

Scalable Semidefinite Programming*

Alp Yurtsever[†], Joel A. Tropp[‡], Olivier Fercoq[§], Madeleine Udell[¶], and Volkan Cevher[†]

Abstract. Semidefinite programming (SDP) is a powerful framework from convex optimization that has striking potential for data science applications. This paper develops a provably correct randomized algorithm for solving large, weakly constrained SDP problems by economizing on the storage and arithmetic costs. Numerical evidence shows that the method is effective for a range of applications, including relaxations of MaxCut, abstract phase retrieval, and quadratic assignment. Running on a laptop equivalent, the algorithm can handle SDP instances where the matrix variable has over 10^{14} entries.

Key words. Augmented Lagrangian, conditional gradient method, convex optimization, dimension reduction, first-order method, randomized linear algebra, semidefinite programming, sketching.

AMS subject classifications. Primary: 90C22, 65K05. Secondary: 65F99.

1. Motivation. For a spectrum of challenges in data science, methodologies based on semidefinite programming offer remarkable performance both in theory and for small problem instances. Even so, practitioners often critique this approach by asserting that it is impossible to solve semidefinite programs (SDPs) at the scale demanded by real-world applications. We would like to argue against this article of conventional wisdom.

This paper proposes a new algorithm, called SketchyCGAL, that can solve very large SDPs to moderate accuracy. The algorithm marries a primal–dual optimization technique [102] to a randomized sketch for low-rank matrix approximation [93]. In each iteration, the primary expense is one low-precision randomized eigenvector calculation [57].

For every standard-form SDP that satisfies strong duality, SketchyCGAL provably converges to a near-optimal low-rank approximation of a solution. The algorithm uses limited arithmetic and minimal storage. It is most effective for weakly constrained problems whose solutions are nearly low-rank. In contrast, given the same computational resources, other methods for this class of problems may fail. In particular, SketchyCGAL needs far less storage

*Submitted to the editors 6 December 2019. Revised on 18 January 2020, 24 July 2020 and 12 November 2020.

Funding: VC and AY have received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program under the grant agreement number 725594 (time-data) and the Swiss National Science Foundation (SNSF) under the grant number 200021_178865/1. JAT gratefully acknowledges ONR Awards N00014-11-1-0025, N00014-17-1-2146, and N00014-18-1-2363. MU gratefully acknowledges DARPA Award FA8750-17-2-0101. Part of this research is conducted while AY is at Massachusetts Institute of Technology, Cambridge, MA, USA. AY acknowledges the Early Postdoc.Mobility Fellowship P2ELP2_187955 from the Swiss National Science Foundation and partial postdoctoral support from the NSF-CAREER grant IIS-1846088.

[†]Laboratory for Information and Inference Systems, Department of Electrical Engineering, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland (alp.yurtsever@epfl.ch, volkan.cevher@epfl.ch, <https://lions.epfl.ch/>).

[‡]Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA, USA (jtropp@cms.caltech.edu, <http://users.cms.caltech.edu/~jtropp>).

[§]Laboratoire Traitement et Communication d’Information, Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France (olivier.fercoq@telecom-paris.fr, <https://perso.telecom-paristech.fr/ofercq/>).

[¶]Department of Operations Research and Information Engineering, Cornell University, Ithaca, NY, USA (udell@cornell.edu, <https://people.orie.cornell.edu/mru8/>).

than the Burer–Monteiro factorization heuristic [24, 20] for certain problem instances [97].

In addition to the theoretical guarantees, we offer evidence that SketchyCGAL is a practical optimization algorithm. For example, on a laptop equivalent, we can solve the MaxCut SDP for a sparse graph with over 20 million vertices, where the matrix variable has over 10^{14} entries. We also tackle large phase retrieval problems arising from Fourier ptychography [48], as well as relaxations [108, 23] of the quadratic assignment problem.

1.1. Example: The maximum cut in a graph. To begin, we derive a fundamental SDP [34, 35, 44] that arises in combinatorial optimization. This example highlights why large SDPs are hard to solve, and it illustrates the potential of our approach.

