIN ATTRIBUTES OF A PRACTICAL IMPLEMENTATION

The goal of this section is to empirically evaluate, in broad terms, what sizes of instances can be solved with an affordable personal computer, the running time needed to do so, as well as the range of certifiable approximation factors one can expect to obtain across instances. We keep our settings comparable to a recent experimental study by Mirka and Williamson [31], which implements several algorithms for the maximum cut problem.

The experiments in this work were performed in the following machines:

(Budget)                    personal laptop with 8 GB of RAM and  
one Intel i3-1115G4 processor 2-core 3.0 GHz (Tiger Lake),

(Laptop)                    personal laptop with 40GB of RAM and  
one AMD Ryzen 7 7730U processor 8-core 4.5 GHz (Ryzen 7000),

and

(Big)                        a shared university server with 768 GB of RAM and  
four Intel Xeon Gold 6230 20-core 2.1 GHz (Cascade Lake).

For solving our SDP relaxations, we used SCS [34] version 3.2.7 [35]. We used Gurobi to solve the LP (15). On ?? and ?? hardware we used Gurobi 12, and on ?? hardware we used Gurobi 10. The implementation of BLAS and LAPACK is provided by OpenBLAS on ?? hardware, and by Intel’s Math Kernel Library (MKL) version 2024.2 on ?? and ?? hardware. In particular, SCS relies on Intel MKL in at least two distinct ways: it uses Intel MKL Pardiso to solve linear systems, and uses LAPACK to project matrices into the positive semidefinite cone. For convenience, our use of BLAS and LAPACK is via the C++ library Armadillo [43] version 12.8. Pseudorandom number generation was managed by the `random` library from C++11, using the Mersenne Twister `std::mt19937` as the underlying engine. Normally distributed samples were obtained by passing the generated bits through `std::normal_distribution<double>`, which implements the standard normal distribution. Our implementation was sequential, with the tools we used employing parallelization. In particular, SCS and Gurobi ran multi-threaded code, either directly or via their use of Intel MKL. Some of our experiments were ran in parallel with the help of GNU Parallel [47]. For each machine, we first observed how many cores were used when running our experiments, and selected the number of parallel experiments so as to avoid contention.

Our maximum cut instances come directly from Mirka and Williamson [31]. They are divided into two sets of instances:

(Dense)                    dense weighted instances from TSPLib [40].

and

(Sparse)                  sparse unweighted (i.e.  $w = 1$ ) instances from the Matrix Market repository [29],

The one exception is `inf-USAir97`, which is an instance from ?? which has weights on its edges. Whereas the ?? instances have their number of vertices as part of the instance name (see Table 0.A), we report the number of vertices in each ?? graph in Table 2.B. The graphs in the ?? data set come from different real-world data, ranging from graphs arising from finite-element models employed to solve partial differential equations (`dwt_209` and `dwt_503`) to graphs arising from email interactions (`email-enrol-only` and `email-univ`). The ?? graphs were selected instances for testing solvers for the Travelling Salesman Problem. There is a small distinction between the instances used in [31] and those used in the broader TSP literature, which we discuss in Section E.

**3.1. A First Implementation and Calibration of Parameters.** In this subsection, we will describe a first implementation of our algorithms. In this first setting, we observe certain difficulties arising from trying to compute  $\beta$ -certificates. We contrast the experimental results with theoretical guarantees, setting the stage for the remainder of this manuscript.

The first batch of experiments were performed on ?? hardware. In these experiments, whose results we display in Tables 1.A and 1.B, we receive as input a graph  $G = (V, E)$  and nonnegative weights  $w \in \mathbb{R}_+^E$ . We normalize the weights to ensure that  $\|w\|_1 := \sum_{ij \in E} |w_{ij}| = 1$ .<sup>1</sup> We formulate (2b) as

$$(27) \quad \inf\{c^\top x : \mathcal{A}(x) \preceq b\} \text{ for } \mathcal{A}(x) := -\text{Diag}(x), \quad b := -\frac{1}{4}\mathcal{L}_G(w), \quad \text{and } c := \mathbb{1},$$

to solve it with SCS, using its default stopping criteria and default tolerances. We then sample  $\lceil 128 \ln m \rceil$  cuts via the random hyperplane technique, from the matrix

$$(28) \quad Y := (1 - \varepsilon)\tilde{Y} + \varepsilon I, \text{ where } \tilde{Y} \text{ is the solver output and } \varepsilon := 10^{-4}.$$

We refer to  $\varepsilon$  as the *perturbation parameter*. Let  $\mathcal{F} \subseteq \mathcal{P}(V)$  denote the set of shores sampled. We solve  $\text{fcc}(\mathcal{F}, z)$  as defined in (15) for  $z := \frac{1}{4}\mathcal{L}_G^*(Y) \in \mathbb{R}_+^E$ . The fact that  $Y$  is a convex combination of the nearly optimal solution with the identity matrix, as in (28), plays an important role throughout this our work. We discuss the role and effect of the perturbation parameter in Section 5.1. For now, however, we highlight that the choice for  $\varepsilon = 10^{-4}$  matches the default accuracy of solutions produced by SCS. Considering the numerical error intrinsic to floating-point arithmetic, both  $\tilde{Y}$  and  $Y$  have similar claims to near optimality, with the latter sidestepping issues arising from numerical rank computations. Table 0.A provides a concrete argument for perturbation. When selecting  $\varepsilon = 0$  in (28), not only do many instances fail to attain feasibility (which is recorded in the RLI column), the output of SCS for the instance `brg180` is rejected by Algorithm 2. In Tables 1.A and 1.B (as well as in Table 0.A), the columns  $\sigma$ ,  $\beta_{\text{mc}}$ , and  $\beta_{\text{fcc}}$  are computed according to (18). Each instance was solved 10 times with different random number seeds for the cut generation procedure. The columns “ $\beta_{\text{mc}}$  avg” and “ $\beta_{\text{mc}}$  std” report the average and standard deviation of all runs, whereas the analogous columns for  $\beta_{\text{fcc}}$  and “time (s)” only consider runs for which  $\text{fcc}(\mathcal{F}, z)$  was finite. Whenever the restricted LP (15) is infeasible, the running time of the pipeline is artificially improved by the LP solver quickly terminating. Hence such runs are excluded from the running time statistics.

Our first results are in Table 1.A, which shows the results for the set of ?? instances. For every instance we obtained a cut we could prove to be at least 94% of the optimal, and in most instances we produced a cut which is at least 99% of the optimal value. Moreover, the objective value found was quite stable, returning the same approximation ratio across the 10 runs for all but one instance. This robustness in stark contrast with the outcomes when computing fractional cut covers. In particular, we failed to produce a feasible cover in all 10 independent runs for the instance `kroA100`, and in 9 of 10 independent runs for the instance `gr137`. In general, we either failed to achieve feasibility, or obtained a solution with  $\beta_{\text{fcc}}$  better than the theoretical guarantee  $\alpha_{\text{GW}}$ . Instances `ch150` and `eil101` stand out: for both, we obtained an infeasible restricted LP in one run, but in the remaining 9 we obtained a certificate of quality at least 87%. These numbers suggest that sampling  $\lceil 128 \ln m \rceil$  cuts is not enough to reliably produce good empirical behaviour on every instance.

These instabilities when computing fractional cut covers are not surprising. Whereas the LP in (4) has exponentially many variables, Carathéodory’s theorem ensures that there exists an optimal solution using at most  $m + 1$  cuts. This result is tight for general linear programming problems, but can be improved for the particular case of approximately solving (4). In [5] we show that  $O(\ln n)$  samples are enough to produce an approximately optimal fractional cut cover with high probability. This is asymptotically best possible. Let  $z \in \mathbb{R}_+^E$  be a positive vector, and let  $\mathcal{F} \subseteq \mathcal{P}(V)$  be a

---

<sup>1</sup> $\|w\|_1 = 1$  actually turned out to be crucial, as SCS otherwise reports the SDP (2b) as infeasible for the instances `bier127`, `brg180`, `d198`, `gr137`, `gr202`, and `gr96`. See the discussion around (42) in Section 6.

subset of shores. An elementary argument shows that

$$(29) \quad |\mathcal{F}| \geq \lceil \log_2(\chi(G)) \rceil \text{ if there exists } y \in \mathbb{R}_+^{\mathcal{F}} \text{ such that } \sum_{S \in \mathcal{F}} y_S \mathbb{1}_{\delta(S)} \geq z,$$

where  $\chi(G)$  denotes the *chromatic number* of  $G$  [18]. Hence, for any family of graphs whose chromatic number is  $\Omega(n)$ , one needs at least  $\log_2 n$  cuts.

Let  $G = (V, E)$  be a graph. Whereas these considerations justify sampling  $\lceil C \ln m \rceil$  cuts for a fixed  $C \in \mathbb{R}_+$ , they say little about how to choose  $C$ . The arguments implying that  $\Theta(\ln m)$  cuts are necessary and sufficient for sampling provide a range of possible choices of  $C$ :

$$(30) \quad 0.7213 \leq C \leq 2.216 \times 10^9.$$

One can adapt the arguments in [5] to show that sampling  $\lceil 2.216 \times 10^9 \ln m \rceil$  shores is enough to compute a  $\beta$ -certificate with  $\beta = 0.99\alpha_{\text{GW}}$  with high probability, i.e., with probability at least  $1 - 1/n$ . The lower bound in (30) follows from (29): as our algorithms must produce feasible solutions on complete graphs, the constant  $C$  must satisfy, for every  $n \in \mathbb{N}$  and  $m := \binom{n}{2}$ ,

$$\lceil C \ln m \rceil = |\mathcal{F}| \geq \lceil \log_2 n \rceil \geq \lceil (2 \ln 2)^{-1} \ln m \rceil.$$

Since  $(2 \ln 2)^{-1} \approx 0.7213$ , if we ignore ceilings for the sake of simplicity we get (30).

In light of (30), Table 1.B is interesting as it shows that, for most of the ?? instances,  $\lceil 128 \ln m \rceil$  cuts are enough to reliably produce good fractional cut covers. Whereas only 3 instances had no **RLI** outcome in Table 1.A, only 2 instances had an **RLI** outcome in some of their runs in Table 1.B, and even then, more often than not, feasibility was attained for these instances. Only **inf-USAir97**, **email-univ**, and **p-hat700-1** did not achieve an average value of  $\beta_{\text{fcc}}$  of at least 0.878, and none of them were off by much. It is crucial to have no **RLI** outcome when trying to tighten the bounds in (30) in practice. Indeed, if we interpret the **RLI** count as estimating the probability of producing a feasible fractional cut cover, we cannot say that  $\lceil 128 \ln m \rceil$  cuts suffice to produce certificates for **ca-netscience**, as with probability 40% we have to sample a new set of  $\lceil 128 \ln m \rceil$  cuts and solve the new restricted LP (15).

**3.2. Working Implementations.** Tables 1.A and 1.B show that  $\lceil 128 \ln m \rceil$  cuts may not be enough to guarantee feasibility of the restricted LP (15) in practice. Considering only the theoretical results from [5], this is not surprising, as 128 is much smaller than the upper bound on  $C$  in (30). On the other hand, Table 1.B hints that  $C := 128$  is almost on target for ?? instances. It is only natural to simply try more samples.

Tables 2.A and 2.B display the results of sampling  $\lceil 256 \ln m \rceil$  cuts. Table Description 1 details how these tables (as well as Tables 0.A to 1.B) were generated. All tables in this text are accompanied by similar descriptions of the experiments they report. When running experiments with  $\lceil 256 \ln m \rceil$  cuts, memory became a bottleneck: in particular, the largest ?? instance **p-hat700-1** could not be solved on ?? hardware, which led us to run this set of experiments on the ?? hardware. As the machine is different from the one where previous experiments were performed, we omit running time information to avoid misleading comparisons.

Unsurprisingly, more samples improved the behavior of our algorithm on both sets of instances. This is less noticeable for the ?? instances (compare Tables 1.B and 2.B), but even there an improvement is noticeable. The two instances that previously had positive **RLI** count now have zero **RLI** count. Instances **inf-USAir97** and **email-univ** are the only instances in Table 2.B where we obtain average values of  $\beta_{\text{fcc}}$  smaller than the theoretical constant  $\alpha_{\text{GW}}$ . More importantly, Table 2.A had no run with an **RLI** outcome, which is quite an improvement over Table 1.A. With feasibility occurring in every run, there was very little variation in the value of  $\beta_{\text{fcc}}$  produced: across both sets of instances the standard deviation among independent runs remains at or below 0.0012, and all the average values of  $\beta_{\text{fcc}}$  are at or above 0.8739, with the majority of them being above the theoretical guarantee of  $\alpha_{\text{GW}} \approx 0.87856$ .

instance	$m$	$1 - \sigma$	$\beta_{mc}$ avg	$\beta_{mc}$ std	RLI	$\beta_{fcc}$ avg	$\beta_{fcc}$ std	time (s)	time std (s)
bayg29	406	.9995	.9934	.0000	9	.8838	–	0.13	–
bays29	406	.9993	.9938	.0000	8	.8819	.0000	0.10	< .01
berlin52	1 326	.9991	.9896	.0000	0	.8904	.0000	0.68	0.11
brazil58	1 653	.9981	.9993	.0000	10	–	–	–	–
gr96	4 560	.9736	.9982	.0000	10	–	–	–	–
kroA100	4 950	.0414	.9979	.0000	10	–	–	–	–
eil101	5 050	.9413	.9809	.0000	8	.8780	.0000	3.23	0.04
gr120	7 140	.9982	.9988	.0000	10	–	–	–	–
bier127	8 001	.9978	.9730	.0000	0	.8772	.0003	9.97	0.90
ch130	8 385	.9963	.9867	.0000	10	–	–	–	–
gr137	9 316	.9252	.9996	.0000	10	–	–	–	–
ch150	11 175	.9954	.9858	.0000	10	–	–	–	–
brg180	16 110	.9983	.9414	.0012	( $\perp_Y$ )	–	–	–	–
d198	19 503	.9981	.9975	.0000	10	–	–	–	–
gr202	20 301	.9983	.9768	.0000	0	.8769	.0001	71.77	6.39
a280	39 060	.9916	.9968	.0000	10	–	–	–	–

TABLE 0.A. ?? instances with  $\lceil 128 \ln m \rceil$  cuts sampled and  $\varepsilon = 0$ , obtained on ?? hardware. See Table Description 1.

**Table Description 1** (Tables 0.A to 3.B). As input, we are given maximum cut instances  $(G, w)$ . We nearly solve the SDPs in (2) using SCS with (27), obtaining  $\tilde{x} \in \mathbb{R}^V$  and  $\tilde{Y} \in \mathbb{S}_+^V$ . We set  $Y := (1 - \varepsilon)\tilde{Y} + \varepsilon I$  as in (28), and  $z := \frac{1}{4}\mathcal{L}_G^*(Y)$  as in STEP 1, where  $\varepsilon \in \{0, 10^{-4}, 1/64\}$ . Set  $(\rho, x, B) := \text{SLACKSANITIZE}(\tilde{x}, w)$  and  $(\mu, R \in \mathbb{R}^{\lfloor k \rfloor \times V}) := \text{REPRESENTATIONSANITIZE}(Y, z)$ . For  $C \in \{128, 256\}$ , we sample  $\lceil C \ln m \rceil$  vectors  $g \in \mathbb{R}^k$  according to the standard multivariate normal distribution, and produce a shore  $S := \{i \in V : g^\top R e_i > 0\}$  for each  $g$ . Let  $\mathcal{F} \subseteq \mathcal{P}(V)$  be the set of shores generated. We compute  $\text{mc}(\mathcal{F}, w)$  directly, and we solve  $\text{fcc}(\mathcal{F}, z)$  using Gurobi. We then define  $\sigma$ ,  $\beta_{mc}$ , and  $\beta_{fcc}$  as in (18) for Tables 0.A to 2.B, and as in (32) (with  $\sigma_\varepsilon$  in place of  $\sigma$ ) for Tables 3.A and 3.B. The whole process is run 10 times with different random generator seeds. Statistics for  $\sigma$  (or  $\sigma_\varepsilon$ ) and  $\beta_{mc}$  use all 10 runs; those for  $\beta_{fcc}$  and running time use only runs that produced a feasible fractional cut cover for  $z$ , thus excluding the infeasible runs counted in the RLI column. For values of  $\varepsilon$  with feasible fractional cut covers in all 10 runs (i.e., zero RLI count), this column is omitted.

Tables 0.A to 3.B report results obtained on ?? hardware with  $C = 128$ , except for Tables 2.A and 2.B which use ?? hardware with  $C = 256$ . Tables 0.A, 1.A, 2.A, and 3.A use ?? instances, whereas Tables 1.B, 2.B, and 3.B use ?? instances. Table 0.A uses  $\varepsilon = 0$ , while Tables 1.A, 1.B, 2.A, and 2.B use  $\varepsilon = 10^{-4}$ . Tables 3.A and 3.B combine the  $\varepsilon = 10^{-4}$  results from Tables 1.A and 1.B with those for  $\varepsilon = 1/64$ , respectively.

instance	$m$	$1 - \sigma$	$\beta_{mc}$ avg	$\beta_{mc}$ std	RLI	$\beta_{fcc}$ avg	$\beta_{fcc}$ std	time (s)	time std (s)
bayg29	406	.9995	.9934	.0000	1	.8838	.0000	0.10	0.01
bays29	406	.9993	.9938	.0000	2	.8837	.0012	0.11	0.01
berlin52	1 326	.9991	.9896	.0000	0	.8903	.0001	0.59	0.09
brazil158	1 653	.9995	.9993	.0000	3	.8888	.0008	0.99	0.19
gr96	4 560	.9994	.9982	.0000	6	.8793	.0000	3.16	0.43
kroA100	4 950	.9977	.9979	.0000	10	–	–	–	–
eil101	5 050	.9992	.9809	.0000	1	.8782	.0000	3.38	0.47
gr120	7 140	.9982	.9988	.0000	5	.9463	.0000	7.23	0.83
bier127	8 001	.9977	.9731	.0000	0	.8773	.0002	9.58	0.66
ch130	8 385	.9982	.9867	.0000	3	.8772	.0000	7.72	0.51
gr137	9 316	.9991	.9996	.0000	9	.9690	–	10.62	–
ch150	11 175	.9977	.9858	.0000	1	.8766	.0002	11.01	1.22
brg180	16 110	.9983	.9421	.0011	5	.8851	.0010	24.19	3.70
d198	19 503	.9981	.9975	.0000	6	.8780	.0003	45.48	4.43
gr202	20 301	.9982	.9768	.0000	0	.8770	.0001	67.70	5.46
a280	39 060	.9953	.9969	.0000	4	.8741	.0003	98.04	4.27

TABLE 1.A. ?? maxcut instances with  $\lceil 128 \ln m \rceil$  cuts sampled and  $\varepsilon = 10^{-4}$ , obtained on ?? hardware. See Table Description 1.

instance	$m$	$1 - \sigma$	$\beta_{mc}$ avg	$\beta_{mc}$ std	RLI	$\beta_{fcc}$ avg	$\beta_{fcc}$ std	time (s)	time std (s)
ENZYMES8	133	.9983	.9785	.0000	0	.9149	.0045	3.30	0.41
soc-dolphins	159	.9979	.9730	.0000	0	.8900	.0000	1.11	0.09
road-chesapeake	170	.9991	.9673	.0000	0	.8880	.0004	0.24	0.01
email-enron-only	623	.9986	.9595	.0023	0	.8882	.0000	30.75	4.42
dwt_209	767	.9977	.9658	.0013	0	.8924	.0000	82.73	10.71
ca-netscience	914	.9963	.9623	.0022	4	.8856	.0000	446.92	32.04
Erdos991	1 417	.9910	.9346	.0024	0	.8796	.0005	1 450.67	65.65
hamming6-2	1 824	.9997	.9997	.0000	0	.9997	.0000	4.42	0.93
inf-USAir97	2 126	.9916	.9668	.0005	0	.8736	.0003	598.55	35.23
ia-infect-hyper	2 196	.9990	.9687	.0010	0	.8898	.0004	16.05	2.12
DD687	2 600	.9979	.9407	.0015	0	.8800	.0005	4 395.64	164.04
dwt_503	2 762	.9984	.9845	.0008	1	.8859	.0000	2 605.63	58.90
ia-infect-dublin	2 765	.9973	.9492	.0009	0	.8805	.0002	852.30	34.11
email-univ	5 451	.9966	.9260	.0012	0	.8690	.0003	29 562.98	2 166.08
johnson16-2-4	5 460	.9997	.9715	.0013	0	.9325	.0006	15.95	2.85
p-hat700-1	60 999	.9985	.9704	.0005	0	.8696	.0002	3 031.69	56.45

TABLE 1.B. ?? maxcut instances with  $\lceil 128 \ln m \rceil$  cuts sampled and  $\varepsilon = 10^{-4}$ , obtained on ?? hardware. See Table Description 1.

From a practical point of view, Tables 2.A and 2.B allow us to consider  $\lceil 256 \ln m \rceil$  as an upper bound for the number of cuts we need to sample in this work, significantly narrowing the range in (30). At least for the set of instances presented here, which were chosen in a previous context for a computational study of the maximum cut problem solely, we have no need to do experiments with  $\lceil 2 \times 10^9 \ln m \rceil$  cuts to obtain reasonable  $\beta$ -certificates. We further wish to avoid resorting to ?? hardware, as results obtained on a machine with 768 GB of RAM have little bearing on what can be accomplished on affordable and readily available hardware. For these reasons, the remaining experiments focus on sampling  $\lceil 128 \ln m \rceil$  cuts.

We conclude this section by presenting a first practical and robust implementation of our primal-dual algorithm. Our experiments are the same as the ones which generated Tables 1.A and 1.B, except that we revisit our choice of perturbation, changing the perturbation parameter  $\varepsilon = 10^{-4}$  in (28) to a much higher  $\varepsilon = 1/64 = 0.015625$ . In particular, we sample  $\lceil 128 \ln m \rceil$  cuts via the random hyperplane technique, and run the experiments in ?? hardware. Tables 3.A and 3.B present the results of this set of experiments. Before discussing the results, we highlight an important aspect of these experiments, that will appear again in this paper. Let  $\tilde{Y} \in \mathbb{S}^V$  be the nearly optimal solution to the SDP (2a) computed by SCS. We call Algorithm 2 with inputs

$$Y_\varepsilon := (1 - \varepsilon)\tilde{Y} + \varepsilon I \quad \text{and} \quad z_\varepsilon := \frac{1}{4}\mathcal{L}_G^*(Y_\varepsilon) = \frac{1-\varepsilon}{4}\mathcal{L}_G^*(\tilde{Y}) + \frac{\varepsilon}{2}\mathbb{1}$$

to sanitize  $Y_\varepsilon$  before proceeding to the rest of our pipeline. This is qualitatively different from what is shown in Tables 1.A and 1.B: by picking  $\varepsilon$  many orders of magnitudes larger than solver accuracy, we are deliberately deviating from the nearly optimal solutions. Thus, even though we started from the same initial  $w \in \mathbb{R}_+^E$  and solved the same SDP relaxation (2), the change of the perturbation parameter  $\varepsilon$  now significantly changes the final  $\beta$ -certificate. In this way, the quality measures we introduced in (18) are all functions of  $\varepsilon \in [0, 1]$ . More precisely, let  $\mathcal{F}_\varepsilon$  be the set of cuts obtained from sampling from  $\text{GW}(Y_\varepsilon)$ , and let  $\mu_\varepsilon$  be the sanitized objective value produced by Algorithm 2. Our quality measures are

$$(31) \quad \sigma_\varepsilon := 1 - \frac{z_\varepsilon^\top w}{\mu_\varepsilon \rho}, \quad \beta_{\text{fcc}} := \frac{(1 - \sigma_\varepsilon)\mu_\varepsilon}{\text{fcc}(\mathcal{F}_\varepsilon, z_\varepsilon)}, \quad \beta_{\text{mc}} := \frac{\text{mc}(\mathcal{F}_\varepsilon, w)}{\rho}; \quad \text{see (18)}.$$

This has an important consequence on the interpretation of Tables 3.A and 3.B: the columns reflect the quality of certification obtained for different pairs  $(w, z_\varepsilon)$ . In particular, whereas distinct values of  $\beta_{\text{mc}}$  really reflect whether one obtained better cuts for the given input  $w \in \mathbb{R}_+^E$  by using  $\mathcal{F}_{10^{-4}}$  or  $\mathcal{F}_{1/64}$ , the distinct values of  $\beta_{\text{fcc}}$  are incomparable, as they correspond to covering distinct vectors  $z_{10^{-4}}$  and  $z_{1/64}$ . In this way, we are taking advantage of the theoretical framework presented in Section 2.1: when given  $w \in \mathbb{R}_+^E$ , we have freedom to select  $z \in \mathbb{R}_+^E$  and to certify  $(w, z) \in H_\sigma(G)$ . By selecting distinct values of  $\varepsilon$ , we attempt to obtain vectors  $z_\varepsilon$  which are easier to cover by repeated sampling, at the cost of having worse pairing quality (i.e., with higher values of  $\sigma$ ). This trade-off, which is most explicit in the definition of  $\beta_{\text{fcc}}$  in (32), will be addressed in Section 5.1.

We now discuss Tables 3.A and 3.B. In none of our instances the perturbation parameter  $\varepsilon = 1/64$  had a significant impact on  $\beta_{\text{mc}}$ : Table 3.A only had improvements in this constant, and Table 3.B had no loss greater than 1%. On the other hand, the improvements on the outcome of solving the restricted LP (15) in Table 3.A is remarkable, with no runs resulting in RLI when  $\varepsilon = 1/64$ . Comparing back to Table 2.A, we see that by increasing the perturbation, we are able to halve the number of samples necessary without losing approximation quality or stability of behavior. Table 3.B illustrates how increasing perturbation decreases the probability of RLI, at the cost of decreasing the value of  $\beta_{\text{fcc}}$ .

Tables 4.A and 4.B display both the average total running time and the average time spent solving the LP for the experiments reported in Tables 1.A, 1.B, 3.A, and 3.B. Since, for each input graph, the SDP being solved is always the same in all of these tables, there is no relevant difference on the running time of the SDP solver, so we highlight only the time for computing the LP (15)

and the total running time. Figures 2 and 3 dissect the running time of our algorithms when  $\varepsilon = 1/64$ , breaking down the running time spent on each of the most significant parts of the pipeline: solving the SDP (2) for input  $(G, w)$ , generating the shores  $\mathcal{F}$ , and solving  $\text{fcc}(\mathcal{F}, z)$ . For each instance, we computed the average across the 10 independent runs of the running time spent on each of these steps, divided them by the average total running time for those instances, and plotted the results. For most instances, the extra work involved in computing the  $\beta$ -certificates does not change the order of magnitude of overall time spent, as solving the SDP remains the main bottleneck. Two of the exceptions are the highly symmetric instances `hamming6-2` and `johnson16-2-4`, where indeed solving the LP completely dominates the running time of the overall procedure. On the other exceptions, `bier127` and `brg180`, solving the SDP still takes between 10% and 20% of the total running time of the procedure, which implies that the whole procedure is between 5 to 10 times slower than just solving the SDP (2).

Before moving on, we make a couple of remarks comparing our experiments to the recent work of Mirka and Williamson [31]. As [31] is only interested in approximately solving maximum cut instances, it is natural to run a *local improvement heuristic* on every cut produced by the random hyperplane technique. Concretely, given a shore  $S$  and a vertex  $i \in V$ , one can always check which of  $S \cup \{i\}$  or  $S \setminus \{i\}$  has better objective value. One can then repeat this local improvement until obtaining a shore where no such exchange improves the objective value. When it comes to using the cuts obtained to cover  $z \in \mathbb{R}_+^E$ , the benefit of the heuristic is not clear. If we add every intermediate cut as a column to our LP, we are significantly increasing the number of cuts considered and thus the running time of our covering algorithm; such cuts, however, are of limited utility, as consecutive cuts are extremely similar to each other. On the other hand, adding only the last produced cut means we take much longer to generate each cut, while also restricting the cuts considered to a subset that is not necessarily better suited for the covering problem. This leads to a noteworthy difference between our works: whereas [31] always sample 10 cuts and run the local improvement heuristic on every cut, we sample  $\lceil C \ln m \rceil$  cuts and do not employ the heuristic.

In part due to this difference, and in part due to the different hardware used, directly comparing running time of our experiments with the ones in [31] is of limited value. It is important to point out, however, that the running time we have obtained in our experiments using ?? hardware is on the same ballpark as the running times they reported. In general, the running times in Tables 1.A, 1.B, 4.A, and 4.B are at most 3 times slower than what was reported in [31]. There are four exceptions: `ENZYMES8`, `email-univ`, `johnson16-2-4`, and `hamming6-2`. `ENZYMES8` is the smallest instance among the sparse ones, but took 24.41s to solve in [31], which is almost 6.8 times slower than what we report in Table 4.B. On the other end, `email-univ` is the slowest instance to solve both for us and for [31], and took 5 times longer in ?? hardware than what is reported by Mirka and Williamson. For this instance, Figure 3 shows almost all of this time was spent solving the SDP relaxation. Our experiments are 7.9 times slower for `hamming6-2` and 37 times slower for `johnson16-2-4`. Table 4.B pinpoints this slowdown to solving the LP: when solving `johnson16-2-4` we spend, on average, 1.07 seconds out of 15.16 seconds on everything but solving the LP (15), whereas [31] reports 0.408 seconds as the total time spent on this instance.

instance	$m$	$1 - \sigma$	$\beta_{mc}$ avg	$\beta_{mc}$ std	$\beta_{fcc}$ avg	$\beta_{fcc}$ std
bayg29	406	.9995	.9934	.0000	.8839	.0000
bays29	406	.9993	.9938	.0000	.8840	.0012
berlin52	1 326	.9991	.9896	.0000	.8904	.0000
brazil58	1 653	.9995	.9993	.0000	.8897	.0009
gr96	4 560	.9994	.9982	.0000	.8794	.0000
kroA100	4 950	.9977	.9979	.0000	.9956	.0001
eil101	5 050	.9992	.9809	.0000	.8782	.0000
gr120	7 140	.9982	.9988	.0000	.9463	.0000
bier127	8 001	.9977	.9732	.0001	.8774	.0002
ch130	8 385	.9982	.9867	.0000	.8773	.0000
gr137	9 316	.9991	.9996	.0000	.9690	.0000
ch150	11 175	.9977	.9858	.0000	.8768	.0000
brg180	16 110	.9983	.9422	.0010	.8869	.0001
d198	19 503	.9981	.9975	.0000	.8790	.0006
gr202	20 301	.9982	.9768	.0000	.8776	.0000
a280	39 060	.9953	.9969	.0000	.8749	.0000

TABLE 2.A. ?? maxcut instances with  $\lceil 256 \ln m \rceil$  cuts sampled and  $\varepsilon = 10^{-4}$ , obtained on ?? hardware. See Table Description 1.

instance	$m$	$n$	$1 - \sigma$	$\beta_{mc}$ avg	$\beta_{mc}$ std	$\beta_{fcc}$ avg	$\beta_{fcc}$ std
ENZYMES8	133	88	.9983	.9785	.0000	.9161	.0008
soc-dolphins	159	62	.9979	.9730	.0000	.8900	.0000
road-chesapeake	170	39	.9991	.9673	.0000	.8883	.0000
email-enron-only	623	143	.9986	.9611	.0016	.8882	.0000
dwt_209	767	209	.9977	.9686	.0019	.8924	.0000
ca-netscience	914	379	.9963	.9643	.0017	.8856	.0000
Erdos991	1 417	492	.9910	.9383	.0026	.8810	.0000
hamming6-2	1 824	64	.9997	.9997	.0000	.9997	.0000
inf-USAir97	2 126	332	.9916	.9670	.0006	.8739	.0000
ia-infect-hyper	2 196	113	.9990	.9702	.0008	.8940	.0000
DD687	2 600	725	.9979	.9419	.0011	.8864	.0001
dwt_503	2 762	503	.9984	.9854	.0005	.8859	.0000
ia-infect-dublin	2 765	410	.9973	.9499	.0009	.8866	.0000
email-univ	5 451	1 133	.9966	.9269	.0009	.8777	.0002
johnson16-2-4	5 460	120	.9997	.9722	.0004	.9459	.0002
p-hat700-1	60 999	700	.9985	.9704	.0003	.8786	.0001

TABLE 2.B. ?? maxcut instances with  $\lceil 256 \ln m \rceil$  cuts sampled and  $\varepsilon = 10^{-4}$ , obtained on ?? hardware. See Table Description 1.

$\varepsilon$	$10^{-4}$	$1/64$	$10^{-4}$	$1/64$	$10^{-4}$	$10^{-4}$	$1/64$	$10^{-4}$	$1/64$
instance	$1 - \sigma_\varepsilon$	$1 - \sigma_\varepsilon$	$\beta_{mc}$	$\beta_{mc}$	RLI	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$ std	$\beta_{fcc}$ std
bayg29	.9995	.9959	.9934	.9934	1	.8838	.8864	.0000	.0000
bays29	.9993	.9957	.9938	.9938	2	.8837	.8873	.0012	.0000
berlin52	.9991	.9960	.9896	.9896	0	.8903	.8922	.0001	.0000
brazil58	.9995	.9958	.9993	.9993	3	.8888	.8920	.0008	.0000
gr96	.9994	.9956	.9982	.9982	6	.8793	.8817	.0000	.0000
kroA100	.9977	.9934	.9979	.9979	10	–	.9902	–	.0001
eil101	.9992	.9959	.9809	.9809	1	.8782	.8809	.0000	.0000
gr120	.9982	.9939	.9988	.9988	5	.9463	.9465	.0000	.0000
bier127	.9977	.9950	.9731	.9731	0	.8773	.8802	.0002	.0004
ch130	.9982	.9948	.9867	.9867	3	.8772	.8799	.0000	.0000
gr137	.9991	.9947	.9996	.9996	9	.9690	.9689	–	.0000
ch150	.9977	.9944	.9858	.9858	1	.8766	.8795	.0002	.0000
brg180	.9983	.9958	.9421	.9434	5	.8851	.8868	.0010	.0005
d198	.9981	.9938	.9975	.9975	6	.8780	.8868	.0003	.0003
gr202	.9982	.9955	.9768	.9768	0	.8770	.8782	.0001	.0002
a280	.9953	.9914	.9969	.9969	4	.8741	.8765	.0003	.0001

TABLE 3.A. ?? instances with  $[128 \ln m]$  cuts sampled and  $\varepsilon = 1/64$ , compared against from Table 1.A, which used  $\varepsilon = 10^{-4}$ . The largest standard deviation of  $\beta_{mc}$  observed was .0011. See Table Description 1.

$\varepsilon$	$10^{-4}$	$1/64$	$10^{-4}$	$1/64$	$10^{-4}$	$10^{-4}$	$1/64$	$10^{-4}$	$1/64$
instance	$1 - \sigma_\varepsilon$	$1 - \sigma_\varepsilon$	$\beta_{mc}$	$\beta_{mc}$	RLI	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$ std	$\beta_{fcc}$ std
ENZYME88	.9983	.9908	.9785	.9785	0	.9149	.9101	.0045	.0076
soc-dolphins	.9979	.9922	.9730	.9730	0	.8900	.8895	.0000	.0000
road-chesapeake	.9991	.9938	.9673	.9673	0	.8880	.8881	.0004	.0000
email-enron-only	.9986	.9940	.9595	.9586	0	.8882	.8887	.0000	.0000
dwt_209	.9977	.9927	.9658	.9654	0	.8924	.8924	.0000	.0000
ca-netscience	.9963	.9917	.9623	.9600	4	.8856	.8861	.0000	.0000
Erdos991	.9910	.9858	.9346	.9323	0	.8796	.8753	.0005	.0002
hamming6-2	.9997	.9985	.9997	.9997	0	.9997	.9982	.0000	.0000
inf-USAir97	.9916	.9869	.9668	.9665	0	.8736	.8744	.0003	.0003
ia-infect-hyper	.9990	.9965	.9687	.9692	0	.8898	.8899	.0004	.0003
DD687	.9979	.9931	.9407	.9376	0	.8800	.8758	.0005	.0004
dwt_503	.9984	.9939	.9845	.9841	1	.8859	.8880	.0000	.0000
ia-infect-dublin	.9973	.9935	.9492	.9471	0	.8805	.8778	.0002	.0006
email-univ	.9966	.9918	.9260	.9230	0	.8690	.8650	.0003	.0003
johnson16-2-4	.9997	.9978	.9715	.9699	0	.9325	.9308	.0006	.0003
p-hat700-1	.9985	.9968	.9704	.9699	0	.8696	.8685	.0002	.0003

TABLE 3.B. ?? instances with  $[128 \ln m]$  cuts sampled and  $\varepsilon = 1/64$ , compared against Table 1.B, which used  $\varepsilon = 10^{-4}$ . The largest standard deviation of  $\beta_{mc}$  observed was 0.0029. See Table Description 1.

$\varepsilon$	$10^{-4}$	$10^{-4}$	$1/64$	$10^{-4}$	$1/64$
instance	RLI	LP (s)	LP (s)	Total (s)	Total (s)
bayg29	1	0.02	0.03	0.10	0.11
bays29	2	0.02	0.03	0.11	0.12
berlin52	0	0.18	0.32	0.59	0.69
brazil158	3	0.12	0.27	0.99	1.03
gr96	6	0.68	1.45	3.16	3.83
kroA100	10	–	3.94	–	6.32
eil101	1	0.87	2.32	3.38	5.00
gr120	5	1.39	1.71	7.23	7.36
bier127	0	5.45	18.72	9.58	22.94
ch130	3	1.85	3.90	7.72	9.00
gr137	9	1.44	2.99	10.62	10.61
ch150	1	2.55	6.97	11.01	15.40
brg180	5	18.12	58.49	24.19	65.00
d198	6	4.39	10.48	45.48	53.48
gr202	0	31.06	61.74	67.70	97.70
a280	4	17.67	38.20	98.04	127.56

TABLE 4.A. Running time information for Tables 1.A and 3.A. (RLI) for  $\varepsilon = 1/64$  is omitted as it was always zero. Columns “LP (s)” and “Total (s)” display the average running time for runs for which  $\text{fcc}(\mathcal{F}, z)$  is finite.




FIGURE 2. Running time breakdown of experiments with  $\varepsilon = 1/64$  in Table 3.A.

$\varepsilon$	$10^{-4}$	$10^{-4}$	$1/64$	$10^{-4}$	$1/64$
instance	RLI	LP (s)	LP (s)	Total (s)	Total (s)
ENZYMES8	0	0.02	0.04	3.30	3.56
soc-dolphins	0	0.03	0.05	1.11	1.05
road-chesapeake	0	0.01	0.02	0.24	0.25
email-enron-only	0	0.89	0.83	30.75	32.93
dwt_209	0	1.40	1.42	82.73	78.08
ca-netscience	4	1.37	2.24	446.92	468.37
Erdos991	0	6.93	6.83	1 450.67	1 460.16
hamming6-2	0	4.15	9.23	4.42	9.49
inf-USAir97	0	3.07	4.61	598.55	623.08
ia-infect-hyper	0	5.98	5.55	16.05	15.17
DD687	0	9.08	9.20	4 395.64	4 445.81
dwt_503	1	8.03	14.95	2 605.63	2 628.42
ia-infect-dublin	0	8.55	9.29	852.30	853.91
email-univ	0	22.74	21.17	29 562.98	30 590.25
johnson16-2-4	0	14.87	14.09	15.95	15.16
p-hat700-1	0	326.55	325.06	3 031.69	3 027.89

TABLE 4.B. Running time information for Tables 1.B and 3.B. (RLI) for  $\varepsilon = 1/64$  is omitted as it was always zero. Columns “LP (s)” and “Total (s)” display the average running time for runs for which  $\text{fcc}(\mathcal{F}, z)$  is finite.




FIGURE 3. Running time breakdown of experiments with  $\varepsilon = 1/64$  in Table 3.B.

## 4. FRACTIONAL CUT-COVERING INSTANCES

Section 3 presents experiments which compute  $\beta$ -certificates given maximum cut instances  $(G, w)$  as input, where  $G = (V, E)$  is a graph and  $w \in \mathbb{R}_+^E$ . In this section, we study the ?? of our algorithms when given fractional cut-covering instances  $(G, z)$  as input, with  $z \in \mathbb{R}_+^E$ .

Let  $G = (V, E)$  be a graph. Let  $\varepsilon \in (0, 1)$ , and set  $\gamma_G := 1 + n + m$ . We formulate (12) as

$$(32) \quad \begin{array}{ll} \text{Minimize} & \gamma_G \cdot \mu, \\ \text{subject to} & \frac{1}{4} \mathcal{L}_G^*(Y + \mu \varepsilon I) \succeq z, \\ & \text{diag}(Y + \mu \varepsilon I) = \mu \mathbb{1}, \\ & \mu \in \mathbb{R}_+, Y \in \mathbb{S}_+^V \end{array} \quad \begin{array}{ll} \text{Maximize} & z^\top w, \\ \text{subject to} & \frac{\varepsilon}{2} \mathbb{1}^\top w + (1 - \varepsilon) \mathbb{1}^\top x \leq \gamma_G, \\ & -\frac{1}{4} \mathcal{L}_G(w) + \text{Diag}(x) \succeq 0, \\ & w \in \mathbb{R}_+^E, x \in \mathbb{R}^V \end{array}$$

Here,  $\varepsilon$  plays a role analogous to the perturbation parameter defined in (28), which justifies our choice of notation. Except for the scaling factor  $\gamma_G$ , formulation (33) was introduced in [5] as a crucial step to prove correctness of our algorithms in exact arithmetic. The scaling factor  $\gamma_G$  is a minor convenience: it ensures that  $\hat{w} := \mathbb{1}$  and  $\hat{x} := \frac{1}{2} \text{deg} + \mathbb{1}$  form a strictly feasible solution to the dual problem, where  $\text{deg} \in \mathbb{R}^V$  denotes the degree vector of  $G$  (ignoring weights). We explored some alternative formulations before settling on (33); see Section B.

The stark contrast between the literature on the maximum cut problem and the literature on the fractional cut-covering problem becomes apparent when considering which fractional cut-covering instances to solve. Whereas we were able to use instances from a recent work on the maximum cut problem in Section 3, we believe this manuscript to be the first one producing fractional cut covers for weighted instances  $(G, z)$ . To the best of our knowledge, the only experiments approximating the fractional cut-covering number appear in [33], which compares distinct bounds for  $\text{fcc}(G) = \text{fcc}(G, \mathbb{1})$  on a set of graphs, without actually obtaining a fractional cut cover. In particular, there was no established benchmark data set for us to experiment with.

A first choice is to use the same instances as Section 3, but to reinterpret its weights as defining a fractional cut-covering instance rather than a maximum cut instance. We present in Tables 5.A and 5.B the results of these experiments. To highlight that we are starting from the fractional cut-covering side, these tables first show information about the quality of the fcc certificates, and then about the maximum cut certificates. Results in Tables 5.A and 5.B are strikingly better than those in Tables 1.A and 1.B. In particular, in every single run of our pipeline we obtained a feasible fractional cut cover. With a comparable running time to Table 1.A, the worst certifiable approximation factor for fractional cut covering in Table 5.A is 89.95%, and most are above 95%. Comparing Tables 1.B and 5.B, we note a significant difference on running time. On average, no instance in Table 5.B took longer than 1,000 seconds, whereas 5 instances took longer than that in Table 1.B.

**Table Description 2** (Tables 5.A and 5.B). Experiments obtained on ?? hardware. Inputs are weighted graphs previously used as maximum cut instances in Tables 0.A to 3.B, but with the weights being interpreted as describing a fractional cut-covering instance  $(G, z)$ . We nearly solve the SDPs in (33) with  $\varepsilon := 10^{-4}$  using SCS (see (61)), thus computing nearly optimal  $\tilde{\mu} \in \mathbb{R}$ ,  $\tilde{Y} \in \mathbb{S}^V$ ,  $\tilde{w} \in \mathbb{R}^E$ , and  $\tilde{x} \in \mathbb{R}^V$ . Set  $Y := \tilde{Y} + \tilde{\mu} \varepsilon I$ , set  $(\rho, x, B) := \text{SLACKSANITIZE}(\tilde{x}, \tilde{w})$ , and set  $(\mu, R \in \mathbb{R}^{[k] \times V}) := \text{REPRESENTATIONSANITIZE}(Y, z)$ . We sample  $\lceil 128 \ln m \rceil$  vectors  $g \in \mathbb{R}^k$  according to the standard multivariate normal distribution, and produce a shore  $S := \{i \in V : g^\top R e_i > 0\}$  for each  $g$ . Let  $\mathcal{F} \subseteq \mathcal{P}(V)$  be the set of shores generated. We compute  $\text{mc}(\mathcal{F}, w)$  directly, and we solve  $\text{fcc}(\mathcal{F}, z)$  using Gurobi. We then define  $\sigma$ ,  $\beta_{\text{mc}}$ , and  $\beta_{\text{fcc}}$  as in (18). The whole process is run 10 times with different random generator seeds. Statistics for  $\sigma$ ,  $\beta_{\text{mc}}$ , and  $\beta_{\text{fcc}}$  use all 10 runs; there were no RLI occurrences. Table 5.A reports ?? instances, and Table 5.B reports ?? instances.

Our duality framework provides an alternative way of generating instances. Central to our work is the idea that the SDP relaxation (2) studied by Goemans and Williamson and the one we have introduced in (12) can produce *pairs* of instances, as well as the necessary information for certifying these objects to be paired via (20). In particular, for every maximum cut instance  $(G, w)$ , by solving the SDP relaxation (2) we compute an associated  $z \in \mathbb{R}_+^E$ , and then certify  $(w, z) \in H_\sigma(G)$ . As a byproduct, we obtain a fractional cut-covering instance  $(G, z)$  paired with  $(G, w)$ . In this way, we can generate what we shall call *paired instances*. We describe in Algorithm 3 how our procedure pairs maximum cut instances  $(G, w)$  with fractional cut-covering instances  $(G, z)$ .

---

**Algorithm 3** Compute an fcc instance  $(G, z)$  paired to an input maxcut instance  $(G, w)$

---

**Parameters:**  $\gamma \in [0, 1)$ .

▷ We have used  $\gamma := 10^{-6}$ .

**Input:** Weighted maximum cut instance  $(G, w)$

**Output:**  $\text{PAIREDINSTANCE}_\gamma(G, w)$  can either fail, returning  $\perp$ , or return an fcc instance  $(G, z)$  for which there exists  $R \in \mathbb{R}^{[k] \times V}$  such that

$$\|\text{diag}(R^\top R) - \mathbb{1}\|_2 \approx 0, z_{ij} \leq \frac{1}{4} \|Re_i - Re_j\|_2^2 \text{ for all } ij \in E, \text{ and } w^\top z \approx \eta(G, w)$$

```

1 procedure PAIREDINSTANCE $_\gamma(G, w)$ 
2   Nearly solve (2) with precision  $10^{-12}$ , obtaining  $\tilde{x} \in \mathbb{R}^V$  and  $\tilde{Y} \in \mathbb{S}_+^V$            ▷ See Section 6
3   Set  $\tilde{z} \leftarrow \frac{1}{4} \mathcal{L}_G^*(\tilde{Y})$ 
4   if SLACKSANITIZE $_{10^{-8}}(\tilde{x}, w) = \perp$  or REPRESENTATIONSANITIZE $_0(\tilde{Y}, \tilde{z}) = \perp$  then
5     return  $\perp$ 
6   else if  $\mu \neq 1$  for  $(\mu, R) \leftarrow \text{REPRESENTATIONSANITIZE}_0(\tilde{Y}, \tilde{z})$  then
7     return  $\perp$ 
8   else
9     for each  $ij \in E$ , set  $z_{ij} \leftarrow \begin{cases} \tilde{z}_{ij}, & \text{if } \tilde{z}_{ij} > \gamma \|\tilde{z}\|_\infty \\ 0, & \text{otherwise} \end{cases}$ 
10    return  $(G, z)$ 

```

---

To obtain high precision, we solved the SDP relaxation in (2) using DDS, an implementation of an interior-point method producing solutions with precision far beyond what SCS can achieve. We postpone the discussion of how we implemented (2) in DDS to Section 6, which is devoted to the impact of the choice of solvers. Tables 6.A and 6.B show the result of running our algorithms on these paired instances. These are much more challenging for our fractional cut-covering algorithm, with only 7 out of the 16 instances producing feasible covers in all 10 independent runs in Table 6.A. In Tables 5.B and 6.B, running time distinguishes easier from harder instances. Indeed, 6 instances in Table 6.B had average running time above 1,000 seconds, whereas no instance in Table 5.B took that long on average. Due to these discrepancies, our subsequent experiments are performed on paired instances.

Analogous to Section 3, perturbation can really improve the behavior of our algorithms. In Tables 7.A and 7.B we compare the results obtained with  $\varepsilon = 1/64$  and  $\varepsilon = 10^{-4}$ , as in Tables 3.A and 3.B. Setting  $\varepsilon = 10^{-4}$  proved insufficient even to reliably reach feasibility: for some instances we observed a 50% probability of failing to produce a fractional cut cover after  $\lceil 128 \ln m \rceil$  samples. With the higher value of  $\varepsilon = 1/64$ , on the other hand, we reliably produced feasible fractional cut covers, with almost no cost on the average value of  $\beta_{\text{fcc}}$  on the runs producing feasible solutions. However, there is one interesting difference between the effect of  $\varepsilon$  in this section and in Section 3. Note that  $\varepsilon$  does not appear in the formulations in (2): there we always solve the same SDP relaxation, and then from the output  $\tilde{Y}$  of the solver we compute a convex combination  $(1 - \varepsilon)\tilde{Y} + \varepsilon I$  with

the identity matrix as in (28). Here, as one can see in (33), the perturbation parameter  $\varepsilon$  appears directly in the relaxation, changing the SDP problem we are solving. In particular, for different choices of  $\varepsilon \in [0, 1]$ , only  $z \in \mathbb{R}_+^E$  is kept fixed, and all the other quality measures change, with

$$(33) \quad \sigma_\varepsilon := 1 - \frac{z^\top w_\varepsilon}{\mu_\varepsilon \rho_\varepsilon}, \quad \beta_{\text{fcc}} := \frac{(1 - \sigma_\varepsilon) \mu_\varepsilon}{\text{fcc}(\mathcal{F}_\varepsilon, z)}, \quad \beta_{\text{mc}} := \frac{\text{mc}(\mathcal{F}_\varepsilon, w_\varepsilon)}{\rho_\varepsilon}; \quad \text{see (18)}.$$

Here,  $\mathcal{F}_\varepsilon \subseteq \mathcal{P}(V)$  is the set of shores obtained from sampling  $\text{GW}(Y_\varepsilon)$  with  $Y_\varepsilon := Y + \mu\varepsilon I$ , and  $\mu_\varepsilon$  is the sanitized objective value produced by Algorithm 2 on input  $(Y_\varepsilon, z)$ . Similar to Tables 3.A and 3.B, we reliably obtain feasible cut covers with certifiably good  $\beta_{\text{fcc}}$  values when  $\varepsilon = 1/64$ . However, here the choice of  $\varepsilon$  may directly impact the running time of the algorithm, as it changes the SDP being solved. There are some instances for which running the algorithm with  $\varepsilon = 1/64$  was significantly faster than with  $\varepsilon = 10^{-4}$  (see `inf-USAir97`, `DD687`, `ia-infect-dublin`), and one instance where it was significantly slower (see `p-hat700-1`).