1.1.1. MaxCut. Consider an undirected graph $G = (V, E)$ comprising a vertex set $V = \{1, \dots, n\}$ and a set E of m edges. The combinatorial Laplacian of the graph is the real positive-semidefinite (psd) matrix

$$(1.1) \quad \mathbf{L} := \sum_{\{i,j\} \in E} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^* \in \mathbb{R}^{n \times n},$$

where $\mathbf{e}_i \in \mathbb{R}^n$ denotes the i th standard basis vector and $*$ refers to the (conjugate) transpose of a matrix or vector. We can search for a maximum-weight cut in the graph by solving

$$(1.2) \quad \text{maximize } \boldsymbol{\chi}^* \mathbf{L} \boldsymbol{\chi} \quad \text{subject to } \boldsymbol{\chi} \in \{\pm 1\}^n.$$

Unfortunately, the formulation (1.2) is NP-hard [55]. One remedy is to relax it to an SDP.

Consider the matrix $\mathbf{X} = \boldsymbol{\chi} \boldsymbol{\chi}^*$ where $\boldsymbol{\chi} \in \{\pm 1\}^n$. The matrix \mathbf{X} is psd; its diagonal entries equal one; and it has rank one. We can express the MaxCut problem (1.2) in terms of the matrix \mathbf{X} by rewriting the objective as a trace. Bringing forward the implicit constraints on \mathbf{X} and dropping the rank constraint, we arrive at the MaxCut SDP:

$$(1.3) \quad \text{maximize } \text{tr}(\mathbf{L} \mathbf{X}) \quad \text{subject to } \text{diag}(\mathbf{X}) = \mathbf{1}, \quad \mathbf{X} \text{ is psd.}$$

As usual, diag extracts the diagonal of a matrix as a vector, and $\mathbf{1} \in \mathbb{R}^n$ is the vector of ones.

The matrix solution \mathbf{X}_* of (1.3) does not immediately yield a cut. Let $\mathbf{x}_* \mathbf{x}_*^*$ be a best rank-one approximation of \mathbf{X}_* with respect to the Frobenius norm. Then the vector $\boldsymbol{\chi}_* = \text{sgn}(\mathbf{x}_*)$ is a valid cut. In many cases, the cut $\boldsymbol{\chi}_*$ yields an excellent solution to the discrete MaxCut problem (1.2). We can also use \mathbf{X}_* to compute a cut that is provably near-optimal via a more involved randomized rounding procedure [44].

1.1.2. What’s the issue? We specify an instance of the MaxCut SDP (1.3) by means of the Laplacian \mathbf{L} of the graph, which has $O(m + n)$ nonzero entries. Our goal is to compute a best rank-one approximation of a solution to the SDP, which has $O(n)$ degrees of freedom. In other words, the total cost of representing the input and output of the problem is $O(m + n)$. Sadly, the matrix variable in (1.3) seems to require storage $O(n^2)$. For example, a graph G with one million vertices leads to an SDP (1.3) with a trillion real variables.

Storage is one of the main reasons that it has been challenging to solve large instances of the MaxCut SDP reliably. Undeterred, we raise a question:

Can we provably find a best rank-one approximation of a solution to the MaxCut SDP (1.3) with storage $O(m + n)$? Can we achieve working storage $O(n)$?

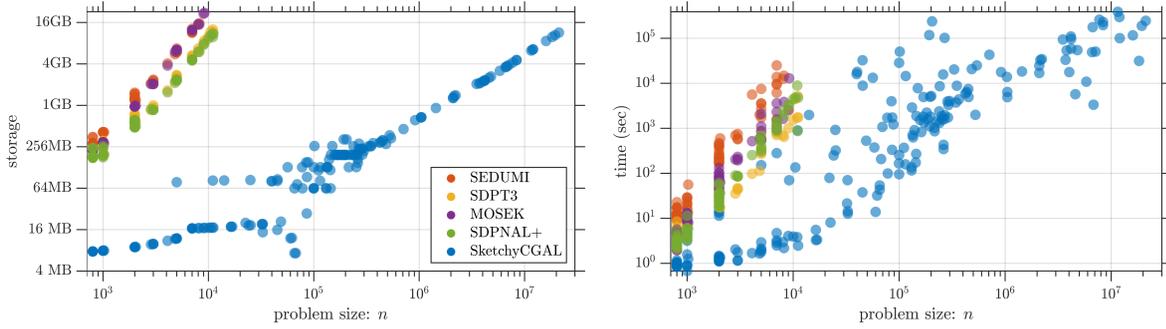


Figure 1.1. MaxCut SDP: Scalability. Storage cost [left] and runtime [right] of SketchyCGAL with sketch size $R = 10$ as compared with four standard SDP solvers. The relative error tolerance for each solver is 10^{-1} . Each marker represents one dataset. See subsection 7.2.3 for details.