**Table Description 3** (Tables 6.A to 7.B). Experiments obtained on ?? hardware. As input, we are given fractional cut-covering instances  $(G, z)$ . In all cases they are **paired** instances, obtained by solving the corresponding maximum cut instance via Algorithm 3; see Section 6.1. In Tables 6.A and 7.A the paired instances were obtained from ?? instances, and in Tables 6.B and 7.B from ?? instances. We nearly solve the SDPs in (33) with  $\varepsilon \in \{10^{-4}, 1/64\}$  using SCS via (61), thus computing nearly optimal  $\tilde{\mu} \in \mathbb{R}$ ,  $\tilde{Y} \in \mathbb{S}^V$ ,  $\tilde{w} \in \mathbb{R}^E$ , and  $\tilde{x} \in \mathbb{R}^V$ . Set  $Y := \tilde{Y} + \varepsilon \tilde{\mu} I$ , set  $(\rho, x, B) := \text{SLACKSANITIZE}(\tilde{x}, \tilde{w})$ , and set  $(\mu, R \in \mathbb{R}^{[k] \times V}) := \text{REPRESENTATIONSANITIZE}(Y, z)$ . We sample  $\lceil 128 \log m \rceil$  vectors  $g \in \mathbb{R}^k$  according to the standard multivariate normal distribution, and produce a shore  $S := \{i \in V : g^\top R e_i > 0\}$  for each  $g$ . Let  $\mathcal{F} \subseteq \mathcal{P}(V)$  be the set of shores generated. We compute  $\text{mc}(\mathcal{F}, w)$  directly, and we solve  $\text{fcc}(\mathcal{F}, z)$  using Gurobi. We then define  $\sigma_\varepsilon$ ,  $\beta_{\text{mc}}$ , and  $\beta_{\text{fcc}}$  as in (34). The whole process is run 10 times with different random generator seeds. Statistics for  $\sigma_\varepsilon$  and  $\beta_{\text{mc}}$  use all 10 runs; those for  $\beta_{\text{fcc}}$  and running time use only runs that produced a feasible fractional cut cover for  $z$ , thus excluding the infeasible runs counted in the **RLI** column. For values of  $\varepsilon$  with feasible fractional cut covers in all 10 runs (i.e., zero **RLI** count), this column is omitted. For ease of comparison, we repeat the relevant columns from Tables 6.A and 6.B in Tables 7.A and 7.B.

instance	$m$	$1 - \sigma$	$\beta_{fcc}$ avg	$\beta_{fcc}$ std	$\beta_{mc}$ avg	$\beta_{mc}$ std	time (s)	time std (s)
bayg29	406	.9999	.9648	.0000	.9996	.0000	0.16	0.02
bays29	406	.9999	.9707	.0000	.9999	.0000	0.15	0.01
berlin52	1 326	.9989	.9217	.0000	.9995	.0000	0.48	0.04
brazil158	1 653	.9997	.9454	.0000	.9876	.0000	0.61	0.06
gr96	4 560	.9994	.9231	.0000	.9995	.0000	3.00	0.40
kroA100	4 950	.9999	.9874	.0000	.9999	.0000	4.39	0.53
eil101	5 050	.9998	.9691	.0000	.9999	.0000	3.39	0.47
gr120	7 140	.9999	.9681	.0000	.9999	.0000	5.65	0.45
bier127	8 001	.9999	.9629	.0000	.9999	.0000	6.68	0.85
ch130	8 385	.9996	.9636	.0000	.9997	.0000	6.27	0.37
gr137	9 316	.9998	.9594	.0000	.9978	.0000	7.64	1.14
ch150	11 175	.9994	.9697	.0000	.9995	.0000	10.66	0.85
brg180	16 110	.9644	.9533	.0007	.9953	.0009	59.13	5.25
d198	19 503	.9999	.9504	.0000	.9996	.0000	28.55	2.44
gr202	20 301	.9998	.8995	.0000	.9830	.0000	21.30	1.64
a280	39 060	.9997	.9606	.0000	.9998	.0000	62.24	6.07

TABLE 5.A. ?? fcc instances with  $\lceil 128 \ln m \rceil$  cuts sampled and  $\varepsilon = 10^{-4}$ , obtained on ?? hardware. See Table Description 2.

instance	$m$	$1 - \sigma$	$\beta_{fcc}$ avg	$\beta_{fcc}$ std	$\beta_{mc}$ avg	$\beta_{mc}$ std	time (s)	time std (s)
ENZYMES8	133	.9999	.8888	.0000	.8888	.0000	0.52	0.06
soc-dolphins	159	.9998	.9598	.0000	.9598	.0000	0.22	0.03
road-chesapeake	170	.9990	.9591	.0000	.9596	.0000	0.15	0.01
email-enron-only	623	.9996	.9805	.0009	.9997	.0000	2.59	0.25
dwt_209	767	.9980	.9190	.0007	.9982	.0000	9.63	1.09
ca-netscience	914	.9999	.9646	.0013	.9875	.0000	42.44	5.84
Erdos991	1 417	.9994	.9534	.0009	.9790	.0000	93.50	3.35
hamming6-2	1 824	.9998	.9681	.0006	.9999	.0000	7.03	1.13
inf-USAir97	2 126	.9997	.9280	.0000	.9957	.0000	61.21	8.50
ia-infect-hyper	2 196	.9974	.9602	.0005	.9966	.0002	6.56	0.70
DD687	2 600	.9999	.9275	.0010	.9999	.0000	177.78	3.96
dwt_503	2 762	.9992	.9343	.0005	.9797	.0023	49.63	2.47
ia-infect-dublin	2 765	.9999	.9652	.0010	.9999	.0000	43.38	4.86
email-univ	5 451	.9998	.9512	.0010	.9998	.0000	945.88	63.76
johnson16-2-4	5 460	.9998	.9325	.0005	.9721	.0008	15.36	1.63
p-hat700-1	60 999	.9970	.9172	.0001	.9794	.0003	471.34	31.92

TABLE 5.B. ?? fcc instances with  $\lceil 128 \ln m \rceil$  cuts sampled and  $\varepsilon = 10^{-4}$ , obtained on ?? hardware. See Table Description 2.

instance	$m$	RLI	$\beta_{fcc}$ avg	$\beta_{fcc}$ std	$\beta_{mc}$ avg	$\beta_{mc}$ std	time (s)	time std (s)
bayg29	405	0	.8849	.0006	.9882	.0000	0.07	0.01
bays29	406	3	.8862	.0008	.9885	.0000	0.06	<0.01
berlin52	1 326	0	.8899	.0000	.9922	.0000	0.32	0.05
brazil58	1 641	2	.8892	.0000	.9993	.0000	0.30	0.05
kroA100	2 500	0	.9999	.0000	.9999	.0000	2.73	0.54
gr96	4 542	1	.8809	.0007	.9981	.0000	1.25	0.17
eil101	5 046	0	.8785	.0000	.9799	.0011	1.44	0.13
gr120	6 942	5	.9701	.0000	.9997	.0000	2.28	0.31
bier127	8 001	0	.8791	.0003	.9842	.0000	6.71	0.84
ch130	8 372	1	.8785	.0002	.9890	.0000	2.86	0.40
gr137	8 803	3	.9864	.0000	.9988	.0000	3.22	0.52
ch150	11 165	0	.8781	.0004	.9862	.0001	3.95	0.23
brg180	16 023	2	.8866	.0007	.9740	.0014	21.70	4.79
d198	19 396	5	.8814	.0010	.9994	.0000	7.58	1.10
gr202	20 301	0	.8784	.0001	.9883	.0002	31.51	3.02
a280	38 874	3	.8794	.0000	.9994	.0000	21.23	2.37

TABLE 6.A. Paired ?? instances with  $\lceil 128 \ln m \rceil$  cuts sampled and  $\varepsilon = 10^{-4}$ , obtained on ?? hardware. See Table Description 3.

instance	$m$	RLI	$\beta_{fcc}$ avg	$\beta_{fcc}$ std	$\beta_{mc}$ avg	$\beta_{mc}$ std	time (s)	time std (s)
ENZYMES8	133	0	.9107	.0060	.9717	.0000	1.17	0.12
soc-dolphins	159	0	.8908	.0000	.9825	.0000	5.66	0.41
road-chesapeake	170	0	.8878	.0002	.9613	.0000	0.14	0.02
email-enron-only	623	0	.8787	.0000	.9639	.0023	4.80	0.48
dwt_209	767	0	.8910	.0000	.9668	.0014	814.51	75.20
ca-netscience	904	2	.8793	.0000	.9829	.0002	3 979.64	128.47
Erdos991	1 417	0	.8668	.0003	.9420	.0016	190.31	8.58
hamming6-2	1 824	0	.9999	.0000	.9999	.0000	0.80	0.12
inf-USAir97	2 126	1	.8776	.0000	.9845	.0005	1 781.84	151.11
ia-infect-hyper	2 196	0	.8894	.0003	.9409	.0012	6.86	0.66
DD687	2 597	0	.8812	.0004	.9419	.0014	76 286.39	3 302.99
dwt_503	2 762	3	.8859	.0000	.9830	.0003	25 857.31	893.64
ia-infect-dublin	2 765	0	.8817	.0003	.9415	.0012	15 269.45	355.02
email-univ	5 450	0	.8626	.0003	.9225	.0008	5 346.54	259.53
johnson16-2-4	5 460	0	.9329	.0007	.9718	.0012	14.91	1.76
p-hat700-1	60 999	0	.8701	.0003	.9447	.0009	457.74	21.99

TABLE 6.B. Paired ?? instances with  $\lceil 128 \ln m \rceil$  cuts sampled and  $\varepsilon = 10^{-4}$ , obtained on ?? hardware. See Table Description 3.

$\varepsilon$	$10^{-4}$	$1/64$	$10^{-4}$	$10^{-4}$	$1/64$	$10^{-4}$	$1/64$	$10^{-4}$	$1/64$
instance	$1 - \sigma_\varepsilon$	$1 - \sigma_\varepsilon$	RLI	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{mc}$	$\beta_{mc}$	time (s)	time (s)
bayg29	.9998	.9909	0	.8849	.8845	.9882	.9972	0.07	0.12
bays29	.9998	.9916	3	.8862	.8868	.9885	.9977	0.06	0.11
berlin52	.9994	.9890	0	.8899	.8872	.9922	.9963	0.32	1.03
brazil58	.9999	.9910	2	.8892	.8883	.9993	.9991	0.30	0.67
kroA100	.9999	.9920	0	.9999	.9998	.9999	.9998	2.73	2.63
gr96	.9998	.9918	1	.8809	.8832	.9981	.9997	1.25	4.06
eil101	.9998	.9915	0	.8785	.8782	.9799	.9965	1.44	4.31
gr120	.9996	.9918	5	.9701	.9700	.9997	.9996	2.28	6.72
bier127	.9997	.9911	0	.8791	.8777	.9842	.9904	6.71	21.33
ch130	.9998	.9915	1	.8785	.8781	.9890	.9992	2.86	8.34
gr137	.9985	.9918	3	.9864	.9873	.9988	.9997	3.22	10.07
ch150	.9998	.9917	0	.8781	.8783	.9862	.9980	3.95	12.85
brg180	.9999	.9896	2	.8866	.8828	.9740	.9771	21.70	71.15
d198	.9998	.9907	5	.8814	.8864	.9994	.9992	7.58	22.06
gr202	.9998	.9851	0	.8784	.8701	.9883	.9932	31.51	66.68
a280	.9998	.9919	3	.8794	.8795	.9994	.9997	21.23	66.74

TABLE 7.A. Effect of  $\varepsilon$  perturbation on **paired ??** instances with  $\lceil 128 \ln m \rceil$  cuts sampled and  $\varepsilon \in \{10^{-4}, 1/64\}$ , obtained on ?? hardware. No standard deviation was larger than .0014. See Table Description 3.

$\varepsilon$	$10^{-4}$	$1/64$	$10^{-4}$	$10^{-4}$	$1/64$	$10^{-4}$	$1/64$	$10^{-4}$	$1/64$
instance	$1 - \sigma_\varepsilon$	$1 - \sigma_\varepsilon$	RLI	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{mc}$	$\beta_{mc}$	time (s)	time (s)
ENZYMES8	.9926	.9849	0	.9107	.9002	.9717	.9734	1.17	1.29
soc-dolphins	.9995	.9917	0	.8908	.8908	.9825	.9988	5.66	2.59
road-chesapeake	.9985	.9914	0	.8878	.8883	.9613	.9994	0.14	0.24
email-enron-only	.9874	.9913	0	.8787	.8883	.9639	.9991	4.80	5.17
dwt_209	.9991	.9920	0	.8910	.8915	.9668	.9991	814.51	681.10
ca-netscience	.9991	.9916	2	.8793	.8794	.9829	.9973	3 979.64	3 239.05
Erdos991	.9763	.9885	0	.8668	.8798	.9420	.9931	190.31	491.03
hamming6-2	.9998	.9921	0	.9999	.9969	.9999	.9999	0.80	1.27
inf-USAir97	.9975	.9902	1	.8776	.8779	.9845	.9986	1 781.84	884.22
ia-infect-hyper	.9988	.9920	0	.8894	.8870	.9409	.9998	6.86	8.64
DD687	.9991	.9917	0	.8812	.8755	.9419	.9995	76 286.39	2 063.97
dwt_503	.9951	.9910	3	.8859	.8862	.9830	.9927	25 857.31	25 974.20
ia-infect-dublin	.9986	.9920	0	.8817	.8773	.9415	.9998	15 269.45	261.17
email-univ	.9878	.9915	0	.8626	.8666	.9225	.9982	5 346.54	6 772.99
johnson16-2-4	.9998	.9979	0	.9329	.9311	.9718	.9715	14.91	16.59
p-hat700-1	.9996	.9833	0	.8701	.8567	.9447	.9111	457.74	1 349.75

TABLE 7.B. Effect of  $\varepsilon$  perturbation on **paired ??** instances with  $\lceil 128 \ln m \rceil$  cuts sampled and  $\varepsilon \in \{10^{-4}, 1/64\}$ , obtained on ?? hardware. No standard deviation was larger than .0093. See Table Description 3.

## 5. OBSTRUCTIONS TO THE FEASIBILITY OF RESTRICTED LP

In Section 3 we introduced the perturbation parameter  $\varepsilon \in (0, 1)$  by sampling cuts using the matrix

$$Y := (1 - \varepsilon)\tilde{Y} + \varepsilon I, \text{ where } \tilde{Y} \text{ is the solver output for the SDP (2a); see (28)}$$

In Section 4, the constraints of (33) include  $\varepsilon$ , and we sample cuts using the matrix

$$Y := \tilde{Y} + \varepsilon\tilde{\mu}I, \text{ where } (\tilde{Y}, \tilde{\mu}) \text{ is the solver output for the primal SDP in (33).}$$

The common language is easily justified: up to scaling, in both cases we only sample cuts from matrices in the set  $\{Y \in \mathbb{S}^V : Y \succeq \varepsilon I, \text{diag}(Y) = \mathbb{1}\}$ . We start this section explaining the theoretical motivation for introducing perturbation. We study distinct choices of perturbation values in Section 5.1, and discuss the interaction between perturbation and our sanitization procedure in Section 5.2. In Section 5.3, we explore an alternative to perturbation based on uniform sampling of shores.

Let  $\gamma \in (0, 2)$ , and consider the graph  $G = K_3$  with  $z \in \mathbb{R}_+^E$  given by  $z_{12} = z_{13} = 1$  and  $z_{23} = \gamma$ . [5, Proposition 31] shows an optimal solution  $Y \in \mathbb{S}_+^V$  to the perturbation-free SDP (12a) such that

$$(34) \quad \mathbb{P}(\{2, 3\} \in \delta(\text{GW}(Y))) = \frac{2\sqrt{\gamma}}{\pi} + O(\gamma^{3/2}).$$

As feasibility is not attained until a cut containing the edge  $\{2, 3\}$  is sampled, the expected number of cuts necessary to attain feasibility grows exponentially with the encoding length of  $z$ . This graph  $G$  has in  $\{2, 3\}$  a “thin edge”, i.e., an edge  $e \in E$  whose positive value  $z_e$  is small relative to  $\|z\|_\infty$ . Figure 4 shows that thin edges raise issues in practice. For each  $\gamma \in \{1, 8^{-1}, \dots, 8^{-5}\}$ , we have a fractional cut-covering instance  $(G, z_\gamma)$ , which we feed to our pipeline. Figure 4 shows in blue the number of samples necessary to obtain feasibility for each of 10 independent runs. In orange, we show  $\frac{\pi}{2\sqrt{\gamma}}$ , motivated by the RHS of (35). Going from  $\gamma = 8^{-1}$  to  $\gamma = 8^{-5}$  corresponds to increasing the input by mere 12 bits, but as Figure 4 shows, this translates into going from an instance that attains feasibility with 10 samples, to an instance that does not attain feasibility after 2000 samples.

The experiments reported by Tables 0.A to 7.B, when performing 10 independent runs per instance, all solve the relevant SDP from scratch. With some instances taking up to 21 hours to be solved (as in Table 6.B), this quite significantly impacts our ability to collect data. To reuse SDP solver output, we introduced two implementations of abstract classes in our pipeline in Figure 1. The `FromFile` solver loads an SDP certificate computed by Algorithms 1 and 2 from a file. These files are generated using `SolverOutputStore`, which is an implementation of `Cover Producer` that simply stores the certificate (20) in a file. Leveraging these implementations, to collect data for 10 independent runs we run our pipeline eleven times: once just to solve the SDP relaxation and store it, and once for each independent run, sampling from the previously computed SDP solution.

**5.1. Perturbations of the Feasible Region.** In this subsection, we first explain how perturbation addresses the issue raised by thin edges. It is natural to ask which other perturbation values produce good behavior. This is explored in Tables 8.A to 9.B.

Let  $G = (V, E)$  be a graph, and let  $\varepsilon \in (0, 1)$  be the perturbation parameter. Our algorithms from Sections 3 and 4 succeed because they do not attempt to cover an adversarially chosen input  $z \in \mathbb{R}_+^E$ , rather covering a vector  $\hat{z} \in \mathbb{R}_+^E$  such that

$$(35) \quad \hat{z} \geq \frac{\varepsilon}{2} \|\hat{z}\|_\infty \mathbb{1} \text{ and } \hat{z} \geq z.$$

Since  $\hat{z} \geq z$ , every fractional cover for  $\hat{z}$  is a fractional cover for  $z$ , and hence we can cover every instance  $\hat{z} \geq z$  by covering instances such that  $\hat{z} \geq \varepsilon/2 \|\hat{z}\|_\infty \mathbb{1}$ . Covering  $\hat{z}$  instead of  $z$  can only weaken




FIGURE 4. Number of samples  $F$  necessary to attain feasibility for  $(G, z)$  where  $G = K_3$ ,  $z_{12} = z_{13} = 1$ , and  $z_{23} = \gamma$ . For  $\gamma = 8^{-5}$  we failed to achieve feasibility after  $2304 = 64 \times \lceil 32 \ln m \rceil$  samples for all 10 independent runs.

our certificates by a multiplicative factor of  $1 - \varepsilon$  [5, Theorem 21]. To prove (36), we use two observations. The first one is that

$$(36) \quad A \preceq B \text{ implies } \frac{1}{4}\mathcal{L}_G^*(A) \leq \frac{1}{4}\mathcal{L}_G^*(B) \text{ for every } A, B \in \mathbb{S}^V,$$

which follows from duality, as  $\mathcal{L}_G(w) \in \mathbb{S}_+^V$  for every  $w \in \mathbb{R}_+^E$ . The second observation is that

$$(37) \quad \begin{aligned} &\text{for every } z \in \mathbb{R}_+^E \text{ and for every } Y \in \mathbb{S}_+^V \text{ such that } \text{diag}(Y) = \mu \mathbb{1}, \\ &\text{if } \frac{1}{4}\mathcal{L}_G^*(Y) \geq z, \text{ then } \|z\|_\infty \leq \mu. \end{aligned}$$

Indeed, for every edge  $ij \in E$ ,  $z_{ij} \leq \langle \frac{1}{4}\mathcal{L}_G^*(Y), e_{ij} \rangle = \frac{1}{4}(Y_{ii} + Y_{jj} - 2Y_{ij}) = \frac{1}{2}(\mu - Y_{ij}) \leq \mu$ , where the last inequality holds since  $\mu^2 - Y_{ij}^2$  is a principal minor of the positive semidefinite matrix  $Y$ . In Sections 3 and 4, to cover  $z \in \mathbb{R}_+^E$  we sampled shores from  $Y \in \mathbb{S}^V$  such that

$$Y \succeq \varepsilon \mu I, \text{diag}(Y) = \mu \mathbb{1} \text{ and } \frac{1}{4}\mathcal{L}_G^*(Y) \geq z$$

for  $\mu \in \mathbb{R}_+$ . Set  $\hat{z} := \frac{1}{4}\mathcal{L}_G^*(Y)$ . It is immediate that  $\hat{z} \geq z$ , and moreover by (37) and (38) we have  $\hat{z} = \frac{1}{4}\mathcal{L}_G^*(Y) \geq \frac{1}{4}\mathcal{L}_G^*(\varepsilon \mu I) = \frac{\varepsilon}{2}\mu \mathbb{1} \geq \frac{\varepsilon}{2}\|\hat{z}\|_\infty \mathbb{1}$ . Hence (36) holds.

**Table Description 4** (Tables 8.A and 8.B). Experiments obtained on ?? hardware. As input, we are given maximum cut instances  $(G, w)$ . In Table 8.A they are ?? instances, and in Table 8.B they are ?? instances. We nearly solve the SDPs in (2) using SCS with (27), obtaining  $\tilde{x} \in \mathbb{R}^V$  and  $\tilde{Y} \in \mathbb{S}^V$ . We set  $Y := (1 - \varepsilon)\tilde{Y} + \varepsilon I$  as in (28), and  $z := \frac{1}{4}\mathcal{L}_G^*(Y)$  as in STEP 1. Set  $(\rho, x, B) := \text{SLACKSANITIZE}(\tilde{x}, w)$  and  $(\mu, R \in \mathbb{R}^{k \times V}) := \text{REPRESENTATIONSANITIZE}(Y, z)$ . We store  $w, z$ , along with the elements  $\rho, \mu, R, B$ , and  $x$  of the SDP certificate (20), in a file. This step is done once for each choice of  $\varepsilon \in \{10^{-4}, 1/64, 1/16, 1/8, 1/4\}$  and  $(G, w)$ .

For each such file, we run our pipeline with 10 different seeds used to sample cuts. In these runs, we first load  $(w, z)$  and the SDP certificate from the file. We sample  $\lceil 128 \ln m \rceil$  vectors  $g \in \mathbb{R}^k$  according to the standard multivariate normal distribution, and produce a shore  $S := \{i \in V : g^\top R e_i > 0\}$  for each  $g$ . Let  $\mathcal{F} \subseteq \mathcal{P}(V)$  be the set of shores generated. We compute  $\text{mc}(\mathcal{F}, w)$  directly, and we solve  $\text{fcc}(\mathcal{F}, z)$  using Gurobi. We then define  $\beta_{\text{mc}}$  and  $\beta_{\text{fcc}}$  as in (32). Statistics for  $\beta_{\text{mc}}$  use all 10 runs; those for  $\beta_{\text{fcc}}$  use only runs that produced a feasible fractional cut cover for  $z$ , thus excluding the infeasible runs counted in the RLI column. For values of  $\varepsilon$  with feasible fractional cut covers in all 10 runs (i.e., zero RLI count), this column is omitted.

We first focus on the approximation quality  $\beta_{\text{mc}}$  obtained on the maximum cut instances in Tables 8.A and 8.B. Table 8.A is remarkable in how big perturbations of the nearly optimal solution do not meaningfully affect the quality of the best cut sampled. Indeed, for all instances in this table, the same  $\beta_{\text{mc}}$  value is obtained to within 2 digits across all values of  $\varepsilon$ . On the other hand, for instances in Table 8.B, each additional increase in  $\varepsilon$  has a negative impact on  $\beta_{\text{mc}}$ . Notice that for `email-univ` the value of  $\beta_{\text{mc}}$  degrades from .9260 to .8942, and for `DD687` it degrades from .9407 to .9065.

Now consider the approximation quality  $\beta_{\text{fcc}}$  on the same tables. Recalling (32), note that each  $\beta_{\text{fcc}}$  column in Tables 8.A and 8.B corresponds to covering a distinct vector  $z_\varepsilon \in \mathbb{R}_+^E$ . This is fundamentally different from the columns reporting  $\beta_{\text{mc}}$ , where the instance  $w \in \mathbb{R}_+^E$  is kept fixed. Perturbation values  $\varepsilon$  larger than or equal to  $1/64$  made the outcome of our algorithm robust: the only RLI outcomes occurred for  $\varepsilon = 10^{-4}$ . In Table 8.A, the value of  $\beta_{\text{fcc}}$  increases with  $\varepsilon$  for most instances. Still, even in this set of examples there are instances where increasing from  $\varepsilon = 1/64$  to  $\varepsilon = 1/4$  degrades  $\beta_{\text{fcc}}$ : see `kroA100`, `gr120`, `gr137`, and `brg180`. This degradation, however, is much more noticeable in Table 8.B, where `email-enron-only`, `inf-USAir97`, and `ia-infect-hyper` are the only instances for which the highest  $\beta_{\text{fcc}}$  is attained for some  $\varepsilon \geq 1/16$ .

Tables 9.A and 9.B have as input fractional cut-covering instances obtained via pairing (see Section 4). For these tables, recall from (34) that  $z \in \mathbb{R}_+^E$  is given an input, which is then paired to  $w_\varepsilon \in \mathbb{R}_+^E$  in a way that is dependent on  $\varepsilon$ . In particular, each  $\beta_{\text{mc}}$  column in Tables 9.A and 9.B certifies the quality of a cut with respect to distinct  $w_\varepsilon \in \mathbb{R}_+^E$ . Whereas changing from  $\varepsilon = 10^{-4}$  to  $\varepsilon = 1/64$  is sufficient to drive RLI to zero (as we had observed in Section 4), subsequently increasing  $\varepsilon$  may actually harm the value of  $\beta_{\text{fcc}}$ , as one can see in Table 9.B. There is no significant improvement on  $\beta_{\text{fcc}}$  for values of  $\varepsilon$  above  $1/16$ . It is noticeable how often in Tables 9.A and 9.B the maximum cut instances obtained via pairing have certifiable cuts with ratios above .99, especially for  $\varepsilon \geq 1/64$ .

Since specific instances respond differently to perturbation, given either a maximum cut or a fractional cut-covering instance, it is worthwhile to experiment with distinct choices of  $\varepsilon$ . Tables 8.A to 9.B suggest it may be reasonable to start such exploration by selecting perturbation parameters in the range  $10^{-4} < \varepsilon \leq 1/16$ .

**Table Description 5** (Tables 9.A and 9.B). Experiments obtained on ?? hardware. We are given fractional cut-covering instances  $(G, z)$  as input. In all cases they are **paired** instances, obtained by solving the corresponding maximum cut instance via Algorithm 3; see Section 6.1. In Table 9.A they are ?? instances, and in Table 9.B they are ?. We nearly solve the SDPs in (33) with  $\varepsilon \in \{10^{-4}, 1/64, 1/16, 1/8, 1/4\}$  using SCS via (61), thus computing nearly optimal  $\tilde{\mu} \in \mathbb{R}$ ,  $\tilde{Y} \in \mathbb{S}^V$ ,  $\tilde{w} \in \mathbb{R}^E$ , and  $\tilde{x} \in \mathbb{R}^V$ . Set  $Y := \tilde{Y} + \varepsilon \tilde{\mu} I$ , set  $(\rho, x, B) := \text{SLACKSANITIZE}(\tilde{x}, \tilde{w})$ , and set  $(\mu, R \in \mathbb{R}^{[k] \times V}) := \text{REPRESENTATIONSANITIZE}(Y, z)$ . We store  $\tilde{w}$ ,  $z$ , along with the elements  $\rho$ ,  $\mu$ ,  $R$ ,  $B$ , and  $x$  of the SDP certificate (20), in a file. This step is done once for each choice of  $\varepsilon \in \{10^{-4}, 1/64, 1/16, 1/8, 1/4\}$  and  $(G, z)$ .

For each such file, we run our pipeline with 10 different seeds used to sample cuts. In these runs, we first load  $(w, z)$  and the SDP certificate from the file. We sample  $\lceil 128 \log m \rceil$  vectors  $g \in \mathbb{R}^k$  according to the standard multivariate normal distribution, and produce a shore  $S := \{i \in V : g^\top R e_i > 0\}$  for each  $g$ . Let  $\mathcal{F} \subseteq \mathcal{P}(V)$  be the set of shores generated. We compute  $\text{mc}(\mathcal{F}, w)$  directly, and we solve  $\text{fcc}(\mathcal{F}, z)$  using Gurobi. We then define  $\beta_{\text{mc}}$ , and  $\beta_{\text{fcc}}$  as in (34). Statistics for  $\beta_{\text{mc}}$  use all 10 runs; those for  $\beta_{\text{fcc}}$  use only runs that produced a feasible fractional cut cover for  $z$ , thus excluding the infeasible runs counted in the **RLI** column. For values of  $\varepsilon$  with feasible fractional cut covers in all 10 runs (i.e., zero **RLI** count), this column is omitted.

**5.2. Interaction between Perturbation and Sanitization.** As one can see from the results for **brg180** in Tables 0.A and 1.A, perturbation was helpful in producing SDP solutions that our sanitization procedures would accept. Whereas the positive effect of perturbation on feasibility was known before this paper, its positive effect on sanitization only became clear from our experiments. We propose two possible explanations for this behavior.

Let  $G = (V, E)$  be a graph, and let  $z \in \mathbb{R}_+^E$ . Using (38) and (37), we have that

$$z \leq \|z\|_\infty \mathbb{1} \leq \mu \mathbb{1} = \frac{1}{2\varepsilon} \mathcal{L}_G^*(\mu\varepsilon I) \leq \frac{1}{2\varepsilon} \mathcal{L}_G^*(Y),$$

so

$$(38) \quad z \leq \frac{1}{2\varepsilon} \mathcal{L}_G^*(Y) \text{ if } Y \succeq \varepsilon \mu I, \text{diag}(Y) = \mu \mathbb{1}, \text{ and } \frac{1}{4} \mathcal{L}_G^*(Y) \geq z.$$

To relate (39) to our sanitization procedure, note that by line 7, Algorithm 2 is working with a Cholesky factorization of  $\tilde{Y}$  with  $\text{diag}(\tilde{Y}) = \mathbb{1}$ . Then step 8 computes  $\inf\{\mu \in \mathbb{R}_+ : \mu \frac{1}{4} \mathcal{L}_G^*(\tilde{Y}) \geq z\}$ . Since

$$\frac{z_{ij}}{\frac{1}{4}(\mathcal{L}_G^*(Y))_{ij}} \leq 2/\varepsilon$$

by (39), the divisions in step 8 of Algorithm 2 are all numerically well-behaved. Of course, there is a disconnect between (39), which is proved using exact arithmetic, and the actual floating-point arithmetic performed in our sanitization procedure. Still, approximate feasibility is enough to provide an upper bound for  $\mu$  not much weaker than  $2/\varepsilon$ , which then ensures that dividing  $z_{ij}$  by  $(\mathcal{L}_G^*(Y))_{ij}$  is well-behaved for every edge  $ij \in E$ , and thus that our computation of  $\mu$  will succeed.

Perturbation could also be helpful for Algorithm 2 in the following way. Let  $Y \in \mathbb{S}^V$  be such that  $\frac{1}{4} \mathcal{L}_G^*(Y) \geq z$ , and let  $\gamma \in (0, 1)$  be the parameter used in Algorithm 2. Let  $P, Z, N \subseteq \mathbb{S}_+^V$  be such that  $Y = P + Z - N$ , with  $\lambda_{\min}(P) > \gamma$  and  $\lambda_{\max}(Z) \leq \gamma$ . Numerically, we consider  $P$  the projection of  $Y$  onto  $\mathbb{S}_+^V$ , thus ignoring  $Z$ . It holds that  $\mathcal{L}_G^*(Z) \leq 2\gamma \mathbb{1}$ , as for every  $ij \in E$ ,

$$(\mathcal{L}_G^*(Z))_{ij} = \langle \mathcal{L}_G(e_{ij}), Z \rangle = (e_i - e_j)^\top Z (e_i - e_j) \leq \lambda_{\max}(Z) \|e_i - e_j\|_2^2 = 2\lambda_{\max}(Z).$$

Since  $Y = P + Z - N \preceq P + Z$ , we can then use (37) to see that

$$z \leq \frac{1}{4} \mathcal{L}_G^*(Y) \leq \frac{1}{4} \mathcal{L}_G^*(P) + \frac{1}{4} \mathcal{L}_G^*(Z) \leq \frac{1}{4} \mathcal{L}_G^*(P) + \frac{1}{2} \gamma \mathbb{1}.$$

Hence  $(\perp_Y)$  could occur, as it is possible that  $z_{ij} > 0 = \frac{1}{4}\mathcal{L}_G^*(P)_{ij}$  if  $\frac{1}{2}\gamma \geq z_{ij}$ . Note that  $Z = N = 0$  if  $\lambda_{\min}(Y) \geq \varepsilon > \gamma$ , so perturbed solutions  $Y$  do not face this issue.

**5.3. Uniform Sampling Shores.** Let  $U: \Omega \rightarrow \mathcal{P}(V)$  be a random variable representing the uniform distribution over subsets of  $V$ . Samples from  $U$  are helpful in covering thin edges since every edge  $e \in E$  is treated uniformly, with  $\mathbb{P}(e \in \delta(U)) = \frac{1}{2}$ . In particular, if one has  $T \in \mathbb{N}$  independent samples  $U_1, \dots, U_T$ , every edge is covered after  $O(\ln m)$  samples with high probability, since

$$\mathbb{P}\left(\bigcup_{t=1}^T \delta(U_t) \neq E\right) \leq \sum_{e \in E} \mathbb{P}\left(e \notin \bigcup_{t=1}^T \delta(U_t)\right) = m2^{-T}.$$

The experiments in this subsection share a common structure: pick a uniform sampling quota  $Q \in \mathbb{N}$  and compute a fractional cut cover using  $\lceil Q \ln m \rceil$  shores sampled from  $U$ , and  $\lceil (128 - Q) \ln m \rceil$  shores sampled from the geometric representation stored by the SDP certificate (20). Our options for  $Q$  arise from the values of  $\varepsilon$  in Tables 8.A to 9.B, as we consider

$$Q \in \{2, 8, 16, 32\} = \{128\varepsilon : \varepsilon \in \{1/64, 1/16, 1/8, 1/4\}\}.$$

As motivated in Section 5.2, we perturb the optimal solution by  $\varepsilon = 10^{-4}$  to obtain SDP solutions acceptable to our sanitization procedures.

Tables 8.A to 9.B have flexibility in picking different pairings, as distinct columns in Tables 8.A and 8.B report  $\beta_{\text{fcc}}$  values for distinct  $z_\varepsilon$ , and distinct columns on Tables 9.A and 9.B report  $\beta_{\text{mc}}$  values for distinct  $w_\varepsilon$ . This is not the case for the upcoming tables in this subsection. In every column of Tables 10.A and 10.B we are certifying the same pair  $(w, z_{10^{-4}}) \in H_\sigma(G)$ , and in every column of Tables 11.A and 11.B we are certifying the same pair  $(w_{10^{-4}}, z) \in H_\sigma(G)$ .

**Table Description 6** (Tables 10.A and 10.B). Experiments obtained on ?? hardware. As input, we are given maximum cut instances  $(G, w)$ . In Table 10.A they are ?? instances, and in Table 10.B they are ?? instances. We nearly solve the SDPs in (2) using SCS with (27), obtaining  $\tilde{x} \in \mathbb{R}^V$  and  $\tilde{Y} \in \mathbb{S}^V$ . We set  $Y := (1 - \varepsilon)\tilde{Y} + \varepsilon I$  as in (28), with  $\varepsilon = 10^{-4}$ , and  $z := \frac{1}{4}\mathcal{L}_G^*(Y)$  as in STEP 1. Set  $(\rho, x, B) := \text{SLACKSANITIZE}(\tilde{x}, w)$  and  $(\mu, R \in \mathbb{R}^{k \times V}) := \text{REPRESENTATIONSANITIZE}(Y, z)$ . We store  $w, z$ , along with the elements  $\rho, \mu, R, B$ , and  $x$  of the SDP certificate (20), in a file. and  $B$  in a file. This step is done once for each  $(G, w)$ .

For each such file, for each number  $Q \in \{2, 8, 16, 32\}$ , and for each of 10 different seeds used to sample cuts, we run our pipeline once. In these runs, we first load  $(w, z)$  and the SDP certificate from the file. We sample  $\lceil Q \ln m \rceil$  shores from the uniform distribution on  $\mathcal{P}(V)$ . We then sample  $\lceil (128 - Q) \ln m \rceil$  vectors  $g \in \mathbb{R}^k$  according to the standard multivariate normal distribution, and produce a shore  $S := \{i \in V : g^\top R e_i > 0\}$  for each  $g$ . Let  $\mathcal{F} \subseteq \mathcal{P}(V)$  be the set of shores generated by both procedures. We compute  $\text{mc}(\mathcal{F}, w)$  directly, and we solve  $\text{fcc}(\mathcal{F}, z)$  using Gurobi. We then define  $\beta_{\text{mc}}$  and  $\beta_{\text{fcc}}$  as in (18). The entries in Tables 11.A and 11.B are the average across the 10 runs for each choice of choice of  $Q$  and input graph. For each  $\beta_{\text{fcc}}$  column with a sampling quota  $Q$ , we display beside it the  $\beta_{\text{fcc}}$  column from Tables 8.A and 8.B with perturbation  $\varepsilon$  such that  $Q = 128\varepsilon$ .

**Table Description 7** (Tables 11.A and 11.B). Experiments obtained on ?? hardware. As input, we are given fractional cut-covering instances  $(G, z)$ . In all cases they are **paired** instances, obtained by solving the corresponding maximum cut instance via Algorithm 3; see Section 6.1. In Table 11.A they are ?? instances, and in Table 11.B they are ??. We nearly solve the SDPs in (33) with  $\varepsilon = 10^{-4}$  using SCS via (61), thus computing nearly optimal  $\tilde{\mu} \in \mathbb{R}$ ,  $\tilde{Y} \in \mathbb{S}^V$ ,  $\tilde{w} \in \mathbb{R}^E$ , and  $\tilde{x} \in \mathbb{R}^V$ . Set  $Y := \tilde{Y} + \tilde{\mu}\varepsilon I$ , set  $(\rho, x, B) := \text{SLACKSANITIZE}(\tilde{x}, w)$ , and set  $(\mu, R \in \mathbb{R}^{k \times V}) := \text{REPRESENTATIONSANITIZE}(Y, z)$ . We store  $\tilde{w}$ ,  $z$ , along with the elements  $\rho$ ,  $\mu$ ,  $R$ ,  $B$ , and  $x$  of the SDP certificate (20), in a file. This step is done once for each  $(G, z)$ .

For each such file, for each number  $Q \in \{2, 8, 16, 32\}$ , and for each of 10 different seeds used to sample cuts, we run our pipeline once. In these runs, we first load  $(w, z)$  and the SDP certificate from the file. We sample  $\lceil Q \ln m \rceil$  shores from the uniform distribution on  $\mathcal{P}(V)$ . We then sample  $\lceil (128 - Q) \ln m \rceil$  vectors  $g \in \mathbb{R}^k$  according to the standard multivariate normal distribution, and produce a shore  $S := \{i \in V : g^\top R e_i > 0\}$  for each  $g$ . Let  $\mathcal{F} \subseteq \mathcal{P}(V)$  be the set of shores generated by both procedures. We compute  $\text{mc}(\mathcal{F}, w)$  directly, and we solve  $\text{fcc}(\mathcal{F}, z)$  using Gurobi. We then define  $\beta_{\text{mc}}$  and  $\beta_{\text{fcc}}$  as in (18). The entries in Tables 11.A and 11.B are the average across the 10 runs for each choice of choice of  $Q$  and input graph. For each  $\beta_{\text{fcc}}$  column with a sampling quota  $Q$ , we display beside it the  $\beta_{\text{fcc}}$  column from Tables 9.A and 9.B with perturbation  $\varepsilon$  such that  $Q = 128\varepsilon$ .

For experiments running the pipeline twice, running time information becomes ambiguous. The *single-run running time* of a run is the sum of

- (1) the time spent in the first run until the completion of the computation of the SDP certificate (20), and
- (2) the time spent in the second run, starting after the SDP certificates (20) are loaded from file.

The single-run running time across Tables 10.A to 11.B never changes too much. This is to be expected, as the time spent in SDP solving is fixed and LP solving should not change too much. The average single-run running time for the maximum cut instance **gr96** is 9.58 seconds when  $Q = 2$ , and 19.68 seconds when  $Q = 32$ , increasing by a factor of 2.05. This is the largest difference on average single-run running time we have found on the experiments in these tables. The highest standard deviation across the 10 seeds we have observed in these tables is 0.0072 for  $\beta_{\text{fcc}}$  in Table 10.B.

We omit  $\beta_{\text{mc}}$  values in Tables 10.A to 11.B since  $Q$  does not noticeably affect  $\beta_{\text{mc}}$ . This is unsurprising, since  $\lceil (128 - Q) \ln m \rceil$  samples from the nearly optimal SDP solution tend to provide good  $\beta_{\text{mc}}$  values. Uniform sampling fulfilled its primary purpose for  $\beta_{\text{fcc}}$ : there are no **RLI** outcomes in Tables 10.A to 11.B, an improvement on Tables 1.A and 1.B corresponding to  $Q = 0$ . The effect of  $Q$  on the average  $\beta_{\text{fcc}}$  is not as strong as the effect of  $\varepsilon$ , and despite the clear benefit from going from  $Q = 0$  to  $Q = 2$ , there is little evidence supporting choices of  $Q$  above 2. These observations further explain the difference, with respect to  $\beta_{\text{fcc}}$ , between uniform sampling and perturbation whenever  $Q = \varepsilon 128$ . When  $Q = 2$  and  $\varepsilon = 1/64$ , instances in Tables 10.B to 11.B evenly split between those where uniform sampling outperforms perturbation and those where it does not. Since higher values of  $Q$  have small effect on  $\beta_{\text{fcc}}$ , to ask which approach provides better results is to ask whether higher perturbation improves or not the value of  $\beta_{\text{fcc}}$ . This question is one of the main topics studied in Section 7. Table 10.A is an exception, where uniform sampling consistently outperforms perturbation.

$\varepsilon$	$10^{-4}$	$1/64$	$1/16$	$1/8$	$1/4$	$10^{-4}$	$10^{-4}$	$1/64$	$1/16$	$1/8$	$1/4$
instance	$\beta_{mc}$	$\beta_{mc}$	$\beta_{mc}$	$\beta_{mc}$	$\beta_{mc}$	RLI	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$
bayg29	.9934	.9934	.9934	.9934	.9934	1	.8838	.8864	.8942	.9015	.9147
bays29	.9938	.9938	.9938	.9938	.9938	2	.8837	.8873	.8946	.9007	.9138
berlin52	.9896	.9896	.9896	.9896	.9896	0	.8903	.8922	.8980	.9059	.9191
brazil158	.9993	.9993	.9993	.9993	.9993	3	.8888	.8920	.8959	.9013	.9129
gr96	.9982	.9982	.9982	.9982	.9982	6	.8793	.8817	.8888	.8988	.9068
kroA100	.9979	.9979	.9979	.9979	.9979	10	—	.9902	.9746	.9548	.9160
eil101	.9809	.9809	.9809	.9808	.9808	1	.8782	.8809	.8887	.8990	.9086
gr120	.9988	.9988	.9988	.9988	.9988	5	.9463	.9465	.9470	.9472	.9131
bier127	.9731	.9731	.9726	.9721	.9709	0	.8773	.8802	.8854	.8911	.8963
ch130	.9867	.9867	.9867	.9867	.9865	3	.8772	.8799	.8876	.8978	.9041
gr137	.9996	.9996	.9996	.9996	.9996	9	.9690	.9689	.9660	.9506	.9127
ch150	.9858	.9858	.9858	.9858	.9855	1	.8766	.8795	.8874	.8976	.9026
brg180	.9418	.9427	.9443	.9444	.9431	4	.8857	.8869	.8860	.8833	.8774
d198	.9975	.9975	.9975	.9975	.9974	7	.8780	.8868	.8907	.8942	.8929
gr202	.9768	.9768	.9763	.9759	.9743	0	.8770	.8782	.8830	.8885	.8923
a280	.9969	.9969	.9969	.9969	.9968	4	.8741	.8765	.8823	.8885	.8867

TABLE 8.A. ?? maximum cut instances with  $\lceil 128 \ln m \rceil$  cuts sampled and  $\varepsilon \in \{10^{-4}, 1/64, 1/16, 1/8, 1/4\}$ , obtained on ?? hardware. See Table Description 4.

$\varepsilon$	$10^{-4}$	$1/64$	$1/16$	$1/8$	$1/4$	$10^{-4}$	$10^{-4}$	$1/64$	$1/16$	$1/8$	$1/4$
instance	$\beta_{mc}$	$\beta_{mc}$	$\beta_{mc}$	$\beta_{mc}$	$\beta_{mc}$	RLI	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$
ENZYMES8	.9785	.9785	.9785	.9785	.9785	0	.9149	.9101	.9043	.8990	.8791
soc-dolphins	.9730	.9730	.9730	.9675	.9643	0	.8900	.8895	.8881	.8861	.8819
road-chesapeake	.9673	.9673	.9673	.9673	.9673	0	.8880	.8881	.8875	.8867	.8850
email-enron-only	.9595	.9586	.9554	.9536	.9457	0	.8882	.8887	.8902	.8846	.8707
dwt_209	.9658	.9654	.9628	.9567	.9491	0	.8924	.8924	.8917	.8837	.8631
ca-netscience	.9637	.9611	.9551	.9464	.9305	2	.8856	.8861	.8843	.8735	.8458
Erdos991	.9357	.9314	.9258	.9180	.9006	0	.8791	.8753	.8639	.8491	.8212
hamming6-2	.9997	.9997	.9997	.9997	.9951	0	.9997	.9954	.9824	.9517	.9314
inf-USAir97	.9668	.9665	.9660	.9648	.9619	0	.8736	.8744	.8777	.8756	.8575
ia-infect-hyper	.9687	.9692	.9679	.9658	.9622	0	.8898	.8899	.8903	.8909	.8904
DD687	.9407	.9388	.9330	.9245	.9065	0	.8801	.8760	.8648	.8508	.8236
dwt_503	.9844	.9841	.9822	.9792	.9740	1	.8859	.8880	.8874	.8808	.8617
ia-infect-dublin	.9491	.9471	.9435	.9386	.9278	0	.8805	.8778	.8707	.8622	.8459
email-univ	.9260	.9230	.9169	.9091	.8942	0	.8690	.8650	.8536	.8398	.8138
johnson16-2-4	.9720	.9713	.9691	.9658	.9582	0	.9325	.9308	.9261	.9206	.9084
p-hat700-1	.9704	.9698	.9689	.9668	.9631	0	.8696	.8685	.8664	.8647	.8627

TABLE 8.B. ?? maximum cut instances with  $\lceil 128 \ln m \rceil$  cuts sampled and  $\varepsilon \in \{10^{-4}, 1/64, 1/16, 1/8, 1/4\}$ , obtained on ?? hardware. See Table Description 4.

$\varepsilon$	$10^{-4}$	$10^{-4}$	$1/64$	$1/16$	$1/8$	$1/4$	$10^{-4}$	$1/64$	$1/16$	$1/8$	$1/4$
instance	RLI	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{mc}$	$\beta_{mc}$	$\beta_{mc}$	$\beta_{mc}$	$\beta_{mc}$
bayg29	0	.8849	.8845	.8849	.8852	.8853	.9882	.9972	.9980	.9987	.9998
bays29	3	.8862	.8868	.8870	.8863	.8867	.9885	.9977	.9990	.9994	.9998
berlin52	0	.8899	.8872	.8876	.8880	.8790	.9922	.9963	.9981	.9986	.9995
brazil158	2	.8892	.8883	.8879	.8892	.8890	.9993	.9991	.9988	.9999	.9997
kroA100	0	.9999	.9998	.9999	.9999	.9999	.9999	.9998	.9999	.9999	.9999
gr96	1	.8809	.8832	.8830	.8828	.8758	.9981	.9997	.9996	.9994	.9998
eil101	0	.8785	.8782	.8780	.8771	.8702	.9799	.9965	.9996	.9985	.9956
gr120	5	.9701	.9700	.9681	.9699	.9703	.9997	.9996	.9976	.9994	.9998
bier127	0	.8791	.8777	.8673	.8601	.8440	.9842	.9904	.9936	.9943	.9996
ch130	1	.8785	.8781	.8781	.8771	.8695	.9890	.9992	.9995	.9985	.9994
gr137	3	.9864	.9873	.9776	.9870	.9867	.9988	.9997	.9898	.9993	.9989
ch150	1	.8781	.8783	.8771	.8779	.8664	.9862	.9981	.9985	.9996	.9983
brg180	4	.8867	.8846	.8742	.8605	.8283	.9738	.9789	.9817	.9849	.9929
d198	5	.8813	.8865	.8868	.8830	.8645	.9994	.9992	.9998	.9998	.9976
gr202	0	.8783	.8702	.8706	.8525	.8448	.9883	.9932	.9985	.9878	.9998
a280	5	.8794	.8795	.8794	.8771	.8695	.9994	.9997	.9997	.9998	.9997

TABLE 9.A. ?? paired fractional cut-covering instances with  $\lceil 128 \ln m \rceil$  cuts sampled and  $\varepsilon \in \{10^{-4}, 1/64, 1/16, 1/8, 1/4\}$ , obtained on ?? hardware. See Table Description 5.

$\varepsilon$	$10^{-4}$	$10^{-4}$	$1/64$	$1/16$	$1/8$	$1/4$	$10^{-4}$	$1/64$	$1/16$	$1/8$	$1/4$
instance	RLI	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{mc}$	$\beta_{mc}$	$\beta_{mc}$	$\beta_{mc}$	$\beta_{mc}$
ENZYMES8	0	.9107	.9002	.9037	.9042	.8908	.9717	.9734	.9943	.9945	.9988
soc-dolphins	0	.8908	.8908	.8899	.8908	.8850	.9825	.9988	.9986	.9997	.9997
road-chesapeake	0	.8878	.8883	.8888	.8885	.8875	.9613	.9994	.9999	.9997	.9997
email-enron-only	0	.8787	.8883	.8856	.8733	.8440	.9639	.9991	.9999	.9999	.9999
dwt_209	0	.8910	.8915	.8881	.8753	.8425	.9667	.9992	.9998	.9997	.9994
ca-netscience	5	.8793	.8794	.8796	.8619	.8184	.9829	.9973	.9971	.9939	.9716
Erdos991	0	.8671	.8796	.8683	.8523	.8161	.9410	.9933	.9965	.9827	.9507
hamming6-2	0	.9999	.9996	.9694	.9175	.8671	.9999	.9997	.9997	.9999	.9936
inf-USAir97	0	.8776	.8779	.8768	.8754	.8566	.9844	.9986	.9984	.9983	.9977
ia-infect-hyper	0	.8894	.8870	.8785	.8673	.8408	.9409	.9998	.9998	.9998	.9994
DD687	0	.8811	.8755	.8563	.8385	.7765	.9410	.9995	.9998	.9999	.9998
dwt_503	2	.8859	.8861	.8818	.8655	.8291	.9829	.9927	.9938	.9940	.9970
ia-infect-dublin	0	.8816	.8776	.8642	.8463	.8092	.9412	.9998	.9999	.9994	.9993
email-univ	0	.8626	.8666	.8504	.8313	.7920	.9222	.9989	.9815	.9621	.9247
johnson16-2-4	0	.9328	.9310	.9264	.9200	.9081	.9718	.9715	.9676	.9641	.9581
p-hat700-1	0	.8701	.8647	.8523	.8365	.7671	.9448	.9195	.9183	.9195	.9388

TABLE 9.B. ?? paired fractional cut-covering instances with  $\lceil 128 \ln m \rceil$  cuts sampled and  $\varepsilon \in \{10^{-4}, 1/64, 1/16, 1/8, 1/4\}$ , obtained on ?? hardware. See Table Description 5.

instance	$\varepsilon = 1/64$	$Q = 2$	$\varepsilon = 1/16$	$Q = 8$	$\varepsilon = 1/8$	$Q = 16$	$\varepsilon = 1/4$	$Q = 32$
	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$
bayg29	.8864	.8838	.8942	.8837	.9015	.8840	.9147	.8839
bays29	.8873	.8842	.8946	.8847	.9007	.8843	.9138	.8848
berlin52	.8922	.8903	.8980	.8904	.9059	.8904	.9191	.8904
brazil58	.8920	.8892	.8959	.8898	.9013	.8904	.9129	.8901
gr96	.8817	.8793	.8888	.8793	.8988	.8794	.9068	.8794
kroA100	.9902	.9953	.9746	.9955	.9548	.9956	.9160	.9957
eil101	.8809	.8781	.8887	.8781	.8990	.8779	.9086	.8779
gr120	.9465	.9463	.9470	.9463	.9472	.9463	.9131	.9463
bier127	.8802	.8774	.8854	.8775	.8911	.8774	.8963	.8774
ch130	.8799	.8773	.8876	.8771	.8978	.8770	.9041	.8768
gr137	.9689	.9690	.9660	.9690	.9506	.9690	.9127	.9690
ch150	.8795	.8766	.8874	.8764	.8976	.8764	.9026	.8762
brg180	.8869	.8853	.8860	.8855	.8833	.8845	.8774	.8841
d198	.8868	.8778	.8907	.8783	.8942	.8783	.8929	.8774
gr202	.8782	.8771	.8830	.8769	.8885	.8766	.8923	.8765
a280	.8765	.8742	.8823	.8741	.8885	.8740	.8867	.8738

TABLE 10.A. ?? maxcut instances with  $\lceil Q \ln m \rceil$  cuts sampled uniformly, and  $\lceil (128 - Q) \ln m \rceil$  cuts sampled from SDP solution  $Y$  from (28), with  $\varepsilon = 10^{-4}$  and for each  $Q \in \{2, 8, 16, 32\}$ . See Table Description 6.

instance	$\varepsilon = 1/64$	$Q = 2$	$\varepsilon = 1/16$	$Q = 8$	$\varepsilon = 1/8$	$Q = 16$	$\varepsilon = 1/4$	$Q = 32$
	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$	$\beta_{fcc}$
ENZYMES8	.9101	.9148	.9043	.9088	.8990	.9134	.8791	.9146
soc-dolphins	.8895	.8900	.8881	.8900	.8861	.8900	.8819	.8900
road-chesapeake	.8881	.8882	.8875	.8883	.8867	.8883	.8850	.8882
email-enron-only	.8887	.8882	.8902	.8882	.8846	.8882	.8707	.8882
dwt_209	.8924	.8924	.8917	.8924	.8837	.8924	.8631	.8924
ca-netscience	.8861	.8856	.8843	.8856	.8735	.8856	.8458	.8856
Erdos991	.8753	.8792	.8639	.8786	.8491	.8777	.8212	.8752
hamming6-2	.9954	.9997	.9824	.9997	.9517	.9997	.9314	.9997
inf-USAir97	.8744	.8738	.8777	.8738	.8756	.8740	.8575	.8738
ia-infect-hyper	.8899	.8898	.8903	.8894	.8909	.8887	.8904	.8873
DD687	.8760	.8797	.8648	.8792	.8508	.8782	.8236	.8758
dwt_503	.8880	.8859	.8874	.8859	.8808	.8859	.8617	.8859
ia-infect-dublin	.8778	.8802	.8707	.8794	.8622	.8788	.8459	.8765
email-univ	.8650	.8690	.8536	.8679	.8398	.8670	.8138	.8642
johnson16-2-4	.9308	.9316	.9261	.9304	.9206	.9289	.9084	.9241
p-hat700-1	.8685	.8693	.8664	.8684	.8647	.8672	.8627	.8642

TABLE 10.B. ?? maxcut instances with  $\lceil Q \ln m \rceil$  cuts sampled uniformly, and  $\lceil (128 - Q) \ln m \rceil$  cuts sampled from SDP solution  $Y$  from (28), with  $\varepsilon = 10^{-4}$  and for each  $Q \in \{2, 8, 16, 32\}$ . See Table Description 6.

instance	$\varepsilon = 1/64$	$Q = 2$	$\varepsilon = 1/16$	$Q = 8$	$\varepsilon = 1/8$	$Q = 16$	$\varepsilon = 1/4$	$Q = 32$
	$\beta_{\text{fcc}}$	$\beta_{\text{fcc}}$	$\beta_{\text{fcc}}$	$\beta_{\text{fcc}}$	$\beta_{\text{fcc}}$	$\beta_{\text{fcc}}$	$\beta_{\text{fcc}}$	$\beta_{\text{fcc}}$
bayg29	.8845	.8852	.8849	.8852	.8852	.8854	.8853	.8855
bays29	.8868	.8863	.8870	.8867	.8863	.8863	.8867	.8870
berlin52	.8872	.8899	.8876	.8899	.8880	.8899	.8790	.8899
brazil58	.8883	.8892	.8879	.8892	.8892	.8892	.8890	.8892
kroA100	.9998	.9999	.9999	.9999	.9999	.9999	.9999	.9999
gr96	.8832	.8810	.8830	.8815	.8828	.8820	.8758	.8816
eil101	.8782	.8785	.8780	.8784	.8771	.8784	.8702	.8783
gr120	.9700	.9701	.9681	.9701	.9699	.9701	.9703	.9701
bier127	.8777	.8792	.8673	.8793	.8601	.8791	.8440	.8792
ch130	.8781	.8784	.8781	.8783	.8771	.8784	.8695	.8780
gr137	.9873	.9864	.9776	.9864	.9870	.9864	.9867	.9864
ch150	.8783	.8781	.8771	.8782	.8779	.8779	.8664	.8777
brg180	.8846	.8871	.8742	.8869	.8605	.8862	.8283	.8851
d198	.8865	.8810	.8868	.8807	.8830	.8815	.8645	.8815
gr202	.8702	.8784	.8706	.8782	.8525	.8780	.8448	.8780
a280	.8795	.8792	.8794	.8794	.8771	.8793	.8695	.8792

TABLE 11.A. ?? paired fractional cut-covering instances with  $\lceil Q \ln m \rceil$  cuts sampled uniformly, and  $\lceil (128 - Q) \ln m \rceil$  cuts sampled from SDP solution  $Y$  for (33), with  $\varepsilon = 10^{-4}$  and for each  $Q \in \{2, 8, 16, 32\}$ . See Table Description 7.

instance	$\varepsilon = 1/64$	$Q = 2$	$\varepsilon = 1/16$	$Q = 8$	$\varepsilon = 1/8$	$Q = 16$	$\varepsilon = 1/4$	$Q = 32$
	$\beta_{\text{fcc}}$	$\beta_{\text{fcc}}$	$\beta_{\text{fcc}}$	$\beta_{\text{fcc}}$	$\beta_{\text{fcc}}$	$\beta_{\text{fcc}}$	$\beta_{\text{fcc}}$	$\beta_{\text{fcc}}$
ENZYMES8	.9002	.9104	.9037	.9073	.9042	.9134	.8908	.9113
soc-dolphins	.8908	.8908	.8899	.8908	.8908	.8908	.8850	.8908
road-chesapeake	.8883	.8879	.8888	.8879	.8885	.8879	.8875	.8879
email-enron-only	.8883	.8787	.8856	.8787	.8733	.8787	.8440	.8787
dwt_209	.8915	.8910	.8881	.8910	.8753	.8910	.8425	.8910
ca-netscience	.8794	.8793	.8796	.8793	.8619	.8827	.8184	.8815
Erdos991	.8796	.8669	.8683	.8659	.8523	.8656	.8161	.8638
hamming6-2	.9996	.9999	.9694	.9999	.9175	.9999	.8671	.9999
inf-USAir97	.8779	.8776	.8768	.8776	.8754	.8776	.8566	.8778
ia-infect-hyper	.8870	.8891	.8785	.8888	.8673	.8882	.8408	.8868
DD687	.8755	.8809	.8563	.8804	.8385	.8794	.7765	.8773
dwt_503	.8861	.8858	.8818	.8859	.8655	.8859	.8291	.8859
ia-infect-dublin	.8776	.8814	.8642	.8810	.8463	.8800	.8092	.8781
email-univ	.8666	.8624	.8504	.8619	.8313	.8607	.7920	.8579
johnson16-2-4	.9310	.9319	.9264	.9306	.9200	.9289	.9081	.9244
p-hat700-1	.8647	.8699	.8523	.8692	.8365	.8678	.7671	.8649

TABLE 11.B. ?? paired fractional cut-covering instances with  $\lceil Q \ln m \rceil$  cuts sampled uniformly, and  $\lceil (128 - Q) \ln m \rceil$  cuts sampled from SDP solution  $Y$  for (33), with  $\varepsilon = 10^{-4}$  and for each  $Q \in \{2, 8, 16, 32\}$ . See Table Description 7.