We are not aware of any correct algorithm that can solve an arbitrary instance of (1.3) with a working storage guarantee better than $\Theta(\min\{m, n^{3/2}\})$; see section 8.

In addition to the limit on storage, a good algorithm should interact with the Laplacian L only through noninvasive, low-cost operations, such as matrix–vector multiplication.

1.1.3. A storage-optimal algorithm for the MaxCut SDP. Surprisingly, it is possible to achieve all the goals announced in the last subsection.

Theorem 1.1 (MaxCut via SketchyCGAL). For any $\varepsilon, \zeta > 0$ and any Laplacian $L \in \mathbb{R}^{n \times n}$, the SketchyCGAL algorithm computes a $(1+\zeta)$ -optimal rank-one approximation of an ε -optimal point of (1.3); see subsection 2.2. The working storage is $O(n/\zeta)$. The algorithm performs at most $\tilde{O}(\varepsilon^{-2.5})$ matrix–vector multiplies with the Laplacian L , plus lower-order arithmetic. The algorithm is randomized; it succeeds with high probability over its random choices.

Theorem 1.1 follows from Theorem 6.3. As usual, \tilde{O} suppresses constants and logarithmic factors. In contrast to Burer–Monteiro methods [97] and to the approximate complementarity paradigm [36], SketchyCGAL provably succeeds for every instance of MaxCut.

In our experience, SketchyCGAL works *better* than the theorem says. Figure 1.1 compares its scalability with four standard SDP solvers on a laptop equivalent (details in subsection 7.2).

1.2. A model problem. SketchyCGAL can solve all standard-form SDPs that satisfy strong duality. To simplify parts of the presentation, we focus on a model problem that includes an extra trace constraint. Section SM4 extends SketchyCGAL to a more expressive problem template that includes standard-form SDPs with additional (conic) inequality constraints.

1.2.1. The trace-constrained SDP. We work over the field $\mathbb{F} = \mathbb{R}$ or $\mathbb{F} = \mathbb{C}$. For each $n \in \mathbb{N}$, define the set $\mathbb{S}_n := \mathbb{S}_n(\mathbb{F})$ of (conjugate) symmetric $n \times n$ matrices with entries in \mathbb{F} .

Introduce the set of $n \times n$ psd matrices with trace one:

$$(1.4) \quad \Delta_n := \{X \in \mathbb{S}_n : \text{tr } X = 1 \text{ and } X \text{ is psd}\}.$$

Our model problem is the following trace-constrained SDP:

$$(1.5) \quad \begin{aligned} & \text{minimize} && \text{tr}(CX) \\ & \text{subject to} && \text{tr}(A_i X) = b_i \quad \text{for } i = 1, \dots, d \quad \text{and} \quad X \in \alpha \Delta_n. \end{aligned}$$

The trace parameter $\alpha > 0$, each matrix $\mathbf{C}, \mathbf{A}_1, \dots, \mathbf{A}_d \in \mathbb{S}_n$, and $b_1, \dots, b_d \in \mathbb{R}$. We always assume that (1.5) satisfies strong duality with its standard-form dual problem.

To solve a general standard-form SDP, we replace the inclusion $\mathbf{X} \in \alpha \Delta_n$ with the constraints that \mathbf{X} is psd and $\text{tr}(\mathbf{X}) \leq \alpha$ for a large enough parameter α .

1.2.2. Applications. The model problem (1.5) has diverse applications in statistics, signal processing, quantum information theory, combinatorics, and beyond. Evidently, the MaxCut SDP (1.3) is a special case. The template (1.5) includes problems in computer vision [49], in microscopy [48], and in robotics [84]. It also supports contemporary machine learning tasks, such as certifying robustness of neural networks [82]. There is a galaxy of other examples.