## 6. SOLVERS

The SDP solver plays a dual role in our pipeline: it provides rigorous simultaneous certificates for both maximum cut and fractional cut-covering problems, while also determining the overall computational cost. As illustrated in Figures 2 and 3, the solver typically accounts for the vast majority of the running time. Given this sensitivity to the solver’s performance, it is natural to ask whether the improved accuracy in the SDP solution and more robust behaviour by the SDP solver would have a significant positive effect on the performance of our pipeline. We therefore reran our experiments using the second-order interior-point method provided by Domain Driven Solver (DDS) version 2.2 [22]. We found that on our SDP instances, DDS was more robust than SeDuMi (as explained in Section C). Another motivation for this choice is forward-looking: our theoretical results extend beyond semidefinite programming and fractional cut covers [6], and DDS offers infrastructure for these conic generalizations of covering problems. Since second-order methods are inherently more memory-intensive, not all instances were solved with our computational resources. To accommodate the increased memory footprint without resorting to the noisy, shared environment of ?? hardware, we performed our solver comparison experiments on ?? hardware.

We revisit some details of our implementation using SCS to provide context for the upcoming results with DDS. Let  $\mathbb{E}$  and  $\mathbb{Y}$  be Euclidean spaces. Let  $\mathcal{A}: \mathbb{E} \rightarrow \mathbb{Y}$  be a linear transformation, let  $b \in \mathbb{Y}$ , let  $c \in \mathbb{E}$ , and let  $K \subseteq \mathbb{Y}$  be a closed convex cone. In the linear conic case considered here, the standard form of an optimization problem solved by SCS is

$$(39) \quad \min\{ \langle c, x \rangle : x \in \mathbb{E}, s \in K, \mathcal{A}(x) + s = b \} \geq \max\{ \langle -b, y \rangle : y \in K^*, -\mathcal{A}^*(y) = c \}.$$

SCS reports an instance to be infeasible according to a parameter  $\varepsilon_{\text{infeas}}$ , whose default value is  $10^{-7}$ . For our formulation of the maxcut SDP (2), we set  $\mathbb{E} := \mathbb{R}^V$ ,  $\mathbb{Y} := \mathbb{S}^V$ ,

$$(40) \quad \mathcal{A} := -\text{Diag}, \quad b := -\frac{1}{4}\mathcal{L}_G(w), \quad c := \mathbb{1}, \quad \text{and} \quad K := \mathbb{S}_+^V.$$

For some instances, SCS reported the SDP (2b) as infeasible or, equivalently, reported (2a) as unbounded. For (41), SCS reports infeasibility by producing  $Y \in \mathbb{S}_+^V$  such that

$$(41) \quad \langle \frac{1}{4}\mathcal{L}_G(w), Y \rangle = 1 \text{ and } \|\text{diag}(Y)\|_\infty < \varepsilon_{\text{infeas}}.$$

Such a certificate exists if and only if  $\eta(G, w) > \varepsilon_{\text{infeas}}^{-1}$ , and for every such instance  $(G, w)$ , SCS may report (2b) as infeasible. In such cases, the output  $Y \in \mathbb{S}_+^V$  meets no stopping criteria other than having objective value larger than  $\varepsilon_{\text{infeas}}^{-1}$ . The maximum cut instances `bier127`, `brg180`, `d198`, `gr137`, `gr202`, and `gr96` have cuts of value larger than  $10^7 = \varepsilon_{\text{infeas}}^{-1}$ , and in our experiments, SCS reports all of them as infeasible. Mirka and Williamson [31] do not address this issue, simply using  $Y$  to sample cuts. Since  $\|w\|_1 \geq \eta(G, w)$  and an SCS infeasibility certificate satisfying (42) implies  $\eta(G, w) > \varepsilon_{\text{infeas}}^{-1}$ , no such certificate exists if  $\|w\|_1 < \varepsilon_{\text{infeas}}^{-1}$ . Hence, scaling the input so that  $\|w\|_1 = 1$  sidesteps this problem.

Whereas normalization imposes structure on the inputs provided to SCS, Algorithms 1 and 2 impose structure on its output, sanitizing it into a standardized numerical certificate (20). Importantly, we can easily try SCS, DDS, and SeDuMi precisely due to (20) being solver agnostic. However, sanitization is necessary even when using a single solver. Let  $(V, E) = C_{301}$  be the cycle on 301 vertices, and set  $w := \mathbb{1} \in \mathbb{R}^E$ . Early experiments solving (2b) had SCS terminating with  $\tilde{x} \in \mathbb{R}^V$  such that  $\mathbb{1}^\top \tilde{x} < 300 = \text{mc}(C_{301}, w)$ . This is impossible if  $\text{Diag}(\tilde{x}) \succeq \frac{1}{4}\mathcal{L}_G(w)$ , but does not contradict the termination criteria for SCS, which merely ensures an upper bound on  $\|\text{Diag}(\tilde{x}) - \frac{1}{4}\mathcal{L}_G(w) - S\|_{\text{SCS}}$  for a norm  $\|\cdot\|_{\text{SCS}}$  on  $\mathbb{S}^V$ . This allows  $\text{Diag}(\tilde{x}) - \frac{1}{4}\mathcal{L}_G(w)$  to be sufficiently far from any positive semidefinite matrix, and thus compromises any attempt at certification which does not sanitize the solver output. We refer the reader to Section B.1 for the stopping criteria (57) used by SCS when given the formulation (41), as well as for the definition of  $\|\cdot\|_{\text{SCS}}$  in (56).

We refrained from exploring distinct formulations of the relevant SDPs (namely, (2) and (33)) to preserve the scope of this work. For each solver, our formulations are straightforward implementations of (2) and (33), with the caveat that we selected the primal problem so as to reduce the number of decision variables. For example, to solve (2), we formulated (2a) as primal in SeDuMi, and (2b) as primal in SCS and DDS. Normalizing the input so that  $\|w\|_1 = 1$  for every maximum cut instance  $(G, w)$  and  $\|z\|_\infty = 1$  for every fractional cut covering instance  $(G, z)$  was the only data pre-processing used.

**6.1. DDS on Maximum Cut Instances.** Having discussed important ways in which our pipeline interacts with the SDP solver selected, we now present the standard formulation used by DDS. With notation as in (41), DDS solves

$$(42) \quad \inf\{\langle c, x \rangle : \mathcal{A}(x) + b \in K\} + \inf\{-\langle y, b \rangle : y \in K^\circ, \mathcal{A}^*(y) = -c\} \geq 0,$$

where  $K^\circ$  is the polar cone of  $K$ . The first infimum is the primal problem, and the second infimum is the dual problem. Assuming a strictly feasible point  $x^{(0)}$  is provided as the initial point, upon successful termination, DDS outputs  $(\tilde{x}, \tilde{y})$  such that

$$(43a) \quad \mathcal{A}(\tilde{x}) + b \in \text{int}(K), \quad \tilde{y} \in \text{int}(K^\circ),$$

$$(43b) \quad |\langle c, \tilde{x} \rangle - \langle \tilde{y}, b \rangle| \leq \varepsilon_{\text{DDS}}(1 + |\langle c, \tilde{x} \rangle| + |\langle \tilde{y}, b \rangle|),$$

$$(43c) \quad \|\mathcal{A}^*(\tilde{y}) + c\| \leq \varepsilon_{\text{DDS}}(1 + \|c\|)$$

where  $\varepsilon_{\text{DDS}} \in (0, 1)$  is a user-specified tolerance. The default choice is  $\varepsilon_{\text{DDS}} = 10^{-8}$ . We provided strictly feasible solutions as starting points for all of our formulations in this manuscript.

Our formulation of (2) takes the form (43) with  $K := \mathbb{S}_+^V$ , and  $\mathcal{A}: \mathbb{R}^V \rightarrow \mathbb{S}^V$  such that

$$\mathcal{A}(x) := \text{Diag}(x), \quad b := -\frac{1}{4}\mathcal{L}_G(w), \quad \text{and} \quad c := \mathbb{1}.$$

In particular, upon successful termination, we have (after a change of variable)  $x \in \mathbb{R}^V$  and  $Y \in \mathbb{S}^V$  such that

$$(44a) \quad \text{Diag}(x) - \frac{1}{4}\mathcal{L}_G(w) \in \text{int}(\mathbb{S}_+^V), \quad Y \in \text{int}(\mathbb{S}_+^V),$$

$$(44b) \quad |\langle \mathbb{1}, x \rangle - \frac{1}{4}\langle \mathcal{L}_G(w), Y \rangle| \leq \varepsilon_{\text{DDS}}(1 + |\langle \mathbb{1}, x \rangle| + |\frac{1}{4}\langle \mathcal{L}_G(w), Y \rangle|),$$

$$(44c) \quad \|\text{diag}(Y) - \mathbb{1}\| \leq \varepsilon_{\text{DDS}}(1 + \sqrt{n}).$$

Whereas for all experiments in this section we work with  $\varepsilon_{\text{DDS}} = 10^{-8}$ , the paired instances computed in Algorithm 3 used feasible solutions that satisfy the above constraints with  $\varepsilon_{\text{DDS}} = 10^{-12}$ .

**Table Description 8** (Tables 12.A and 12.B). Experiments obtained on ?? hardware. As input, we are given maximum cut instances  $(G, w)$ . In Table 12.A they are ?? instances, and in Table 12.B they are ?? instances. We use DDS on the SDPs in (2), obtaining  $\tilde{x} \in \mathbb{R}^V$  and  $\tilde{Y} \in \mathbb{S}^V$ , from which we set  $Y := (1 - \varepsilon)\tilde{Y} + \varepsilon I$  as in (28) and  $z := \frac{1}{4}\mathcal{L}_G^*(Y)$  as in STEP 1. Set  $(\rho, x, B) := \text{SLACKSANITIZE}(\tilde{x}, w)$ , set  $(\mu, R \in \mathbb{R}^{[k] \times V}) := \text{REPRESENTATIONSANITIZE}(Y, z)$ . We store  $w, z, \rho, \mu, x, R$ , and  $B$  in a file. This step is done once for each  $\varepsilon \in \{0, 10^{-8}, 1/64\}$  and input graph  $(G, w)$ .

For each such file, we run our pipeline with 10 different seeds used to sample cuts. In these runs, we first load our SDP certificates from the file. We sample  $\lceil 128 \log m \rceil$  vectors  $g \in \mathbb{R}^k$  according to the standard multivariate normal distribution, and produce a shore  $S := \{i \in V : g^\top R e_i > 0\}$  for each  $g$ . Let  $\mathcal{F} \subseteq \mathcal{P}(V)$  be the set of shores generated. We compute  $\text{mc}(\mathcal{F}, w)$  directly, and solve  $\text{fcc}(\mathcal{F}, z)$  using Gurobi. We then define  $\beta_{\text{mc}}$ , and  $\beta_{\text{fcc}}$  as in (32). Statistics for  $\sigma$  and  $\beta_{\text{mc}}$  use all 10 runs; those for  $\beta_{\text{fcc}}$  use only runs that produced a feasible fractional cut cover for  $z$ , thus excluding the infeasible runs counted in the RLI column; tables in which all runs produced a feasible solution omit this column.

$\varepsilon$	0	$10^{-8}$	$1/64$	0	$10^{-8}$	$1/64$	0	0	$10^{-8}$	$10^{-8}$	$1/64$
instance	$1 - \sigma$	$1 - \sigma$	$1 - \sigma$	$\beta_{mc}$ avg	$\beta_{mc}$ avg	$\beta_{mc}$ avg	RLI	$\beta_{fcc}$ avg	RLI	$\beta_{fcc}$ avg	$\beta_{fcc}$ avg
bayg29	.9999	.9999	.9963	.9938	.9938	.9938	8	.8848	8	.8848	.8880
bays29	.9999	.9999	.9963	.9943	.9943	.9943	6	.8845	6	.8845	.8895
berlin52	.9999	.9999	.9968	.9904	.9904	.9904	0	.8900	0	.8900	.8919
brazil158	.9999	.9999	.9961	.9996	.9996	.9996	10	–	10	–	.8905
gr96	.9999	.9999	.9961	.9985	.9985	.9985	10	–	10	–	.8857
kroA100	.9999	.9999	.9955	.9999	.9999	.9999	5	.9999	5	.9999	.9926
eil101	.9999	.9999	.9965	.9813	.9813	.9813	8	.8785	8	.8785	.8813
gr120	.9999	.9999	.9956	.9998	.9998	.9998	10	–	10	–	.9702
bier127	.9999	.9999	.9972	.9753	.9753	.9752	0	.8791	0	.8791	.8818
ch130	.9999	.9999	.9965	.9881	.9881	.9881	10	–	10	–	.8813
gr137	.9999	.9999	.9954	.9999	.9999	.9999	10	–	10	–	.9863
ch150	.9999	.9999	.9965	.9877	.9877	.9877	10	–	10	–	.8813
brg180	.9999	.9999	.9974	.9438	.9438	.9453	10	–	10	–	.8882
d198	.9999	.9999	.9956	.9985	.9985	.9985	10	–	10	–	.8891
gr202	.9999	.9999	.9971	.9783	.9783	.9782	0	.8783	0	.8783	.8796
a280	.9999	.9999	.9958	.9995	.9995	.9995	10	–	10	–	.8818

TABLE 12.A. Dense maximum cut instances solved with DDS  $[128 \ln m]$  cut sampled. See Table Description 8. The largest standard deviation for  $\beta_{mc}$  was 0.0016, and the largest for  $\beta_{fcc}$  was 0.0003. No experiment with  $\varepsilon = 1/64$  had an (RLI) outcome.

$\varepsilon$	0	$10^{-8}$	$1/64$	0	$10^{-8}$	$1/64$	0	0	$10^{-8}$	$10^{-8}$	$1/64$
instance	$1 - \sigma$	$1 - \sigma$	$1 - \sigma$	$\beta_{mc}$ avg	$\beta_{mc}$ avg	$\beta_{mc}$ avg	RLI	$\beta_{fcc}$ avg	RLI	$\beta_{fcc}$ avg	$\beta_{fcc}$ avg
ENZYMES8	.9999	.9999	.9924	.9795	.9795	.9795	0	.8873	0	.8873	.9034
soc-dolphins	.9999	.9999	.9942	.9744	.9744	.9720	0	.8911	0	.8911	.8906
road-chesapeake	.9999	.9999	.9945	.9679	.9679	.9679	0	.8888	0	.8888	.8888
email-enron-only	.9999	.9999	.9953	.9601	.9601	.9594	0	.8890	0	.8890	.8896
dwt_209	.9999	.9999	.9948	.9682	.9682	.9680	0	.8917	0	.8917	.8917
ca-netscience	.9999	.9999	.9952	.9682	.9682	.9666	10	–	10	–	.8814
Erdos991	.9999	.9999	.9945	.9488	.9488	.9468	0	.8876	0	.8876	.8841
hamming6-2	.9999	.9999	.9986	.9999	.9999	.9999	0	.9999	0	.9999	.9983
inf-USAir97	.9999	.9999	.9951	.9731	.9731	.9731	8	.8797	8	.8797	.8808
ia-infect-hyper	.9999	.9999	.9973	.9710	.9710	.9707	0	.8902	0	.8902	.8903
DD687	.9999	.9999	.9950	.9441	.9441	.9411	7	.8815	7	.8815	.8782
dwt_503	.9999	.9999	.9953	.9876	.9876	.9867	10	–	10	–	.8892
ia-infect-dublin	.9999	.9999	.9961	.9515	.9515	.9509	0	.8827	0	.8827	.8803
email-univ	.9999	.9999	.9950	.9306	.9306	.9274	0	.8734	0	.8734	.8690
johnson16-2-4	.9999	.9999	.9979	.9717	.9720	.9711	0	.9329	0	.9327	.9308
p-hat700-1	.9999	.9999	.9982	.9725	.9725	.9720	0	.8703	0	.8703	.8693

TABLE 12.B. Sparse maximum cut instances solved with DDS and  $[128 \ln m]$  cut sampled. See Table Description 8. The largest standard deviation for  $\beta_{mc}$  was 0.0038, and the largest for  $\beta_{fcc}$  was 0.0007. No experiment with  $\varepsilon = 1/64$  had an (RLI) outcome.

We now consider Tables 12.A and 12.B, which replicate the experiments from Section 3 using the DDS solver and sampling  $\lceil 128 \ln m \rceil$  cuts. As before, we evaluate three values of the perturbation parameter  $\varepsilon$ : no perturbation, a perturbation equal to the solver’s default tolerance, and  $1/64$ . For DDS, whose default tolerance is  $10^{-8}$ , this corresponds to  $\varepsilon \in \{0, 10^{-8}, 1/64\}$ .

We see improved pairing quality (i.e., higher  $1 - \sigma$ ) when comparing  $\varepsilon = 10^{-8}$  in Tables 12.A and 12.B with  $\varepsilon = 10^{-4}$  in Tables 1.A and 1.B. In other words, whenever perturbation matches solver accuracy, we obtain better pairings. This is to be expected, as the second-order method produces more accurate solutions, and this directly impacts the value of  $\sigma$ . For the same perturbation value of  $\varepsilon = 1/64$ , we can compare Tables 12.A and 12.B to Tables 3.A and 3.B and see that we also consistently get better values of  $\sigma$ .

The only instances in Table 12.A which had a worse  $\beta_{\text{mc}}$  value for  $\varepsilon = 1/64$  than for  $\varepsilon = 0$  were `bier127` and `gr202`, but even for those instances the change was only on the fourth decimal place. This fits within a common pattern observed in previous sections, where perturbation has little impact on  $\beta_{\text{mc}}$  for ?? instances. Also analogous to previous results, Table 12.B has a clearer impact of higher perturbation on the value of  $\beta_{\text{mc}}$ . More interestingly, using DDS has an overall positive impact on  $\beta_{\text{mc}}$  compared to SCS. There are three instances (`soc-dolphins`, `email-enron-only`, and `johnson-16-2-4`) where, with  $\varepsilon = 10^{-4}$ , SCS obtained a better  $\beta_{\text{mc}}$  than DDS with  $\varepsilon = 1/64$ . This is noticeable in that, even with much higher perturbation, DDS still produces better  $\beta_{\text{mc}}$  value than SCS for most instances.

As in Section 3, our algorithm was better at producing fractional cut-covers for ?? instances than for ?? instances. This difference is most evident on the RLI count in Tables 12.A and 12.B. In contrast, DDS always produced certificates accepted by our sanitization procedure, even when  $\varepsilon = 0$ . Tables 12.A and 12.B do show that choosing  $\varepsilon = 1/64$  has a positive impact on the outcome of our procedure, once again producing feasible fractional cut covers reliably. For the same perturbation of  $\varepsilon = 1/64$ , there were 5 instances where  $\beta_{\text{fcc}}$  was higher when using SCS compared to DDS: `berlin52`, `brazil58`, `ENZYMES8`, `dwt_209`, and `ca-netscience`. Thus, in general, DDS produced better  $\beta_{\text{fcc}}$  results.

**6.2. DDS on Fractional Cut-Covering Instances.** We now turn our attention to experiments which take as input a fractional cut-covering instance. We refer the reader to Section B for the exact formulation of (33) using DDS. The higher memory requirements for DDS prevented us from solving the fractional cut-covering instances obtained via pairing from `a280` and `p-hat700-1`. DDS took considerably longer to solve (33) for the ?? instances as opposed to the ?? instances; see Tables 16.A and 16.B. This is the opposite behavior to SCS, which struggled more with the sparse instances (Table 6.B) than it did with dense instances (Table 6.B).

**Table Description 9** (Tables 13.A and 13.B). Experiments obtained on ?? hardware. As input, we are given fractional cut-covering instances  $(G, z)$ . In Table 13.A they are paired ?? instances, and in Table 13.B they are paired ?? instances — see Section 4. We use DDS on the SDPs in (33), obtaining  $\tilde{\mu} \in \mathbb{R}_+$ ,  $\tilde{Y} \in \mathbb{S}^V$ ,  $w \in \mathbb{R}_+^E$ , and  $\tilde{x} \in \mathbb{R}^V$ . Set  $Y := \tilde{Y} + \varepsilon \tilde{\mu} I$ , set  $(\rho, x, B) := \text{SLACKSANITIZE}(\tilde{x}, w)$ , and set  $(\mu, R \in \mathbb{R}^{|k| \times V}) := \text{REPRESENTATIONSANITIZE}(Y, z)$ . We store  $w, z, \rho, \mu, x, R$ , and  $B$  in a file. This step is done once for each choice of  $\varepsilon \in \{0, 10^{-8}, 1/64\}$  and input graph  $(G, z)$ .

For each file produced in such a manner, we run our pipeline with 10 different seeds used to sample cuts. In these runs, we first load our SDP certificates from the file. We sample  $\lceil 128 \log m \rceil$  vectors  $g \in \mathbb{R}^k$  according to the standard multivariate normal distribution, and produce a shore  $S := \{i \in V : g^\top R e_i > 0\}$  for each  $g$ . Let  $\mathcal{F} \subseteq \mathcal{P}(V)$  be the set of shores generated. We compute  $\text{mc}(\mathcal{F}, w)$  directly, and solve  $\text{fcc}(\mathcal{F}, z)$  using Gurobi. We then define  $\beta_{\text{mc}}$ , and  $\beta_{\text{fcc}}$  as in (34). Statistics for  $\sigma$  and  $\beta_{\text{mc}}$  use all 10 runs; those for  $\beta_{\text{fcc}}$  use only runs that produced a feasible fractional cut cover for  $z$ , thus excluding the infeasible runs counted in the RLI column; tables in which all runs produced a feasible solution omit this column.

Tables 13.A and 13.B replicate the experiments from Section 4 on paired fcc instances, and are analogous to Tables 7.A and 7.B, except that we additionally include the case  $\varepsilon = 0$ , and use the DDS default tolerance  $\varepsilon = 10^{-8}$  (instead of  $10^{-4}$  for SCS). Once again, no perturbation and low perturbation lead to bad outcomes when computing fractional cut covers, which are evident by the high RLI counts. Similar to the results with SCS, ?? instances were harder to cover than the ?? instances. For  $\varepsilon = 1/64$ , using DDS rather than SCS produces better certifiable bounds  $\beta_{\text{fcc}}$  and  $\beta_{\text{mc}}$ . The instances gr96 and ch130 are the only exceptions, where the estimated value of  $\beta_{\text{mc}}$  is slightly worse in Table 13.A than in Table 7.A.

A noticeable difference lies in the presence of remarkably low  $1 - \sigma$  values in instances kroA100, gr96, brg180, and ca-netscience in Tables 13.A and 13.B. In all of these cases, DDS produced certificates which were accepted by our sanitization procedure, albeit with high values of  $\mu$  computed by Algorithm 2. Recalling (34), higher values of  $\mu$  directly increase the value of  $\sigma$ . These instances show that even when using second-order methods perturbation may still help improve the behaviour of our algorithms.

$\varepsilon$	0	$10^{-8}$	$1/64$	0	0	$10^{-8}$	$10^{-8}$	$1/64$	0	$10^{-8}$	$1/64$
instance	$1 - \sigma$	$1 - \sigma$	$1 - \sigma$	RLI	$\beta_{fcc}$ avg	RLI	$\beta_{fcc}$ avg	$\beta_{fcc}$ avg	$\beta_{mc}$ avg	$\beta_{mc}$ avg	$\beta_{mc}$ avg
bayg29	.9996	.9994	.9921	9	.8848	7	.8845	.8854	.9951	.9921	.9999
bays29	.9995	.9997	.9921	9	.8845	7	.8847	.8870	.9939	.9952	.9999
berlin52	.9999	.9999	.9922	0	.8900	0	.8900	.8900	.9920	.9927	.9989
brazil158	.9999	.9999	.9921	10	–	10	–	.8892	.9987	.9986	.9999
kroA100	.0744	.1044	.9921	9	.9999	9	.9999	.9999	.9999	.9999	.9999
gr96	.8212	.8408	.9921	10	–	10	–	.8834	.9979	.9979	.9996
eil101	.9995	.9996	.9921	8	.8785	9	.8785	.8786	.9849	.9823	.9991
gr120	.9971	.9949	.9921	10	–	10	–	.9703	.9999	.9999	.9999
bier127	.9999	.9999	.9921	0	.8791	0	.8791	.8785	.9798	.9808	.9969
ch130	.9682	.9790	.9921	10	–	10	–	.8786	.9904	.9881	.9981
gr137	.9419	.9615	.9921	10	–	10	–	.9873	.9999	.9999	.9999
ch150	.9999	.9999	.9921	10	–	10	–	.8786	.9872	.9871	.9984
brg180	.8302	.8525	.9925	10	–	10	–	.8857	.9685	.9694	.9812
d198	.9982	.9986	.9921	10	–	10	–	.8878	.9981	.9981	.9993
gr202	.9991	.9993	.9921	0	.8782	0	.8783	.8764	.9821	.9818	.9977

TABLE 13.A. Dense paired fractional cut-covering instances solved on ?? hardware with DDS and  $\varepsilon \in \{0, 10^{-8}, 1/64\}$  and  $\lceil 128 \ln m \rceil$  cut sampled. See Table Description 9.

$\varepsilon$	0	$10^{-8}$	$1/64$	0	0	$10^{-8}$	$10^{-8}$	$1/64$	0	$10^{-8}$	$1/64$
instance	$1 - \sigma$	$1 - \sigma$	$1 - \sigma$	RLI	$\beta_{fcc}$ avg	RLI	$\beta_{fcc}$ avg	$\beta_{fcc}$ avg	$\beta_{mc}$ avg	$\beta_{mc}$ avg	$\beta_{mc}$ avg
ENZYMES8	.9999	.9999	.9921	0	.8873	0	.8886	.9041	.9805	.9789	.9999
soc-dolphins	.9999	.9999	.9921	0	.8911	0	.8911	.8911	.9712	.9690	.9999
road-chesapeake	.9999	.9999	.9921	0	.8888	0	.8888	.8889	.9660	.9650	.9999
email-enron-only	.9999	.9999	.9921	0	.8890	0	.8890	.8890	.9574	.9579	.9999
dwt_209	.9998	.9996	.9921	0	.8917	0	.8917	.8915	.9680	.9684	.9999
ca-netscience	.6347	.9863	.9921	10	–	10	–	.8800	.9679	.9674	.9998
Erdos991	.9998	.9999	.9921	0	.8874	0	.8875	.8835	.9509	.9502	.9999
hamming6-2	.9999	.9999	.9921	0	.9999	0	.9999	.9999	.9999	.9999	.9999
inf-USAir97	.9895	.9868	.9921	8	.8797	9	.8797	.8795	.9735	.9738	.9997
ia-infect-hyper	.9999	.9999	.9921	0	.8902	0	.8902	.8873	.9618	.9618	.9999
DD687	.9996	.9997	.9921	6	.8819	6	.8818	.8776	.9439	.9440	.9999
dwt_503	.9790	.9801	.9921	10	–	10	–	.8871	.9860	.9863	.9998
ia-infect-dublin	.9998	.9999	.9921	0	.8826	0	.8827	.8785	.9477	.9475	.9999
email-univ	.9999	.9999	.9921	0	.8734	0	.8731	.8684	.9319	.9325	.9996
johnson16-2-4	.9999	.9999	.9979	0	.9330	0	.9324	.9313	.9718	.9718	.9717

TABLE 13.B. Sparse paired fractional cut-covering instances solved on ?? hardware with DDS and  $\varepsilon \in \{0, 10^{-8}, 1/64\}$  and  $\lceil 128 \ln m \rceil$  cut sampled. See Table Description 9.

## 7. COMPARISON WITH THE AVERAGING-AND-SCALING ALGORITHM

The framework introduced in [5] works as follow:

STEP 1: depending on whether the input is a maximum cut instance  $(G, w)$  or a fractional cut-covering instance  $(G, z)$ , solve either (2) or (12). Let  $Y \in \mathbb{S}_+^V$  and  $\mu \in \mathbb{R}_+$  be the nearly optimal for (2a) or (12a), whichever SDP was solved;

STEP 2: let  $S_1, \dots, S_T$  be a finite set of shores obtained from independent samples from  $\text{GW}(Y)$ ;

STEP 3: output a shore  $S \subseteq V$  that maximizes  $w^\top \mathbb{1}_{\delta(S)}$  over  $\mathcal{F} := \{S_1, \dots, S_T\}$  and output

$$y = \text{avg}((S_i)_{i \in [T]}, z) \hat{p}$$

where

$$\hat{p} := \frac{1}{T} \sum_{i=1}^T \mathbb{1}_{\delta(S_i)} \text{ and } \text{avg}((S_i)_{i \in [T]}, z) := \inf\{\mu \in \mathbb{R}_+ : \mu \hat{p} \geq z\}.$$

We refer to this approach as the *averaging-and-scaling algorithm*. We abuse notation and write  $\text{avg}(\mathcal{F}, z)$  whenever  $\mathcal{F} = \{S_1, \dots, S_T\}$  and the samples  $S_1, \dots, S_T$  are clear from context. [5] proves that, for every  $\beta \in (0, \alpha_{\text{GW}})$ , one may take  $T \in \Theta(\ln m)$  such that, with high probability,

$$\text{avg}(\mathcal{F}, z) \leq \frac{1}{\beta} \text{fcc}(G, z).$$

Together with the trivial bound  $\text{fcc}(\mathcal{F}, z) \leq \text{avg}(\mathcal{F}, z)$ , we obtain  $\text{fcc}(\mathcal{F}, z) \leq (1/\beta) \text{fcc}(G, z)$ . This section compares the averaging-and-scaling algorithm to our refinement based on the restricted LP (15). We first derive the value of  $\text{avg}((S_i)_{i \in [T]}, z)$  as the number of samples  $T$  approaches  $+\infty$ , which we call the *asymptotic objective value* of the averaging-and-scaling algorithm. We then present experiments measuring how many samples are needed for the LP-based approach to certify bounds better than the asymptotic objective value of the averaging-and-scaling algorithm. We also measure how many samples are needed for feasibility to be attained, and provide an overview of these variables for the dataset we have at hand.

Let  $G = (V, E)$  be a graph, and let  $z \in \mathbb{R}_+^E$ . Let  $Y \in \mathbb{S}_+^V$  and  $\mu \in \mathbb{R}_+$  be such that  $\text{diag}(Y) = \mu \mathbb{1}$  and  $\frac{1}{4} \mathcal{L}_G^*(Y) \geq z$ . Let  $p \in \mathbb{R}_+^E$  be given by  $p_{ij} = \mathbb{P}(ij \in \delta(\text{GW}(Y)))$  for every  $ij \in E$ . By the Strong Law of Large Numbers,

$$p = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{\delta(S_t)}$$

almost surely. Whenever this convergence happens, the objective value of the averaging-and-scaling algorithm converges to

$$(45) \quad \text{avg}(Y, z) := \max_{ij \in E} \frac{z_{ij}}{p_{ij}} = \max_{ij \in E} \frac{\pi z_{ij}}{\arccos(\mu^{-1} Y_{ij})}.$$

In particular, this objective value satisfies

$$\text{avg}(Y, z) \leq \frac{\mu}{\alpha_{\text{GW}}},$$

since  $\frac{1}{4} \mathcal{L}_G^*(Y) \geq z$  and the definition of  $\alpha_{\text{GW}}$  imply

$$\text{avg}(Y, z) = \max_{ij \in E} \frac{z_{ij} \pi}{\arccos(\mu^{-1} Y_{ij})} \leq \max_{ij \in E} \frac{\mu - Y_{ij}}{2} \frac{\pi}{\arccos(\mu^{-1} Y_{ij})} = \mu \max_{ij \in E} \frac{\pi}{2} \frac{1 - \mu^{-1} Y_{ij}}{\arccos(\mu^{-1} Y_{ij})} \leq \frac{\mu}{\alpha_{\text{GW}}}.$$

The value  $\text{avg}(Y, z)$  can be easily computed from  $Y$  and  $z$ , thus providing a meaningful alternative to running the averaging-and-scaling algorithm.

Let  $G = (V, E)$  be a graph, and let  $(z, w) \in H_\sigma(G)$  for a given  $\sigma \in (0, 1)$ . We shift our perspective on the algorithms described in Section 2: we now study them as stochastic processes, where at each step a new sample is taken and the values  $\text{mc}(\mathcal{F}, w)$  and  $\text{fcc}(\mathcal{F}, z)$  are computed for the set of shores  $\mathcal{F} \subseteq \mathcal{P}(V)$  sampled so far. Since  $(w, z) \in H_\sigma(G)$ , there exist  $(Y, \mu)$  feasible in (12a) and

$(\rho, x)$  feasible in (17a) such that  $w^\top z \geq (1 - \sigma)\rho\mu$ . Let  $(S_i(Y))_{i \in \mathbb{N}}$  be shores independently sampled from  $\text{GW}(Y)$ . For every  $t \in \mathbb{N}$ , set  $\mathcal{F}_t(Y) := \{S_i(Y) : i \in [t]\}$ . If context allows, we omit  $Y$  and write  $\mathcal{F}_t$  and  $S_t$  for  $t \in \mathbb{N}$ . We are interested in the following stopping times:

$$(46) \quad \begin{aligned} F &:= \inf\{t \in \mathbb{N} : \text{fcc}(\mathcal{F}_t(Y), z) < +\infty\}, \\ A &:= \inf\{t \in \mathbb{N} : (1 - \sigma) \text{avg}(Y, z) \geq \text{fcc}(\mathcal{F}_t(Y), z)\}. \end{aligned}$$

Both random variables describe properties of the restricted LP (15):  $F$  counts how many samples were taken until it was feasible, and  $A$  counts how many samples were taken until the resulting certificate had better value than the averaging-and-scaling algorithm. Indeed, recalling (18), the condition defining  $A$  is equivalent to

$$\beta_{\text{mc}} \geq \beta_{\text{fcc}} = \frac{(1 - \sigma)\mu}{\text{fcc}(\mathcal{F}_t(Y), z)} \geq \frac{\mu}{\text{avg}(Y, z)},$$

In experiments, we report

$$(47) \quad \tilde{F}_T := \frac{\min\{F, T\}}{\ln m} \quad \text{and} \quad \tilde{A}_T := \frac{\min\{A, T\}}{\ln m}.$$

The truncation by  $T$  arises from practical necessity, and normalization by  $\ln m$  allows comparison between instances. If context allows, we omit  $T$  and write  $\tilde{F}$  and  $\tilde{A}$ .

Before presenting experimental data, we establish some basic properties of the random variables  $F$  and  $A$ . Let  $G = (V, E)$  be a graph. We denote by  $\mathcal{L}_G(\mathbb{R}_+^E) := \{\mathcal{L}_G(w) : w \in \mathbb{R}_+^E\}$  the *Laplacian cone*, with all matrices arising as Laplacians of weights on the graph  $G$ . We have that  $\mathcal{L}_G(\mathbb{R}_+^E) \subseteq \mathbb{S}_+^V$ , and hence that

$$\mathbb{S}_+^V \subseteq (\mathcal{L}_G(\mathbb{R}_+^E))^* = \{Y \in \mathbb{S}^V : (e_i - e_j)^\top Y (e_i - e_j) \geq 0, \forall ij \in E\}.$$

These cones geometrically capture the asymptotic behaviour of our algorithms.

**Proposition 4.** Let  $G = (V, E)$  be a graph, let  $z \in \mathbb{R}_+^E$ , and let  $Y \in \mathbb{S}_+^V$ . If  $Y \in \text{int}((\mathcal{L}_G(\mathbb{R}_+^E))^*)$ , then

$$\lim_{t \rightarrow \infty} \text{fcc}(\mathcal{F}_t(Y), z) < +\infty$$

almost surely.

*Proof.* Let  $Y \in \text{int}(\mathcal{L}_G(\mathbb{R}_+^E)^*)$ , and set  $\rho := \min_{ij \in E} \langle \frac{1}{4}\mathcal{L}_G(e_{ij}), Y \rangle > 0$ . Let  $ij \in E$ . Since

$$\mathbb{P}(ij \in \delta(\text{GW}(Y))) \geq \alpha_{\text{GW}} \langle Y, \frac{1}{4}\mathcal{L}_G(e_{ij}) \rangle \geq \alpha_{\text{GW}} \rho > 0,$$

the Borel-Cantelli lemma implies  $ij \in \delta(S_t(Y))$  for infinitely many  $t \in \mathbb{N}$  almost surely. In particular,

$$\mathbb{P}(\forall t \in \mathbb{N}, \text{fcc}(\mathcal{F}_t(Y), z) = +\infty) \leq \mathbb{P}(\exists ij \in E \text{ s.t. } \forall t \in \mathbb{N}, ij \notin \delta(S_t(Y))) = 0.$$

Thus, almost surely, there exists  $t \in \mathbb{N}$  such that  $\text{fcc}(\mathcal{F}_t(Y), z)$  is finite. Whenever this happens,  $\lim_{t \rightarrow \infty} \text{fcc}(\mathcal{F}_t(Y), z)$  exists and is finite.  $\square$

**Proposition 5.** Let  $G = (V, E)$  be a graph, let  $z \in \mathbb{R}_+^E$ , and let  $Y \in \mathbb{S}_+^V$ . If  $Y \in \text{int}(\mathbb{S}_+^V)$ , then

$$\text{fcc}(G, z) = \lim_{t \rightarrow +\infty} \text{fcc}(\mathcal{F}_t(Y), z)$$

almost surely.

*Proof.* Let  $Y \in \mathbb{S}_+^V$ . Let  $S \subseteq V$  be such that  $\mathbb{P}(\text{GW}(Y) = S) = 0$ , and set  $s := 2\mathbb{1}_S - \mathbb{1}$ . We claim that

$$(48) \quad \{x \in \mathbb{R}^V : \text{Diag}(s)Y^{1/2}x > 0\} = \{x \in \mathbb{R}^V : \langle Y^{1/2}e_i, x \rangle s_i > 0, \forall i \in V\} = \emptyset.$$

Indeed, if we denote by  $\gamma$  the Gaussian measure on  $\mathbb{R}^V$ , we have

$$0 = \mathbb{P}(\text{GW}(Y) = S) \geq \gamma(\{h \in \mathbb{R}^V : \langle Y^{1/2}e_i, h \rangle s_i > 0, \forall i \in V\}) \geq 0.$$

Whence the set  $\{x \in \mathbb{R}^V : \langle Y^{1/2}e_i, x \rangle_{s_i} > 0, \forall i \in V\}$  is an open set with zero Gaussian measure, so (49) holds. Applying Gordan’s Lemma (see, e.g., [46, Section 7.8]) to (49), there exists  $y \in \mathbb{R}_+^V$  with  $y \neq 0$  such that  $Y^{1/2} \text{Diag}(s)y = (\text{Diag}(s)Y^{1/2})^\top y = 0$ . Since  $Y(\text{Diag}(s)y) = Y^{1/2}Y^{1/2} \text{Diag}(s)y = 0$  and  $\text{Diag}(s)y \neq 0$ , we conclude  $\text{rank}(Y) < n$ .

Now let  $Y \in \text{int}(\mathbb{S}_+^V)$ , and let  $S \subseteq V$ . Since  $\mathbb{P}(S = \text{GW}(Y)) > 0$ , the Borel-Cantelli lemma implies  $S_t(Y) = S$  for infinitely many  $t$  almost surely. Hence

$$\mathbb{P}(\forall t \in \mathbb{N}, \text{fcc}(\mathcal{F}_t(Y), z) \neq \text{fcc}(G, z)) \leq \mathbb{P}(\exists S \subseteq V \text{ s.t. } \forall t \in \mathbb{N}, S_t(Y) \neq S) = 0.$$

Thus, almost surely, there exists  $t \in \mathbb{N}$  such that  $\text{fcc}(\mathcal{F}_t(Y), z) = \text{fcc}(G, z)$ . Whenever this happens,  $\lim_{t \rightarrow \infty} \text{fcc}(\mathcal{F}_t(Y), z)$  exists and is equal to  $\text{fcc}(G, z)$ .  $\square$

Let  $G = (V, E)$  be a graph, let  $\sigma \in (0, 1)$ , and let  $(w, z) \in H_\sigma(G)$ . There exist  $(\mu, Y)$  feasible in (12a) and  $(\rho, x)$  feasible in (17a) such that  $(1 - \sigma)\rho\mu \leq w^\top z$ . Whereas Corollary 5 ensures that  $F$  is finite almost surely, Corollary 6 implies that, almost surely,  $\text{avg}(Y, z) \geq \text{fcc}(G, z) = \lim_{t \rightarrow \infty} \text{fcc}(\mathcal{F}_t(Y), z)$ , which is weaker than

$$(1 - \sigma) \text{avg}(Y, z) \geq \lim_{t \rightarrow \infty} \text{fcc}(\mathcal{F}_t(Y), z),$$

needed to prove that  $A$  is finite almost surely. This is natural: since  $A$  records the number of samples necessary to obtain a certificate better than  $\text{avg}(Y, z)$ , the pairing quality captured by  $\sigma$  plays a crucial role in its definition.

**7.1. Overview of  $\tilde{F}$  and  $\tilde{A}$  Behavior.** All figures in this subsection arise from the same set of experiments, which we now describe. The experiments were run on ?? hardware. As input, we are given maximum cut instances  $(G, w)$ , both ?? and ??. We use SCS to approximately solve the problems in (2), obtaining  $\tilde{x} \in \mathbb{R}^V$  and  $\tilde{Y} \in \mathbb{S}_+^V$ , from which we set  $Y := (1 - \varepsilon)\tilde{Y} + \varepsilon I$  as in (28), and  $z := \frac{1}{4}\mathcal{L}_G^*(Y)$ . Set  $(\rho, x, B) := \text{SLACKSANITIZE}(\tilde{x}, w)$  and  $(\mu, R \in \mathbb{R}^{[k] \times V}) := \text{REPRESENTATIONSANITIZE}(Y, z)$ . We sample  $T := \lceil 256 \ln m \rceil$  vectors  $g_t \in \mathbb{R}^k$  according to the standard multivariate normal distribution, and produce a shore  $S_t := \{i \in V : g_t^\top R e_i > 0\}$  for each  $g_t$ . For each  $t \in [T]$ , we use Gurobi to compute the value of  $\text{fcc}(\mathcal{F}_t, z)$ . This process is done once for each input graph, for each  $\varepsilon \in \{10^{-4}, 1/32, 3/64, 1/16\}$ , and for each of 10 different seeds used for the random number generation.

The good results in Tables 2.A and 2.B motivate sampling  $\lceil 256 \ln m \rceil$  cuts. Figure 5 provides an overview of the values of  $\tilde{F}$  and  $\tilde{A}$  in our dataset: each point represents an input instance, with the average value of  $\tilde{F}$  and  $\tilde{A}$  among the 10 runs providing its coordinates. None of the 4 plots in Figure 5 share the same  $x$ -axis. There is a clear difference between the behavior of the instances with  $\varepsilon = 10^{-4}$  and  $\varepsilon \in \{1/32, 3/64, 1/16\}$ . This reflects the choice of these values, with  $10^{-4}$  being a minimal perturbation. A distinctive feature of  $\varepsilon = 10^{-4}$  are the instances on the line  $\tilde{F} = \tilde{A}$ , reflecting that whenever feasibility was first attained, the objective value obtained was better than the averaging-and-scaling algorithm. We also have, on the line  $\tilde{A} = 256$ , both sparse and dense instances where the averaging-and-scaling algorithm would eventually provide better covers than the ones obtained after  $\lceil 256 \ln m \rceil$  samples. The comparison of  $\varepsilon = 10^{-4}$  with  $\varepsilon = \frac{1}{32}$  captures the effect of perturbation really well. Note that on the lower value of  $\varepsilon = 10^{-4}$ , we have values of  $\tilde{F}$  going up to 200, whereas on the higher value of  $\varepsilon = 1/32 = 0.03125$  there is no instance whose average  $\tilde{F}$  is larger than 12. Accordingly, the average values of  $\tilde{F}$  decrease with every subsequent increase in  $\varepsilon$ . The effect of  $\varepsilon$  on  $\tilde{A}$  can also be seen in Figure 5. Whereas for most instances, increasing  $\varepsilon$  decreases  $\tilde{A}$ , we can identify 7 instances in our dataset where higher perturbations either do not lower  $\tilde{A}$  or actually make them higher, and these instances become evident in the plot for  $\varepsilon = 0.0625 = \frac{1}{16}$ . They are: DD687, Erdos991, email-univ, ia-infect-dublin, inf-USAir97, johnson16-2-4, and p-hat700-1. Figure 6, instead of showing averaged values, shows the result of  $\tilde{A}$  and  $\tilde{F}$  for each run for a subset of the instances. In both Figures 5 and 6, the effect of increasing  $\varepsilon$  on  $\tilde{F}$  is visible.

Figure 7 display data for all 10 runs for 4 instances exhibiting distinct behaviour with respect to changes in perturbation. The color of each data point report its perturbation:  $\varepsilon = 1/32 = 0.03125$  in blue,  $\varepsilon = 3/64 = 0.046875$  in orange, and  $\varepsilon = 1/16 = 0.0625$  in green. Each of the 4 graphs have different axis. For **a280**, increasing the value of  $\varepsilon$  improves both  $\tilde{F}$  and  $\tilde{A}$ , as can be seen from the green points clustering at smaller coordinates the orange points, which themselves cluster at smaller coordinates than the blue points. The instance **gr120** is one where feasibility is hard to achieve, which is most evident by how many points were close to the line  $\tilde{A} = \tilde{F}$  when  $\varepsilon = 1/32$ . The change on  $\tilde{A}$  for increasing  $\varepsilon$  in this instances is hard to notice. For **dwt\_209**, the effect  $\varepsilon$  on  $\tilde{F}$  is negligible, whereas  $\tilde{A}$  increases for each increase in perturbation. For the instance **dwt\_503** there is a trade-off: perturbation improves  $\tilde{F}$  at the cost of worsening  $\tilde{A}$ .

Instances **gr120** and **dwt\_503** are particularly interesting in that the values observed for  $\tilde{F}$  are close to a probabilistic bound we can formally prove. Note that both instances predate works on fractional cut covers, and are in our data set because they were used in [31]. In short, neither **gr120** nor **dwt\_503** was created or selected due to how our algorithm behaves when solving it. Let  $G = (V, E)$  be a graph, and let  $\varepsilon \in (0, 1)$ . Let  $Y \in \mathbb{S}_+^V$  be such that  $\text{diag}(Y) = \mu \mathbb{1}$  and  $Y \succeq \varepsilon \mu I$  for  $\mu > 0$ . Let  $T \in \mathbb{N}$ , and let  $X := \sum_{t=1}^T \mathbb{1}_{\delta(S_t)}$ , for  $S_1, \dots, S_T$  independently sampled from  $\text{GW}(Y)$ . We claim that

$$(49) \quad \text{if } T \geq \frac{1}{-\ln(1 - \frac{\sqrt{2\varepsilon}}{\pi})} (\ln m + \ln(20)) \text{ then } \mathbb{P}(\exists ij \in E, X_{ij} = 0) \leq 0.05.$$

This result is a numerical improvement over Corollary 5 using a hypothesis stronger than the one in Corollary 6. Since  $Y \succeq \varepsilon \mu I$ , from [5, (63)] we have that

$$(50) \quad \mathbb{P}(ij \in \text{GW}(\mu^{-1}Y)) \geq \frac{\sqrt{2\varepsilon}}{\pi} \text{ for every } ij \in E.$$

The assumption on  $T$  may be rewritten as  $\ln(1/20) \geq \ln m + T \ln\left(1 - \frac{\sqrt{2\varepsilon}}{\pi}\right)$ , and hence (50) follows, since

$$\mathbb{P}(\exists ij \in E, X_{ij} = 0) \leq m \left(1 - \frac{\sqrt{2\varepsilon}}{\pi}\right)^T = \exp\left(\ln m + T \ln\left(1 - \frac{\sqrt{2\varepsilon}}{\pi}\right)\right) \leq \frac{1}{20}.$$

The instance **gr120** has  $m = 7140$  edges. If we divide the lower bound in (50) by  $\ln m$  so it matches the  $x$ -axis in Figure 7, we get

$$T/\ln m \geq 16.12 \text{ for } \varepsilon = 1/32, T/\ln m \geq 13.04 \text{ for } \varepsilon = 3/64, T/\ln m \geq 11.20 \text{ for } \varepsilon = 1/16.$$

For **dwt\_503**, we have that  $m = 126253$ , and that

$$T/\ln m \geq 15.13 \text{ for } \varepsilon = 1/32, T/\ln m \geq 12.24 \text{ for } \varepsilon = 3/64, T/\ln m \geq 10.51 \text{ for } \varepsilon = 1/16.$$

These bounds are comparable with the values of  $\tilde{F}$  for instances **gr120** and **dwt\_503** shown in Figure 7, substantiating the claim that this is an instance where feasibility is hard to achieve.

instance	$\tilde{A} = 256$ count	instance	$\tilde{A} = 256$ count
email-univ	40	eil101	10
johnson16-2-4	40	gr202	10
inf-USAir97	38	ch150	10
p-hat700-1	31	bier127	8
Erdos991	21	kroA100	1
DD687	20	d198	1
a280	10	gr96	1
ch130	10		

TABLE 14. Instances where we were unable to beat the averaging-and-scaling algorithm after  $\lceil 256 \ln m \rceil$  samples. Each instance was solved 10 times for each of 4 possible choices of  $\varepsilon$ .




FIGURE 5. Overview of  $\tilde{A}$  and  $\tilde{F}$ . Each point represents an instance, with the average value of  $\tilde{A}$  and  $\tilde{F}$  defining its coordinates. In orange we have ?? instances, whereas in blue we have ?? instances.

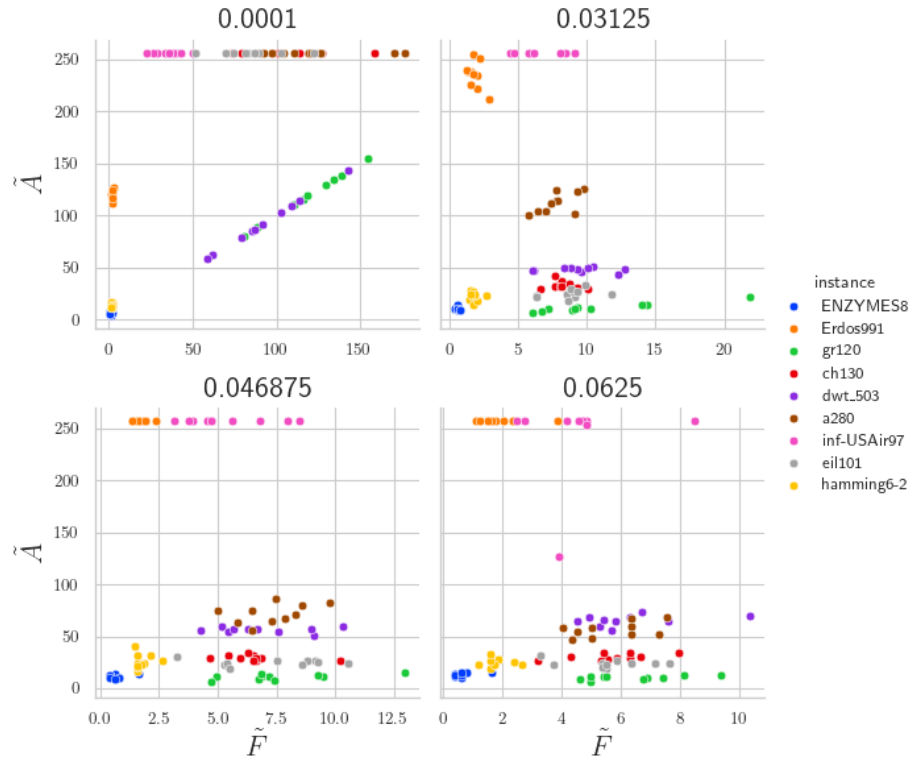


FIGURE 6. Each of the 10 recorded values of  $\tilde{A}$  and  $\tilde{F}$  for a subset of the instances.