1.3. Complexity of SDP formulations and solutions. This section describes some special features that commonly appear in large, real-world SDPs. Our algorithm will take advantage of these features, even as it provides guarantees for every instance of (1.5).

1.3.1. Structure of the problem data. The matrices \mathbf{C} and \mathbf{A}_i that appear in (1.5) are often highly structured, or sparse, or have low-rank. As such, we can specify the SDP using a small amount of information. In our work, we exploit this property by treating the problem data for the SDP (1.5) as a collection of black boxes that support specific linear algebraic operations. The algorithm for solving the SDP only needs to access the data via these black boxes, and we can insist that these subroutines are implemented efficiently.

1.3.2. Low-rank solutions of SDPs. We will also capitalize on the fact that SDPs frequently have low-rank solutions, or the solutions are approximated well by low-rank matrices. There are several reasons why we can make this surmise.

Weakly-constrained SDPs. First, many SDPs have low-rank solutions just because they are weakly constrained. That is, the number d of linear equalities in (1.5) is much smaller than the number n^2 of components in the matrix variable. This situation often occurs in signal processing and statistics problems, where d reflects the amount of measured data. SketchyCGAL is designed for weakly constrained SDPs, but it does not require this property.

A weakly constrained SDP has at least one low-rank solution because of the geometry of the set of psd matrices; see [16, Prop. II(13.4) and Prob. II.14.5] and [79].

Fact 1.2 (Barvinok–Pataki). *When $\mathbb{F} = \mathbb{R}$, the SDP (1.5) has a solution with rank $r \leq \sqrt{2(d+1)}$. When $\mathbb{F} = \mathbb{C}$, there is a solution with rank $r \leq \sqrt{d+1}$.*

For example, the MaxCut SDP (1.3) admits a solution with rank $\sqrt{2(n+1)}$.

Although a weakly-constrained SDP can have solutions with high rank, a *generic* weakly-constrained SDP has a unique solution, which must be low-rank [5].

Fact 1.3 (Alizadeh et al.). *Let $\mathbb{F} = \mathbb{R}$. Except for a set of matrices $\{\mathbf{C}, \mathbf{A}_1, \dots, \mathbf{A}_d\}$ with measure zero, the solution set of the SDP (1.5) is a unique matrix with rank $r \leq \sqrt{2(d+1)}$.*

Matrix rank minimization. Second, some SDPs are *designed* to produce a low-rank matrix that satisfies a system of linear matrix equations. This idea can be traced to the control theory literature [68, 78], and it was explored thoroughly in Fazel’s thesis [37]. Early applications include Euclidean distance matrix completion [2] and collaborative filtering [89]. Extensive empirical work indicates that these SDPs often produce low-rank solutions.

Structural properties. There are other reasons that an SDP must have a low-rank solution. For instance, consider the optimal power flow SDPs developed by Lavaei and Low [60], where the rank of the solution is controlled by the geometry of the power grid.

1.3.3. Algorithms? To summarize, many realistic SDPs have structured data, and they admit solutions that are (close to) low rank. Are there algorithms that can exploit these features? Although there are a number of methods that attempt to do so, none can provably solve every SDP with limited arithmetic and minimal storage. See [section 8](#) for related work.

Why has it been so difficult to develop provably correct algorithms for finding low-rank solutions to structured SDPs? Most approaches that try to control the rank run headlong into a computational complexity barrier: For any fixed rank parameter r , it is NP-hard to solve the model problem (1.5) if the variable \mathbf{X} is also constrained to be a rank- r matrix [37, p. 7].

To escape this sticky fact, we revise the computational goal, following [106]. The key insight is to seek a rank- r matrix that *approximates a solution* to (1.5). See [subsection 2.2](#) for a detailed explanation. This shift in perspective opens up new algorithmic prospects.

1.4. Contributions. Inspired by [106], we derive an algorithm that harnesses the favorable properties common in large SDPs. The SketchyCGAL algorithm solves the SDP (1.5) using a primal–dual optimization method [102] developed by a subset of the authors. Each iteration requires one coarse eigenvector computation [57] and leads to a rank-one update of the psd matrix variable. Instead of storing this matrix, we maintain a compressed representation by means of a matrix sketching technique [93]. After the optimization algorithm terminates, we extract from the sketch a low-rank approximation of a solution of the SDP. This idea leads to a practical, provably correct SDP solver that economizes on storage and arithmetic.