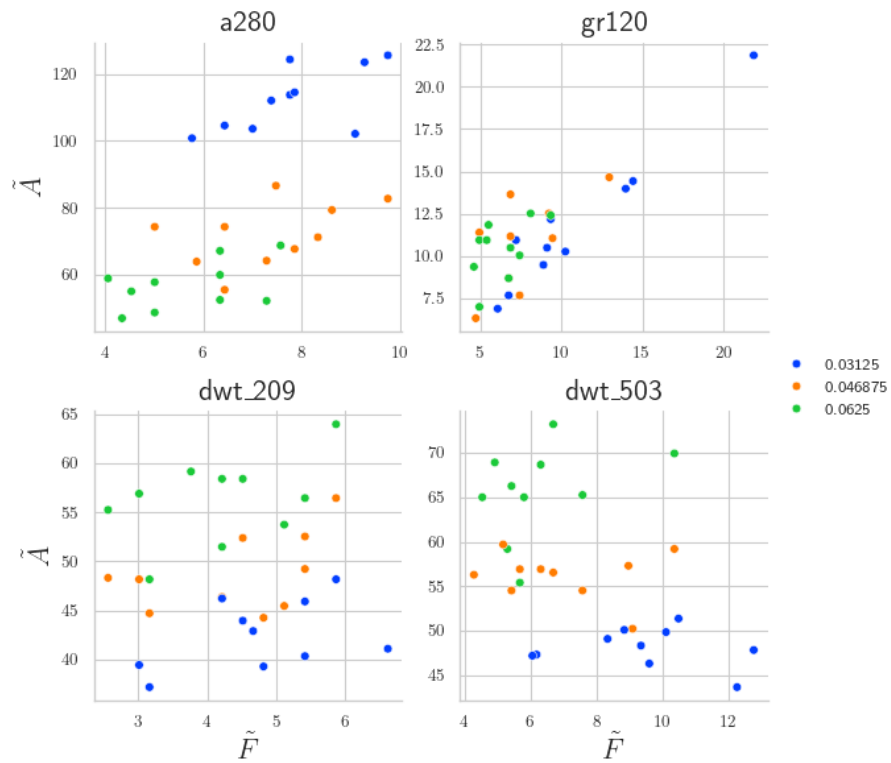


FIGURE 7. Interesting instances

**7.2. More Samples.** Table 14 singles out 3 instances where  $\lceil 256 \ln m \rceil$  cuts were not sufficient for the LP-based algorithm to certify a bound better than the asymptotic objective value of the averaging-and-scaling algorithm. They are `email-univ`, `johnson16-2-4`, and `inf-USAir97`. For the other instances, there was at least one choice of perturbation  $\varepsilon$  in which LP-solving consistently outperformed averaging-and-scaling. For the instances `email-univ` and `inf-USAir97` we repeat the experiments from the previous subsection, now sampling  $\lceil 1024 \ln m \rceil$ , whereas for `johnson16-2-4` we sample  $\lceil 2275 \ln m \rceil$  cuts. We plot the results in Figures 8 to 11.

Figures 8 to 11 show the value of the provable certificate for  $\beta_{\text{fcc}}$  obtained at each sample. More precisely, let  $G = (V, E)$  be a graph, and let  $w, z \in \mathbb{R}_+^E$ . Let  $(\rho, \mu, R, B, x)$  be a tuple satisfying (20), and set  $\sigma := 1 - z^\top w / (\rho\mu)$ . Let  $T \in \mathbb{N}$  denote the total number of cuts sampled. The figures present the value of the certificate quality for  $\beta_{\text{fcc}}$  obtained after each sample:

$$\beta_t := \frac{(1 - \sigma)\mu}{\text{fcc}(\mathcal{F}_t(Y), z)} \text{ for every } t \in 1, \dots, T := \lceil 1024 \ln m \rceil.$$

Given the small variation on the value of  $\beta_t$  across independent runs, we present the results of a single run. The curves are color coded by the perturbation value:  $\varepsilon = 10^{-4}$  in blue,  $\varepsilon = 1/32$  in orange,  $\varepsilon = 3/64$  in green, and  $\varepsilon = 1/16$  in red. Following the same color coding, the colorful dashed horizontal lines report the value of  $\text{avg}(Y_\varepsilon, z_\varepsilon)$  for  $\varepsilon \in \{10^{-4}, 1/32, 3/64, 1/16\}$ . Whereas the dashed lines overlap in Figures 8 and 9, they are clearly separated in Figures 10 and 11. The  $x$ -axis is scaled so it represents multiples of  $\ln m$ . For example, “400” indicates that  $400 \ln m$  cuts have been sampled. In Figures 8 and 9, we show two vertical black lines: one marking our usual choice of sampling of  $128 \ln m$  samples, and one marking  $256 \ln m$  samples as in the previous subsection.

Compare the effect of increased perturbation on the instance `email-univ` in Figure 8 with the effect on the instance `inf-USAir97` in Figure 9. For both instances, perturbation does not noticeably changes the asymptotic objective value of the averaging-and-scaling algorithm, causing the dashed lines to overlap. For `email-univ`, increasing perturbation consistently reduces the value of  $\beta_{\text{fcc}}$ , with  $\varepsilon = 1/32$  having a certificate quality almost 1% lower than  $\varepsilon = 10^{-4}$ . For `inf-USAir97`, when  $\varepsilon = 10^{-4}$ , the quality of the certificate improves up until around  $190 \ln m$  samples, and then no noticeable improvement is observed, even after  $1024 \ln m$  samples. This is not the case for the higher values of perturbation, where the certificate quality keeps increasing as we keep sampling. In particular,  $\varepsilon = 1/16$  obtains a result better than the averaging-and-scaling algorithm around  $256 \ln m$  samples, and  $\varepsilon = 3/64$  does so around  $500 \ln m$  samples.

Figures 10 and 11 show the same data but with the  $y$ -axis restricted to different ranges. Whereas Figure 10 shows both the horizontal line  $\alpha_{\text{GW}}$  and  $\beta_{\text{fcc}}^* \approx 0.9728$ , Figure 11 restricts the  $y$ -axis to emphasize the effect of perturbation on the behavior of our algorithm. For the specific instance `johnson16-2-4`, we know its maximum cut and fractional cut-covering values, and we can prove (via Corollary 6) that the lines in Figures 10 and 11 eventually achieve the value of  $\beta_{\text{fcc}}^*$  plotted in the figures. The instance `johnson16-2-4` is similar to `email-univ` in that perturbation has a negative impact on the quality of the certificates. Interestingly, perturbation improves the asymptotic objective value of the averaging-and-scaling algorithm. For the largest perturbation value  $\varepsilon = 1/16$ , after  $\lceil 2275 \ln m \rceil$  samples, the value of  $\beta_{\text{fcc}}$  is just below 0.955, whereas the asymptotic value of the averaging-and-scaling algorithm is 0.9573. For the other choices of perturbation,  $\lceil 2275 \ln m \rceil$  seem to be enough: we outperform the averaging-and-scaling algorithm after approximately  $500 \ln m$  samples for  $\varepsilon = 10^{-4}$ , and after approximately  $1200 \ln m$  samples for  $\varepsilon = 1/32$ . For  $\varepsilon = 3/64$ , around  $2300 \ln m$  samples seem to suffice to produce a certificate better than the averaging-and-scaling algorithm.

We conclude this section computing the values  $\text{mc}(G, w)$  and  $\text{fcc}(G, z)$  for `johnson16-2-4`. Let  $G = (V, E)$  be a graph, and let  $w, z \in \mathbb{R}_+^E$ . Let  $(\rho, \mu, R, B, x)$  be a tuple satisfying (20). Set

$$(51) \quad \sigma := 1 - \frac{w^\top z}{\rho\mu}, \beta_{\text{mc}}^* := \frac{\text{mc}(G, w)}{\rho}, \text{ and } \beta_{\text{fcc}}^* := \frac{(1 - \sigma)\mu}{\text{fcc}(G, z)}.$$

Since  $(1 - \sigma)\rho\mu = z^\top w \leq \text{fcc}(G, z) \text{mc}(G, z)$ , we have that  $\beta_{\text{mc}}^* \geq \beta_{\text{fcc}}^*$ . From (16) we have  $\beta_{\text{mc}}^* \geq \beta_{\text{mc}}$ , and  $\beta_{\text{fcc}}^* \geq \beta_{\text{fcc}}$ . Corollary 6 states that, as long as the matrix  $Y \in \mathbb{S}_+^V$  used for sampling is positive definite, our LP-based algorithm eventually computes an optimal fractional cut cover, in which case  $\beta_{\text{fcc}}^*$  is the limit, as the number of samples goes to infinity, of the certifiable quality  $\beta_{\text{fcc}}$  plotted in Figures 8 to 11.

Let  $G = (V, E)$  be an edge-transitive graph. Since  $G$  is edge-transitive, a simple symmetrization argument (see, e.g., [42, Lemma 2.1]) implies that

$$(52) \quad m = \text{fcc}(G) \text{mc}(G).$$

Set  $w := \mathbb{1}$ , set  $z := \gamma\mathbb{1}$  for some  $\gamma \in \mathbb{R}_+$ , and let  $(\rho, \mu, R, B, x)$  satisfy (20). We claim that

$$(53) \quad \beta_{\text{fcc}}^* = \beta_{\text{mc}}^*.$$

By definition of  $\sigma$ , we have that  $(1 - \sigma)\rho\mu = z^\top w = \gamma m$ . Then (52), the definition of  $z$ , and (53) imply (54):

$$\beta_{\text{fcc}}^* = \frac{(1 - \sigma)\mu}{\text{fcc}(G, z)} = \frac{(1 - \sigma)\mu}{\gamma \text{fcc}(G)} = \frac{(1 - \sigma)\mu \text{mc}(G)}{\gamma m} = \frac{(1 - \sigma)\mu\rho \text{mc}(G)}{\gamma m \rho} = \beta_{\text{mc}}^*.$$

Note that if  $z = \gamma\mathbb{1} = \frac{1}{4}\mathcal{L}_G^*(Y)$ , then

$$\frac{1}{4}\mathcal{L}_G^*((1 - \varepsilon)Y + \varepsilon I) = (1 - \varepsilon)z + \frac{\varepsilon}{2}\mathbb{1} = \left( (1 - \varepsilon)\gamma + \frac{\varepsilon}{2} \right) \mathbb{1}$$

for any  $\varepsilon \in [0, 1]$ . Therefore, for any perturbation  $\varepsilon$  of the nearly optimal solution, (54) still holds. From (54), we can conclude that for any perturbation  $\varepsilon$ , we have

$$\beta_{\text{fcc}}^* = \frac{\text{mc}(G)}{\rho} \text{ if } G \text{ edge-transitive and } z = \gamma\mathbb{1} \text{ for } \gamma \in \mathbb{R}_+,$$

as Corollary 6 implies that, with enough samples,  $\beta_{\text{fcc}} = \beta_{\text{fcc}}^*$ , and hence  $\beta_{\text{mc}}^* = \beta_{\text{mc}} = \beta_{\text{fcc}} = \beta_{\text{fcc}}^*$ .

We can now apply these results to johnson16-2-4. The instance johnson16-2-4 comes from the second DIMACS challenge [21], and represents a graph where each vertex is a subset of [16] with two elements, with two sets adjacent if their Hamming distance is at least 4. Note that, for this vertex set, we may equivalently say that two vertices are adjacent if the respective sets are disjoint. Hence johnson16-2-4 is the *Kneser graph*  $\text{Kn}(16, 2)$ . For every  $n \in \mathbb{N}$ , Poljak and Tuza [39] prove that a given set of shores of  $\text{Kn}(n, 2)$  contains a maximum cut. Enumerating this set of shores and picking the best one, we computed  $\text{mc}(\text{Kn}(16, 2)) = 3036$  (see Section D). When solving the SDP relaxation with  $\varepsilon = 10^{-4}$ , we obtained  $\rho = 3120.57$ . Thus

$$\beta_{\text{fcc}}^* = \frac{3036}{3120.57} \approx 0.9728.$$

In practice, we obtained  $z \in \mathbb{R}_+^E$  such that  $\|z - \gamma\mathbb{1}\|_\infty \leq 10^{-13}$  for a given  $\gamma \in \mathbb{R}_+$ . This implies that, with enough sampling, all the lines in Figures 10 and 11 will reach 0.9728. One can prove that  $\eta(\text{Kn}(16, 2)) = 3120$ , and indeed after sanitizing the output of DDS we were able to obtain  $\rho = 3120$ . This slightly improves the value of  $\beta_{\text{fcc}}^*$  to  $3036/3120 \approx 0.9730$ .




FIGURE 8. Value of  $\text{fcc}(\mathcal{F}_t, z)$  for **email-univ** across  $\lceil 1024 \ln m \rceil$  samples.




FIGURE 9. Value of  $\text{fcc}(\mathcal{F}_t, z)$  for **inf-USAir97** across  $\lceil 1024 \ln m \rceil$  samples.




FIGURE 10. Value of  $\text{fcc}(\mathcal{F}_t, z)$  for johnson16-2-4 across  $\lceil 2275 \ln m \rceil$  samples.




FIGURE 11. Value of  $\text{fcc}(\mathcal{F}_t, z)$  for johnson16-2-4 across  $\lceil 2275 \ln m \rceil$  samples.

## 8. CONCLUSIONS

Our experiments provide ample evidence of the ?? of the primal-dual framework described in Section 2. We have presented multiple algorithms which compute  $\beta$ -certificates given either a maximum cut or a fractional cut-covering instance as input. A theme throughout the paper is how, with the aid of perturbation, we obtain approximation ratios exceeding the theoretical guarantee of  $\alpha_{\text{GW}}$  with a modest  $\lceil 128 \ln m \rceil$  sample size. Computing a  $\beta$ -certificate from the nearly optimal solutions to either of the SDPs (2) or (12) is typically a cheap computational task relative to nearly solving these SDPs.

Our positive empirical results build on several aspects. Convex duality theory guides both theoretical and computational choices. Sanitization, which uses Algorithms 1 and 2 to produce the short certificate in Corollary 2, is crucial in making our results robust in the presence of floating-point error, which consequently allows distinct SDP solvers to be employed. The LP and SDP instances arising from our experiments can be reliably solved by current LP and SDP solvers. Perturbation is helpful not only in improving feasibility and properly addressing thin edges, but also in guaranteeing feasible solutions certifiable with Algorithm 2.

The experimental results provide validation for the theoretical duality framework proposed in [5]. Perhaps the most unusual property of the duality relationship presented in Section 2 is that to obtain the dual problem for a given input, one must solve an SDP to optimality. This is potentially challenging in the Turing machine model and, indeed, [5] assumes a real-number machine model [7] with access to oracles capable of computing exact square roots and sampling from a standard normal distribution. For our actual experiments, we have been working with approximate pairings, whose quality is captured by the parameter  $\sigma$  as defined in (18). Indeed, we have been pairing instances by mapping

$$\begin{aligned} z \in \mathbb{R}_+^E &\mapsto \{ w \in \mathbb{R}_+^E : w^\top z \geq (1 - \sigma)\eta(G, w)\eta^\circ(G, z) \}, \text{ and} \\ w \in \mathbb{R}_+^E &\mapsto \{ z \in \mathbb{R}_+^E : w^\top z \geq (1 - \sigma)\eta(G, w)\eta^\circ(G, z) \}. \end{aligned}$$

When  $\sigma = 0$ , we recover the pairing proposed in [5]. The final quality of the  $\beta$ -certificates computed are evidence that our framework is able to obtain good results by working with floating-point arithmetic. This robustness to approximate pairings is important even when one assumes a machine capable of doing exact real-number arithmetic, since [5] employs perturbation of the feasible region to address issues arising from thin edges.

Since we were able to reliably produce certificates with sample sizes much smaller than predicted by theory, some natural questions arise.

- (1) What results could one obtain for significantly larger graphs? The gap between  $C := 128$  and the upper bound in (30) leaves open the possibility that larger graphs might demand more samples to achieve good approximation ratios. Studying such larger instances requires solving significantly larger SDP and LP instances.
- (2) Conversely, can one prove improved bounds on sample sizes, perhaps on restricted families of graphs? Theory on the interplay between perturbation, LP and SDP solver behavior, and the randomized aspects of our algorithms could potentially guide such improvements.
- (3) The theoretical framework in [5] has been vastly extended into a conic setting in [6]. Does the ?? in computing certificates for the maximum cut problem and its dual generalize to computing certificates for the problem of maximizing convex quadratic forms on sign vectors [32] and its associated dual? What about other Boolean Constraint Satisfaction Problems on two variables?

## REFERENCES

- [1] M. F. Anjos and F. Liers. “Global approaches for facility layout and VLSI floorplanning”. In: *Handbook on semidefinite, conic and polynomial optimization*. Volume 166. Internat. Ser. Oper. Res. Management Sci. Springer, New York, 2012, pages 849–877.
- [2] A. Anshu, D. Gosset, and K. Morenz. “Beyond product state approximations for a quantum analogue of max cut”. In: *15th Conference on the Theory of Quantum Computation, Communication and Cryptography*. Volume 158. LIPIcs. Leibniz Int. Proc. Inform. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2020, Art. No. 7, 15 pages.
- [3] F. Barahona, M. Jünger, and G. Reinelt. “Experiments in quadratic 0-1 programming”. In: *Math. Programming* 44.2 (1989), pages 127–137.
- [4] F. Barahona. “On the computational complexity of Ising spin glass models”. In: *J. Phys. A* 15.10 (1982), pages 3241–3253.
- [5] N. Benedetto Proença, M. K. de Carli Silva, C. M. Sato, and L. Tunçel. “A Primal-Dual Extension of the Goemans–Williamson Algorithm for the Weighted Fractional Cut-Covering Problem”. In: *Math. Programming* 215 (2026), pages 1–56.
- [6] N. Benedetto Proença, M. K. de Carli Silva, C. M. Sato, and L. Tunçel. *Generalized Cuts and Grothendieck Covers: a Primal-Dual Approximation Framework Extending the Goemans–Williamson Algorithm*. June 26, 2024. arXiv: [2406.18670](https://arxiv.org/abs/2406.18670) [cd.DS].
- [7] L. Blum, F. Cucker, M. Shub, and S. Smale. *Complexity and Real Computation*. New York, NY: Springer, 1998. xvi+453 pages.
- [8] T. Bonato, M. Jünger, G. Reinelt, and G. Rinaldi. “Lifting and separation procedures for the cut polytope”. In: *Math. Programming* 146.1-2 (2014), pages 351–378.
- [9] G. Chen, G. Jing, and W. Zang. “Proof of the Goldberg-Seymour conjecture on edge-colorings of multigraphs”. In: *J. Comb. Optim.* 50.3 (2025), Paper No. 23, 91 pages.
- [10] J. Cheriyan, W. H. Cunningham, L. Tunçel, and Y. Wang. “A linear programming and rounding approach to max 2-sat”. In: *Cliques, coloring, and satisfiability (New Brunswick, NJ, 1993)*. Volume 26. DIMACS Ser. Discrete Math. Theoret. Comput. Sci. Amer. Math. Soc., Providence, RI, 1996, pages 395–414.
- [11] C. S. Edwards. “An improved lower bound for the number of edges in a largest bipartite subgraph”. In: *Recent advances in graph theory (Proc. Second Czechoslovak Sympos., Prague, 1974)*. Academia, Prague, 1975, pages 167–181.
- [12] P. Erdős. “On some extremal problems in graph theory”. In: *Israel J. Math.* 3 (1965), pages 113–116.
- [13] D. R. Fulkerson. “Anti-blocking polyhedra”. In: *J. Combinatorial Theory Ser. B* 12 (1972), pages 50–71.
- [14] D. R. Fulkerson. “Blocking and anti-blocking pairs of polyhedra”. In: *Math. Programming* 1 (1971), pages 168–194.
- [15] S. Gharibian and O. Parekh. “Almost optimal classical approximation algorithms for a quantum generalization of Max-Cut”. In: *Approximation, randomization, and combinatorial optimization. Algorithms and techniques*. Volume 145. LIPIcs. Leibniz Int. Proc. Inform. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019, Art. No. 31, 17 pages.
- [16] M. X. Goemans and D. P. Williamson. “Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming”. In: *J. ACM* 42.6 (1995), pages 1115–1145.
- [17] M. Grötschel, L. Lovász, and A. Schrijver. “The ellipsoid method and its consequences in combinatorial optimization”. In: *Combinatorica* 1.2 (1981), pages 169–197.
- [18] F. Harary, D. Hsu, and Z. Miller. “The biparticity of a graph”. In: *J. Graph Theory* 1.2 (1977), pages 131–133.

- [19] C. Helmberg and F. Rendl. “Solving quadratic  $(0, 1)$ -problems by semidefinite programs and cutting planes”. In: *Math. Programming* 82.3 (1998), pages 291–315.
- [20] T. Y. Ho and L. H. Hsu. “A note on the minimum cut cover of graphs”. In: *Oper. Res. Lett.* 15.4 (1994), pages 193–194.
- [21] D. S. Johnson and M. A. Trick, editors. *Cliques, coloring, and satisfiability*. Volume 26. DIMACS Series in Discrete Mathematics and Theoretical Computer Science. Papers from the workshop held as part of the 2nd DIMACS Implementation Challenge in New Brunswick, NJ, October 11–13, 1993. American Mathematical Society, Providence, RI, 1996, pages xii+657.
- [22] M. Karimi and L. Tunçel. “Efficient implementation of interior-point methods for quantum relative entropy”. In: *INFORMS J. Comput.* 37.1 (2025), pages 3–21.
- [23] R. M. Karp. “Reducibility among combinatorial problems”. In: *Complexity of computer computations (Proc. Sympos., IBM Thomas J. Watson Res. Center, Yorktown Heights, N.Y., 1972)*. Plenum, New York, 1972, pages 85–103.
- [24] T. Kelly and L. Postle. “Fractional coloring with local demands and applications to degree-sequence bounds on the independence number”. In: *Journal of Combinatorial Theory, Series B* 169 (2024), pages 298–337.
- [25] S. Khot, G. Kindler, E. Mossel, and R. O’Donnell. “Optimal inapproximability results for MAX-CUT and other 2-variable CSPs?”. In: *SIAM J. Comput.* 37.1 (2007), pages 319–357.
- [26] S. A. Khot and N. K. Vishnoi. “The unique games conjecture, integrability gap for cut problems and embeddability of negative-type metrics into  $\ell_1$ ”. In: *J. ACM* 62.1 (2015), Art. 8, 39 pages.
- [27] N. Krislock, J. Malick, and F. Roupin. “BiqCrunch: a semidefinite branch-and-bound method for solving binary quadratic problems”. In: *ACM Trans. Math. Software* 43.4 (2017), Art. 32, 23 pages.
- [28] R. Loulou. “Minimal cut cover of a graph with an application to the testing of electronic boards”. In: *Oper. Res. Lett.* 12.5 (1992), pages 301–305.
- [29] *Matrix Market*. National Institute of Standards and Technology. URL: <https://math.nist.gov/MatrixMarket/> (visited on 04/19/2026).
- [30] D. W. Matula. “ $k$ -components, clusters and slicings in graphs”. In: *SIAM J. Appl. Math.* 22 (1972), pages 459–480.
- [31] R. Mirka and D. P. Williamson. “An experimental evaluation of semidefinite programming and spectral algorithms for max cut”. In: *ACM J. Exp. Algorithmics* 28 (2023), Art. 2.1, 18 pages.
- [32] Y. Nesterov. “Semidefinite relaxation and nonconvex quadratic optimization”. In: *Optim. Methods Softw.* 9.1-3 (1998), pages 141–160.
- [33] J. Neto and W. Ben-Ameur. “On fractional cut covers”. In: *Discrete Appl. Math.* 265 (2019), pages 168–181.
- [34] B. O’Donoghue, E. Chu, N. Parikh, and S. Boyd. “Conic Optimization via Operator Splitting and Homogeneous Self-Dual Embedding”. In: *Journal of Optimization Theory and Applications* 169.3 (June 2016), pages 1042–1068.
- [35] B. O’Donoghue, E. Chu, N. Parikh, and S. Boyd. *SCS: Splitting Conic Solver*. Version 3.2.7. August 2024.
- [36] P. M. Pardalos and G. P. Rodgers. “Parallel branch and bound algorithms for quadratic zero-one programs on the hypercube architecture”. In: volume 22. 1-4. Supercomputers and large-scale optimization: algorithms, software, applications (Minneapolis, MN, 1988). 1990, pages 271–292.
- [37] S. Poljak and F. Rendl. “Solving the max-cut problem using eigenvalues”. In: volume 62. 1-3. Partitioning and decomposition in combinatorial optimization. 1995, pages 249–278.
- [38] S. Poljak and D. Turzík. “A polynomial time heuristic for certain subgraph optimization problems with guaranteed worst case bound”. In: *Discrete Math.* 58.1 (1986), pages 99–104.

- [39] S. Poljak and Z. Tuza. “Maximum bipartite subgraphs of Kneser graphs”. In: *Graphs Combin.* 3.2 (1987), pages 191–199.
- [40] G. Reinelt. *TSPLIB95*. URL: <http://comopt.ifi.uni-heidelberg.de/software/TSPLIB95/> (visited on 03/09/2025).
- [41] F. Rendl, G. Rinaldi, and A. Wiegele. “Solving max-cut to optimality by intersecting semidefinite and polyhedral relaxations”. In: *Math. Programming* 121.2 (2010), pages 307–335.
- [42] R. Šámal. “Cubical coloring—fractional covering by cuts and semidefinite programming”. In: *Discrete Math. Theor. Comput. Sci.* 17.2 (2015), pages 251–266.
- [43] C. Sanderson and R. Curtin. “Armadillo: an Efficient Framework for Numerical Linear Algebra”. In: *2025 17th International Conference on Computer and Automation Engineering (ICCAE)*. 2025, pages 303–307.
- [44] C. Sanderson and R. Curtin. *Armadillo: C++ Library for Linear Algebra and Scientific Computing*. Version 12.8.x. Git commit [43c41ab7](https://github.com/ArmadilloLibrary/Armadillo/commit/43c41ab7).
- [45] E. R. Scheinerman and D. H. Ullman. *Fractional graph theory*. Dover Publications, Inc., Mineola, NY, 2011. xviii+211 pages.
- [46] A. Schrijver. *Theory of linear and integer programming*. Wiley-Interscience Series in Discrete Mathematics. A Wiley-Interscience Publication. John Wiley & Sons, Ltd., Chichester, 1986, pages xii+471.
- [47] O. Tange. *GNU Parallel*. Version 20221122. November 2022.
- [48] Y. Ye, M. J. Todd, and S. Mizuno. “An  $O(\sqrt{n}L)$ -iteration homogeneous and self-dual linear programming algorithm”. In: *Math. Oper. Res.* 19.1 (1994), pages 53–67.