Theorem 1.4 (The model problem via SketchyCGAL). *Assume that the model problem (1.5) satisfies strong duality. For any $\varepsilon, \zeta > 0$ and any rank parameter r , the SketchyCGAL algorithm computes a $(1 + \zeta)$ -optimal rank- r approximation of an ε -optimal point of (1.5); see [subsection 2.2](#). The storage cost is $O(d + rn/\zeta)$. Most of the arithmetic consists in $\tilde{O}(\varepsilon^{-2.5})$ matrix–vector products with each matrix $\mathbf{C}, \mathbf{A}_1, \dots, \mathbf{A}_d$ from the problem data. The algorithm succeeds with high probability.*

[Theorem 6.3](#) contains full theoretical details. Note that the storage $\Theta(d + rn)$ is the minimum possible for any primal–dual algorithm that returns a rank- r solution to (1.5). [Section 7](#) contains numerical evidence that the algorithm is effective for a range of examples.

1.5. Roadmap. [Section 2](#) presents an abstract framework for studying SDPs that exposes the challenges associated with large problems. [Section 3](#) outlines a primal–dual algorithm, called CGAL, for the model problem (1.5). [Sections 4](#) and [5](#) introduce methods from randomized linear algebra that we use to control storage and arithmetic costs. [Section 6](#) develops the SketchyCGAL algorithm, its convergence theory, and its resource usage guarantees. [Section 7](#) contains a numerical study of SketchyCGAL, and [section 8](#) covers related work.

1.6. Notation. The symbol $\|\cdot\|_{\mathbb{F}}$ denotes the Frobenius norm, while $\|\cdot\|_*$ is the nuclear norm (i.e., Schatten-1). The unadorned norm $\|\cdot\|$ refers to the ℓ_2 norm of a vector, the spectral norm of a matrix, or the operator norm of a linear map from $(\mathbb{S}_n, \|\cdot\|_{\mathbb{F}})$ to $(\mathbb{R}^d, \|\cdot\|)$. We write $\langle \cdot, \cdot \rangle$ for both the ℓ_2 inner product on vectors and the trace inner product on matrices.

The map $\llbracket \mathbf{M} \rrbracket_r$ returns an r -truncated singular-value decomposition of the matrix \mathbf{M} , which is a best rank- r approximation with respect to every unitarily invariant norm [69].

We use the standard computer science interpretation of the asymptotic notation $\mathcal{O}, \tilde{\mathcal{O}}, \Theta$.

2. Scalable semidefinite programming. To solve the model problem (1.5) efficiently, we need to exploit structure inherent in the problem data. This section outlines an abstract approach that directs our attention to the core computational difficulties.

2.1. Abstract form of the model problem. Let us instate compact notation for the linear constraints in the model problem (1.5). Define a linear map \mathcal{A} and its adjoint \mathcal{A}^* via

$$(2.1) \quad \begin{aligned} \mathcal{A} : \mathbb{S}_n &\rightarrow \mathbb{R}^d & \text{where } \mathcal{A}\mathbf{X} &= [\langle \mathbf{A}_1, \mathbf{X} \rangle \ \dots \ \langle \mathbf{A}_d, \mathbf{X} \rangle]; \\ \mathcal{A}^* : \mathbb{R}^d &\rightarrow \mathbb{S}_n & \text{where } \mathcal{A}^* \mathbf{z} &= \sum_{i=1}^d z_i \mathbf{A}_i. \end{aligned}$$

We bound the linear map with the operator norm $\|\mathcal{A}\| := \|\mathcal{A}\|_{\mathbb{F} \rightarrow \ell_2}$. Form the vector $\mathbf{b} := (b_1, \dots, b_d) \in \mathbb{R}^d$ of constraint values. In this notation, (1.5) becomes

$$(2.2) \quad \text{minimize } \langle \mathbf{C}, \mathbf{X} \rangle \quad \text{subject to } \mathcal{A}\mathbf{X} = \mathbf{b}, \quad \mathbf{X} \in \alpha \Delta_n.$$

Problem instances are parameterized by the tuple $(\mathbf{C}, \mathcal{A}, \mathbf{b}, \alpha)$.

2.2. Approximate solutions. Let \mathbf{X}_* be a solution to the model problem (2.2). For $\varepsilon \geq 0$, we say that a matrix \mathbf{X}_ε is ε -optimal for (2.2) when

$$\|\mathcal{A}\mathbf{X}_\varepsilon - \mathbf{b}\| \leq \varepsilon \quad \text{and} \quad \langle \mathbf{C}, \mathbf{X}_\varepsilon \rangle - \langle \mathbf{C}, \mathbf{X}_* \rangle \leq \varepsilon.$$

Many optimization algorithms aim to produce such ε -optimal points; cf. section 8.

As we saw in subsection 1.3.2, there are many situations where the solutions to (2.2) have low rank or they admit accurate low-rank approximations. The ε -optimal points inherit these properties for sufficiently small ε . This insight suggests a new computational goal.

For a rank parameter r , we will seek a rank- r matrix $\widehat{\mathbf{X}}$ that approximates an ε -optimal point \mathbf{X}_ε . More precisely, for a fixed suboptimality parameter $\zeta > 0$, we want

$$(2.3) \quad \|\mathbf{X}_\varepsilon - \widehat{\mathbf{X}}\|_* \leq (1 + \zeta) \cdot \|\mathbf{X}_\varepsilon - \llbracket \mathbf{X}_\varepsilon \rrbracket_r\|_* \quad \text{where } \text{rank } \widehat{\mathbf{X}} \leq r \text{ and } \mathbf{X}_\varepsilon \text{ is } \varepsilon\text{-optimal.}$$

Given (2.3), if the ε -optimal point \mathbf{X}_ε is close to *any* rank- r matrix, then the rank- r approximate solution $\widehat{\mathbf{X}}$ is also close to the ε -optimal point \mathbf{X}_ε . This formulation is advantageous because it is easier to compute and to store the low-rank matrix $\widehat{\mathbf{X}}$.

2.3. Black-box presentation of problem data. To develop scalable algorithms for (2.2), it is productive to hide the internal complexity of the problem instance from the algorithm [106]. To do so, we treat \mathbf{C} and \mathcal{A} as black boxes that support three primitive computations:

$$(2.4) \quad \begin{array}{|c|} \hline \textcircled{1} \quad \mathbf{u} \mapsto \mathbf{C}\mathbf{u} \\ \mathbb{R}^n \rightarrow \mathbb{R}^n \\ \hline \end{array} \quad \begin{array}{|c|} \hline \textcircled{2} \quad (\mathbf{u}, \mathbf{z}) \mapsto (\mathcal{A}^* \mathbf{z})\mathbf{u} \\ \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}^n \\ \hline \end{array} \quad \begin{array}{|c|} \hline \textcircled{3} \quad \mathbf{u} \mapsto \mathcal{A}(\mathbf{u}\mathbf{u}^*) \\ \mathbb{R}^n \rightarrow \mathbb{R}^d \\ \hline \end{array}$$

The vectors $\mathbf{u} \in \mathbb{R}^n$ and $\mathbf{z} \in \mathbb{R}^d$ are arbitrary. Although these functions may seem abstract, they are often quite natural and easy to implement well. For example, see subsection 2.4.

We will formulate algorithms for (2.2) that interact with the problem data only through the operations (2.4). We tacitly assume that the primitives require minimal storage and arithmetic; otherwise, it may be impossible to develop a truly efficient algorithm.

2.4. Example: The MaxCut SDP. To provide a concrete example, let us summarize the meaning of the primitives and the desired storage costs for solving the MaxCut SDP (1.3).

- The primitive ① : $\mathbf{u} \mapsto -\mathbf{L}\mathbf{u}$. In the typical case that the Laplacian \mathbf{L} is sparse, this amounts to a sparse matrix–vector multiply.
- The primitive ② : $(\mathbf{u}, \mathbf{z}) \mapsto (\text{diag}^* \mathbf{z})\mathbf{u} = (z_1 u_1, \dots, z_n u_n)$.
- The primitive ③ : $\mathbf{u} \mapsto \text{diag}(\mathbf{u}\mathbf{u}^*) = (|u_1|^2, \dots, |u_n|^2)$.
- The MaxCut SDP has n linear constraints, and we seek a rank-one approximation of the solution. Thus, we desire an algorithm that operates with $\Theta(n)$ working storage.