## APPENDIX A. USING THE SOFTWARE

**A.1. Basic Dependencies.** Before compiling our code, you need to install its dependencies. Our core dependencies are: git, pkg-config, Boost, CMake, BLAS and LAPACK. On Ubuntu 24.04.2 LTS, these can be installed with

```
sudo apt-get install git pkg-config libboost-all-dev \
  cmake libblas-dev liblapack-dev
```

The files in the folders `linux` and `macos` allow you to describe where dependencies are found. We continue this description assuming you are using `linux`. The file `linux/boost.pc`, as distributed, should be able to detect where one can find the headers and shared libraries to run Boost. The file `linux/gurobi.pc` requires modifications to indicate where Gurobi is installed on your system. It should be enough to simply change the line defining `prefix` to point to where Gurobi is installed in your system, and the last line `Libs` to load the correct libraries for compilation. The libraries need for compilation change according to your Gurobi version. For example, for Gurobi 10 on ?? hardware, the file is

```
prefix=/opt/uw/gurobi/10.0.0/linux64
includedir=${prefix}/include
libdir=${prefix}/lib
```

Name: Gurobi

Description: The Gurobi optimizer

Version: 10.0.0

Cflags: -I\${includedir}

Libs: -Wl,-rpath=\${libdir} -L\${libdir} -lgurobi\_g++5.2 -lgurobi100

and for Gurobi 12 on ?? hardware it is

```
prefix=/home/nathan/Documents/gurobi1202/linux64
includedir=${prefix}/include
libdir=${prefix}/lib
```

Name: Gurobi

Description: The Gurobi optimizer

Version: 12.0.2

Cflags: -I\${includedir}

Libs: -Wl,-rpath=\${libdir} -L\${libdir} -lgurobi\_g++8.5 -lgurobi120

Gurobi requires a license; read more about it in its [getting started page](#). MATLAB is an optional dependency (see Section [A.3](#)); if you do not wish to use it, simply remove the file `linux/matlab.pc`.

Makefile handles the remaining dependencies. You can just go into the directory and run

```
make
```

If no issues arise, congratulations, you have a minimal working installation!

**A.2. Running Experiments.** To run an experiment, you need to create a file describing which components of the pipeline you want to use, and select their configuration parameters.

For convenience, the input instances used in this manuscript are available with the code. You do need to convert the files into the format used by our pipeline. You can do so with the following commands:

```
./mmconverter instances/*.mtx
./tspconverter instances/*.tsp
```

A file called `instances/ENZYMES8.mtx.in` (among other files) should now exist.

With the instance files available, you can create a file called `experiment.cfg` with the following contents:

```
instance = instances/ENZYMES8.mtx.in
scaler = w_1_norm
solver = nu_scs
coverer = gurobi
sampler = hyperplane
presampler = null
scs_eps_rel = 1e-4
scs_eps_abs = 1e-4
scs_pert_eps = 0.0
round_eps = 0.0
rank_cutoff = 1e-8
sampler_C = 128
sampler_seed = 1337
```

You can then run this experiment with `./main experiments.cfg`. This file describes a run for the instance ENZYMES8 in Table [0.A](#). The options `scaler`, `solver`, `coverer`, `sampler`, and `presampler` correspond to the options in our pipeline Figure [1](#). Each options further requires its own parameters to be set. For example, the `nu_scs` solver formulates [\(2\)](#) as in [\(41\)](#). You can specify SCS termination parameters with `scs_eps_rel` and `scs_abs_rel`, and the perturbation parameter  $\varepsilon$  used as in [\(28\)](#) with `round_eps`.

The experiments we have used to generate the tables in this paper are available in the Makefile. For example, you can run

```
make first-table
```

to create a folder `first-table` with several configurations of our pipeline, corresponding to Table [0.A](#). You can check the options with, for example,

```
cat first-table/brg180-nu-null.scs-0000-0000-0000-0000
```

And then run

```
./main first-table/brg180-nu-null.scs-0000-0000-0000-0000
```

You can then look at the file again and see the results:

```
cat first-table/brg180-nu-null.scs-0000-0000-0000-0000
```

**A.3. Optional dependencies: Matlab.** We offer the option to integrate our pipeline with Matlab to use DDS and SeDuMi to solve the SDP relaxations (2) and (12). Since we interact with matlab from C++, our code needs to load the MATLAB Data API for C++, and for this reason you should edit the file `linux/matlab.pc` to indicate where we can find these files. Changing the line `prefix` should suffice.

To install the specific versions of the solvers we work, you can run

```
make DDS-2.2 sedumi-1.3.7
```

You can then run either

```
make dds-mm
```

or

```
make sedumi-mm
```

to create the files describing maximum cut experiments with `??` instances being solved by DDS or SeDuMi.

**A.4. Adding Input Instances.** To add an input instance, create a text file where

- (i) the first line has  $n$  and  $m$ , denoting the number of vertices and edges in the graph, respectively;
- (ii) the following  $m$  lines have  $u, v$ , and  $w$ , where  $u$  and  $v$  are numbers in  $\{1, \dots, n\}$  representing a vertex and  $w \in \mathbb{R}$  is a floating point number representing the weight of the edge.

For example, we can represent a path on three vertices with edge weights  $3/2$  and  $5/2$  with

```
3 2
1 2 1.5
2 3 2.5
```

If we save this file as `path.in`, one can run our pipeline on this input instance by changing the instance line to `instance = path.in`.

We provide two helpers to generate these standardized files from other formats, both of which are compiled when running `make`. The executable `mmconverter` accepts sparse matrices in some of the formats described by Matrix Market. In particular, it accepts `coordinate pattern symmetric` and `coordinate real symmetric`. Running `./mmconverter instances/ENZYMES8.mtx`, one creates a file `instances/ENZYMES8.mtx.in` that is accepted by our pipeline. Similarly, the executable `tspconverter` accepts some of the formats described by TSPLib. In particular, it accepts those with `EDGE_WEIGHT_TYPE` given by `EXPLICIT`, `GEO`, `EUC_2D`. Running `./tspconverter instances/brg180.tsp` creates a file `instances/brg180.tsp.in` which is accepted by our pipeline. Our `GEO` instances may not match those used by the TSP literature, see Section E.

## APPENDIX B. FORMULATIONS

Throughout this section,  $\mathbb{E}$  and  $\mathbb{Y}$  are Euclidean spaces,  $\mathcal{A}: \mathbb{E} \rightarrow \mathbb{Y}$  is a linear transformation, and  $b \in \mathbb{Y}$  and  $c \in \mathbb{E}$  are vectors.

**B.1. SCS.** Let  $K \subseteq \mathbb{Y}$  be a closed convex cone. In the linear conic case considered here, the standard form of an optimization problem solved by SCS is

$$(40) \quad \min\{\langle c, x \rangle : x \in \mathbb{E}, s \in K, \mathcal{A}(x) + s = b\} \geq \max\{\langle -b, y \rangle : y \in K^*, -\mathcal{A}^*(y) = c\}.$$

SCS has a stopping criterion based on two parameters:  $\varepsilon_{\text{abs}}$  and  $\varepsilon_{\text{rel}}$ . Upon successful termination, SCS returns  $(x, s, y)$  such that  $x \in K$ ,  $s \in K^*$ , and

$$(54a) \quad \|\mathcal{A}(x) + s - b\|_\infty \leq \varepsilon_{\text{abs}} + \varepsilon_{\text{rel}} \max\{\|\mathcal{A}(x)\|_\infty, \|s\|_\infty, \|b\|_\infty\}$$

$$(54b) \quad \|\mathcal{A}^*(y) + c\|_\infty \leq \varepsilon_{\text{abs}} + \varepsilon_{\text{rel}} \max\{\|\mathcal{A}^*(y)\|_\infty, \|c\|_\infty\}$$

$$(54c) \quad |c^\top x + b^\top y| \leq \varepsilon_{\text{abs}} + \varepsilon_{\text{rel}} \max\{|c^\top x|, |b^\top y|\}.$$

Note that SCS represents all of its cones as embedded in  $\mathbb{R}^d$ . In particular,

$$(55) \quad \|Y\|_{\text{SCS}} = \max\left\{\max_{i \in [n]} |Y_{ii}|, \max_{i \neq j \in [n]} \sqrt{2}|Y_{ij}|\right\} \text{ for every } Y \in \mathbb{S}^V.$$

For the formulation (41), on successful termination SCS outputs solutions  $x \in \mathbb{R}^V$ ,  $Y \in \mathbb{S}_+^V$ , and  $S \in \mathbb{S}_+^V$  satisfying

$$(56) \quad \begin{aligned} \|S - (\text{Diag}(x) - \tfrac{1}{4}\mathcal{L}_G(w))\|_{\text{SCS}} &\leq \varepsilon_{\text{abs}} + \varepsilon_{\text{rel}} \max\{\|x\|_\infty, \|S\|_{\text{SCS}}, \|\tfrac{1}{4}\mathcal{L}_G(w)\|_{\text{SCS}}\} \\ \|\text{diag}(Y) - \mathbb{1}\|_\infty &\leq \varepsilon_{\text{abs}} + \varepsilon_{\text{rel}} \cdot \max\{\|\text{diag}(Y)\|_\infty, 1\} \\ |\mathbb{1}^\top x - \langle \tfrac{1}{4}\mathcal{L}_G(w), Y \rangle| &\leq \varepsilon_{\text{abs}} + \varepsilon_{\text{rel}} \max\{|\mathbb{1}^\top x|, |\langle \tfrac{1}{4}\mathcal{L}_G(w), Y \rangle|\} \\ \langle Y, S \rangle &= 0, \text{ with } Y \succeq 0, S \succeq 0 \end{aligned}$$

**B.2. SeDuMi.** Let  $K \subseteq \mathbb{E}$  be a closed convex cone. The standard form of an optimization problem solved by SeDuMi is

$$(57) \quad \min\{\langle c, x \rangle : x \in K, \mathcal{A}(x) = b\} \geq \max\{\langle b, y \rangle : y \in K^*, \mathcal{A}^*(y) + s = c\}.$$

Note that, in the previous formulation with SCS and DDS, we had  $K \subseteq \mathbb{Y}$ , whereas now we assume  $K \subseteq \mathbb{E}$ . Let  $G = (V, E)$  be a graph, and let  $w \in \mathbb{R}_+^E$  be so  $(G, w)$  is a maximum cut instance. The SeDuMi formulation of (2b) is given by (58) with  $K := \mathbb{S}_+^V$  and  $\mathcal{A}: \mathbb{S}^V \rightarrow \mathbb{R}^V$  given by

$$(58) \quad \mathcal{A}(Y) := \text{diag}(Y), b := \mathbb{1}, c := -\tfrac{1}{4}\mathcal{L}_G(w).$$

**B.3. Formulations of polar SDP.** Let  $G = (V, E)$  be a graph. Set  $\gamma_G := 1 + n + m$  which we use as a uniform scaling factor. Let  $\varepsilon, \varphi, \xi \in (0, 1)$ . The most general formulation for (12) we studied was the following:

$$(59) \quad \begin{array}{ll} \text{Minimize } \gamma_G \cdot (\mu - \varphi \text{Tr}(Y)), & \text{Maximize } \hat{z}^\top w, \\ \text{subject to } \tfrac{1}{4}\mathcal{L}_G^*(Y + \mu\varepsilon I) \geq z, & \text{subject to } \tfrac{\varepsilon}{2}\mathbb{1}^\top w + (1 - \varepsilon)\mathbb{1}^\top x \leq \gamma_G, \\ \text{diag}(Y + \mu\varepsilon I) = \mu\mathbb{1}, & -\tfrac{1}{4}\mathcal{L}_G(w) + \text{Diag}(x) \succeq \gamma\varphi I, \\ \mu \in \mathbb{R}_+, & w \in \mathbb{R}_+^E, \\ Y \in \mathbb{S}_+^V, & x \in \mathbb{R}^V. \end{array}$$

where  $\hat{z}_e := \max\{z_e, \frac{1}{2}\xi\|z\|_\infty\}$  for every  $e \in E$ . The parameter  $\varphi$  strengthens the numerical robustness of the ‘ $\text{Diag}(x) \succeq \frac{1}{4}\mathcal{L}_G(w)$ ’ constraint. The parameters  $\xi$  and  $\varepsilon$  are related to thin-edges: The parameter  $\xi$  ensures  $\hat{z}_{ij} \geq \frac{1}{2}\xi\|\hat{z}\|_\infty$ , and  $\varepsilon$  ensures the eigenvalues of the Gram matrix used for

sampling are at least  $\mu\varepsilon$ . Preliminary experiments suggested it was reasonable to keep both  $\xi = 0$  and  $\varphi = 0$ , which explains the simplified formulation in (33).

The SCS formulation of (60) is given by (40) with

$$(60) \quad \begin{aligned} \mathbb{E} &:= \mathbb{R}^E \oplus \mathbb{R}^V, \mathbb{Y} := \mathbb{R} \oplus \mathbb{R}^E \oplus \mathbb{S}^V, K := \mathbb{R}_+ \oplus \mathbb{R}_+^E \oplus \mathbb{S}_+^V \subseteq \mathbb{Y}, \\ \mathcal{A}\left(\begin{bmatrix} w \\ x \end{bmatrix}\right) &:= \begin{bmatrix} \frac{\varepsilon}{2}\mathbb{1}^\top w + (1-\varepsilon)\mathbb{1}^\top x \\ -w \\ \frac{1}{4}\mathcal{L}_G(w) - \text{Diag}(x) \end{bmatrix}, b := \gamma_G \begin{bmatrix} 1 \\ 0 \\ -\varphi I \end{bmatrix}, c := \begin{bmatrix} -\hat{z} \\ 0 \end{bmatrix}. \end{aligned}$$

For both DDS and SeDuMi, we assume  $\varphi = \xi = 0$ , and thus formulate (33). The DDS formulation of (33) is given by (43) with

$$(61) \quad \begin{aligned} \mathbb{E} &:= \mathbb{R}^E \oplus \mathbb{R}^V, \mathbb{Y} := \mathbb{R} \oplus \mathbb{R}^E \oplus \mathbb{S}^V, K := \mathbb{R}_+ \oplus \mathbb{R}_+^E \oplus \mathbb{S}_+^V \subseteq \mathbb{Y}, \\ \mathcal{A}\left(\begin{bmatrix} w \\ x \end{bmatrix}\right) &:= \begin{bmatrix} -\frac{\varepsilon}{2}\mathbb{1}^\top w - (1-\varepsilon)\mathbb{1}^\top x \\ w \\ -\frac{1}{4}\mathcal{L}_G(w) + \text{Diag}(x) \end{bmatrix}, b := \gamma_G \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \text{ and } c := \begin{bmatrix} -z \\ 0 \end{bmatrix}. \end{aligned}$$

The SeDuMi formulation of (33) is given by (58) with

$$(62) \quad \begin{aligned} \mathbb{E} &:= \mathbb{R} \oplus \mathbb{R}^E \oplus \mathbb{S}^V, \mathbb{Y} := \mathbb{R}^E \oplus \mathbb{R}^V, K := \mathbb{R}_+ \oplus \mathbb{R}_+^E \oplus \mathbb{S}_+^V \subseteq \mathbb{E}, \\ \mathcal{A}\left(\begin{bmatrix} \mu \\ u \\ Y \end{bmatrix}\right) &:= \begin{bmatrix} \frac{1}{4}\mathcal{L}_G^*(Y + \mu\varepsilon I) - u \\ (1-\varepsilon)\mu\mathbb{1} - \text{diag}(Y) \end{bmatrix}, b := \begin{bmatrix} z \\ 0 \end{bmatrix}, c := \gamma_G \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}. \end{aligned}$$

Formulations (61), (62), and (63) make concrete our claim that we selected the most straightforward implementations of (33). Furthermore, we always formulated our problem over the same cone  $K = \mathbb{R}_+ \oplus \mathbb{R}_+^E \oplus \mathbb{S}_+^V$ .

We compare the stopping criteria for SCS and DDS. Note that (44a), by actually asserting feasibility, is stronger than (55a), which only states that the primal solution is near a feasible solution. The inequality  $\text{Diag}(x) \succeq \frac{1}{4}\mathcal{L}_G(w)$  is part of primal feasibility in our formulations for both (2) and (33). The other guarantees are ensured in a similar way, be it the linear constraint on the dual variables or the duality gap bound. Note that the 1 in the RHS of the bounds (44b) and (44c) imply that the single parameter  $\varepsilon_{\text{DDS}}$  is both an absolute-error and a relative-error bound. SeDuMi does not report its stopping criteria in its manual. However, enough is known of its implementation that we can comment on it. It is an Interior-Point Method implementation using Self-Dual Embedding Technique of [48], which does not accept starting primal or dual solution. For example, for (63), on a successful termination, SeDuMi computes  $w \in \mathbb{R}^E$ ,  $x \in \mathbb{R}^V$ , and  $S \in \mathbb{R}_+ \oplus \mathbb{R}_+^E \oplus \mathbb{S}_+^V$  such that

$$\mathcal{A}^*\left(\begin{bmatrix} w \\ x \end{bmatrix}\right) + S - \begin{bmatrix} \gamma_G \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} \frac{\varepsilon}{2}\mathbb{1}^\top w + (1-\varepsilon)\mathbb{1}^\top x \\ -w \\ \frac{1}{4}\mathcal{L}_G(w) - \text{Diag}(x) \end{bmatrix} + S - \begin{bmatrix} \gamma_G \\ 0 \\ 0 \end{bmatrix} \approx 0.$$

In other words, similar to SCS, SeDuMi only guarantees that  $\text{Diag}(x) - \frac{1}{4}\mathcal{L}_G(w)$  is near a positive semidefinite matrix. In practice, due to the tighter constraint of  $10^{-8}$ , the output of SeDuMi behaved similarly to the output of DDS.

## APPENDIX C. COMPARISON WITH SEDUMI

Experiments in this appendix were run on ??.

**C.1. Maximum Cut Instances.** As input, we are given maximum cut instances  $(G, w)$ . We solve (2) with SeDuMi using formulation (59). We record the SDP solving time, and compare it to the SDP solving time needed by DDS when collecting the data in Tables 12.A and 12.B. As perturbation does not change the SDP being solved, we show timing information for  $\varepsilon = 0$ .

instance	DDS	SeDuMi	instance	DDS	SeDuMi
bayg29	1.32	1.14	ENZYMES8	1.83	1.22
bays29	1.37	1.16	soc-dolphins	1.60	1.17
berlin52	1.42	1.12	road-chesapeake	1.45	1.15
brazil58	1.55	1.18	email-enron-only	2.24	1.41
gr96	1.87	1.27	dwt_209	3.53	1.76
kroA100	2.01	1.27	ca-netscience	12.40	4.39
eil101	1.87	1.29	Erdos991	13.58	7.63
gr120	2.21	1.45	hamming6-2	1.42	1.13
bier127	2.10	1.36	inf-USAir97	10.00	3.45
ch130	2.17	1.41	ia-infect-hyper	2.11	1.30
gr137	2.29	1.56	DD687	38.50	14.87
ch150	2.35	1.54	dwt_503	16.69	8.57
brg180	2.82	1.62	ia-infect-dublin	10.23	4.57
d198	3.74	1.98	email-univ	149.19	53.30
gr202	3.35	1.82	johnson16-2-4	1.67	1.19
a280	4.90	2.70	p-hat700-1	30.02	15.88

TABLE 15. Comparing SDP (2) solving times (in seconds) between DDS and SeDuMi. For DDS, data corresponds to Table 12.A with  $\varepsilon = 0$ .

**C.2. Fractional Cut-Covering Instances.** As input, we are given paired fractional cut-covering instances  $(G, z)$ . In Table 16.A, they are ??, and in Table 16.B they are ??. We solve (33) with SeDuMi (using formulation (63)) for each of  $\varepsilon \in \{0, 10^{-8}, 1/64\}$ . We record the SDP solving time, and compare it to the SDP solving time needed by DDS when collecting the data in Tables 13.A and 13.B. SeDuMi reported it ran “into numerical problems” for the instances `ch130`, `ch150`, `gr120`, `gr137`, and `gr96`. The output was, however, accepted by our sanitization procedure. For the instances `brg180`, `d198`, and `gr202`, SeDuMi was terminated after 8 hours (28 800 seconds).

As shown in Table 15, SeDuMi outperforms DDS by a factor in the range  $[1.15, 3]$  on our maxcut instances. This finding is consistent with other experiments reported in the literature comparing DDS and SeDuMi. One contributing reason for DDS’s slower performance is its use of a more conservative and sophisticated step-size strategy relative to other commonly used interior-point based software.

The picture changes considerably when we turn to the fractional cut-covering instances. For the small and easy SDPs reported in Tables 16.A and 16.B, SeDuMi remains faster than DDS by a factor of up to 3. However, on many of the larger and more challenging instances, SeDuMi either takes substantially longer than DDS to solve the instances or encounters numerical difficulties. Taken together, these experiments suggest that on the instances used in our experiments, DDS performs more robustly than SeDuMi on harder instances, both in terms of computation time and numerical reliability.

instance	DDS	SeDuMi	DDS	SeDuMi	DDS	SeDuMi
	$\varepsilon = 0$	$\varepsilon = 0$	$\varepsilon = 10^{-8}$	$\varepsilon = 10^{-8}$	$\varepsilon = 1/64$	$\varepsilon = 1/64$
bayg29	1.26	3.09	1.19	2.38	1.66	2.22
bays29	1.19	1.39	1.20	0.90	2.14	0.87
berlin52	3.62	4.97	4.05	3.32	8.18	4.16
brazil58	4.96	5.56	5.87	5.64	8.32	7.03
kroA100	5.03	17.57	5.59	17.86	7.01	16.39
gr96	30.68	702.20	39.45	725.17	65.54	357.24
eil101	33.73	762.26	45.04	757.99	90.42	508.79
gr120	106.50	872.12	131.97	878.18	312.76	1 208.93
bier127	248.74	2 355.25	314.32	2 344.48	420.88	2 613.99
ch130	121.89	2 834.15	152.12	2 839.44	280.58	1 924.76
gr137	239.31	2 050.61	274.13	2 049.40	420.01	2 543.89
ch150	239.72	7 509.22	287.64	7 509.28	662.26	4 884.11
brg180	563.15	–	658.13	–	1 565.76	–
d198	4 135.89	–	4 880.60	–	5 992.67	–
gr202	4 039.98	–	4 819.71	–	5 088.82	–

TABLE 16.A. Comparing SDP (33) solve time (in seconds) between DDS and SeDuMi.

instance	DDS	SeDuMi	DDS	SeDuMi	DDS	SeDuMi
	$\varepsilon = 0$	$\varepsilon = 0$	$\varepsilon = 10^{-8}$	$\varepsilon = 10^{-8}$	$\varepsilon = 1/64$	$\varepsilon = 1/64$
ENZYMES8	1.51	0.99	1.41	0.95	1.89	1.02
soc-dolphins	1.41	1.04	1.34	1.06	1.56	0.88
road-chesapeake	1.15	0.91	1.12	0.96	1.16	0.81
email-enron-only	2.52	2.38	3.03	2.09	3.47	2.07
dwt_209	5.70	3.88	4.77	3.93	5.93	3.46
ca-netscience	16.59	8.77	14.03	9.07	12.98	13.59
Erdos991	18.35	22.27	21.66	18.86	24.44	22.52
hamming6-2	3.29	6.50	3.78	6.50	6.93	6.58
inf-USAir97	23.00	46.14	27.33	44.24	34.38	39.34
ia-infect-hyper	6.34	25.29	7.81	32.11	10.76	21.20
DD687	74.06	112.32	69.24	119.05	90.38	117.06
dwt_503	38.28	109.08	46.53	121.54	74.42	119.09
ia-infect-dublin	31.35	89.37	30.18	86.36	37.34	86.94
email-univ	281.62	1 146.86	258.74	1 250.39	364.52	1 184.11
johnson16-2-4	24.45	424.85	33.77	495.51	34.81	397.62

TABLE 16.B. Comparing SDP (33) solve time (in seconds) between DDS and SeDuMi.

## APPENDIX D. MAXIMUM CUT FOR KNESER GRAPH

Let  $n \in \mathbb{N}$ . [39, Theorem 5] prove that, for some  $p \in \{1, \dots, n\}$ , the shore

$$\left\{ S \in \binom{[n]}{2} : S \cap \{1, \dots, p\} \neq \emptyset \right\}$$

defines a maximum cut of  $\text{Kn}(n, 2)$ . By enumerating such shores for  $\text{Kn}(16, 2)$ , we computed its maximum cut value. This subsection describes the code used to do so.

We defined a structure to represent vertices of the Kneser graph and quickly test if they belong to the shore we are testing.

```

1 size_t p = 0;
2
3 struct vertex_t {
4     size_t a, b;
5 };
6
7 struct vertex_t vertex(size_t i, size_t j)
8 {
9     assert(i < j);
10    return (struct vertex_t) {
11        i, j
12    };
13 }
14
15 int in_cut(struct vertex_t S)
16 {
17     return (S.a <= p);
18 }

```

We then simply try all choices of  $p$ , counting how many edges are cut for each shore. To enumerate edges, we enumerate all subsets of  $[n]$  with 4 elements  $\{i, j, k, \ell\}$ , and then test the three different edges involving these points:

$$\{\{i, j\}, \{k, \ell\}\}, \{\{i, k\}, \{j, \ell\}\}, \{\{i, \ell\}, \{j, k\}\}.$$

```

1 void process_edges(size_t i, size_t j, size_t k, size_t l)
2 {
3     cut_size += in_cut(vertex(i, j)) != in_cut(vertex(k, l));
4     cut_size += in_cut(vertex(i, k)) != in_cut(vertex(j, l));
5     cut_size += in_cut(vertex(i, l)) != in_cut(vertex(j, k));
6 }
7
8 int main()
9 {
10    size_t n = 16;
11    for (p = 1; p <= n; p++) {
12        cut_size = 0;
13        for (size_t i = 1; i <= n; i++)
14            for (size_t j = i + 1; j <= n; j++)
15                for (size_t k = j + 1; k <= n; k++)
16                    for (size_t l = k + 1; l <= n; l++)
17                        process_edges(i, j, k, l);
18        printf("p = %zu:\t%zu\n", p, cut_size);
19    }
20 }

```

## APPENDIX E. TSPLIB INSTANCES

In this short section, we discuss how the ?? maximum cut instances we have used may not match the instances used in the broader TSP literature, where they originate from.

TSPLib [40] is a collection of weighted graphs encoded in a dozen different ways, among them the GEO format. For this format, each vertex is given latitude and longitude coordinates, and the distance between vertices is the distance on the three-dimensional sphere. The TSPLib documentation provides a “a (simplified) C-implementation for computing the distances from the input coordinates”. The code for computing latitude and longitude for every vertex  $i$  is

```
PI = 3.141592;
deg = nint( x[i] );
min = x[i] - deg;
latitude[i] = PI * (deg + 5.0 * min / 3.0 ) / 180.0;
deg = nint( y[i] );
min = y[i] - deg;
longitude[i] = PI * (deg + 5.0 * min / 3.0 ) / 180.0
```

and the code for computing distances between vertices  $i$  and  $j$  is

```
RRR = 6378.388;
q1 = cos( longitude[i] - longitude[j] );
q2 = cos( latitude[i] - latitude[j] );
q3 = cos( latitude[i] + latitude[j] );
dij = (int) ( RRR * acos( 0.5*((1.0+q1)*q2 - (1.0-q1)*q3) ) + 1.0);
```

The documentation states “`nint(x)` can be replaced by `(int) (x+0.5)`”, which computes the nearest integer for positive floating point numbers.

Whereas in the first part, `deg` is computed with `nint`, in the second part, `dij` is computed directly from an integer conversion. [31] access TSPLib instances using a Julia library. The library uses the Julia function `trunc`, which rounds the floating point towards zero, and thus is equivalent to computing the floor for non-negative floating point numbers. Importantly, the library uses `trunc` both when computing `deg` and when computing `dij`, and thus amounts to changing the lines in the first block of code above into `deg = (int) (x[i])` and `deg = (int) (y[i])`. The library maintainer has some discussion on the impact of these changes on the actual graphs produced. We adopt the same convention as the Julia library, so our graphs match the ones used in [31].