Note that we still need $\Theta(m)$ numbers to store the Laplacian \mathbf{L} of a generic graph with m edges. But we do not charge the optimization algorithm for this storage because the algorithm only interacts with \mathbf{L} through the primitives (2.4).

3. An algorithm for the model problem. We will develop a scalable method for the model problem (2.2) by enhancing an existing algorithm, called CGAL [102], developed by a subset of the authors. This method works well, but it is not as scalable as we would like because it lacks storage and arithmetic guarantees. This section summarizes the CGAL method and its convergence properties. Subsequent sections introduce additional ideas that we need to control resource usage, culminating with the SketchyCGAL algorithm in section 6.

3.1. The augmented problem. For a parameter $\beta > 0$, we revise the problem (2.2) by introducing an extra term in the objective:

$$(3.1) \quad \text{minimize} \quad \langle \mathbf{C}, \mathbf{X} \rangle + \frac{\beta}{2} \|\mathcal{A}\mathbf{X} - \mathbf{b}\|^2 \quad \text{subject to} \quad \mathcal{A}\mathbf{X} = \mathbf{b}, \quad \mathbf{X} \in \alpha\Delta_n.$$

The original problem (2.2) and the augmented problem (3.1) share the same optimal value and optimal set. But the new formulation has several benefits: the augmented objective cooperates with the affine constraint to penalize infeasible points, and the dual of the augmented problem is smooth, whereas the dual of the original problem is not. See [18, Ch. 2] for background.

3.2. The augmented Lagrangian. We construct the Lagrangian L_β of the augmented problem (3.1) by introducing a dual variable $\mathbf{y} \in \mathbb{R}^d$ and promoting the affine constraint:

$$(3.2) \quad L_\beta(\mathbf{X}; \mathbf{y}) := \langle \mathbf{C}, \mathbf{X} \rangle + \langle \mathbf{y}, \mathcal{A}\mathbf{X} - \mathbf{b} \rangle + \frac{\beta}{2} \|\mathcal{A}\mathbf{X} - \mathbf{b}\|^2 \quad \text{for } \mathbf{X} \in \alpha\Delta_n \text{ and } \mathbf{y} \in \mathbb{R}^d.$$

For reference, note that the partial derivatives of the augmented Lagrangian L_β satisfy

$$(3.3) \quad \partial_{\mathbf{X}} L_\beta(\mathbf{X}; \mathbf{y}) = \mathbf{C} + \mathcal{A}^* \mathbf{y} + \beta \mathcal{A}^*(\mathcal{A}\mathbf{X} - \mathbf{b}) \quad \text{and} \quad \partial_{\mathbf{y}} L_\beta(\mathbf{X}; \mathbf{y}) = \mathcal{A}\mathbf{X} - \mathbf{b}.$$

We attempt to minimize the augmented Lagrangian L_β with respect to the primal variable \mathbf{X} , while we attempt to maximize with respect to the dual variable \mathbf{y} .

Algorithm 3.1 CGAL for the model problem (2.2)**Input:** Problem data for (2.2) implemented via the primitives (2.4); number T of iterations**Output:** Approximate solution \mathbf{X}_T to (2.2)

```

1 function CGAL( $T$ )
2   Scale problem data (subsection 7.1.1) ▷ [opt] Recommended!
3    $\beta_0 \leftarrow 1$  and  $K \leftarrow +\infty$  ▷ Default parameters
4    $\mathbf{X} \leftarrow \mathbf{0}_{n \times n}$  and  $\mathbf{y} \leftarrow \mathbf{0}_d$ 
5   for  $t \leftarrow 1, 2, 3, \dots, T$  do
6      $\beta \leftarrow \beta_0 \sqrt{t+1}$  and  $\eta \leftarrow 2/(t+1)$ 
7      $(\xi, \mathbf{v}) \leftarrow \text{ApproxMinEvec}(\mathbf{C} + \mathcal{A}^*(\mathbf{y} + \beta(\mathcal{A}\mathbf{X} - \mathbf{b})); q_t)$  ▷ Algorithm 4.1 with  $q_t = t^{1/4} \log n$ 
▷ Implement with primitives (2.4)!
8      $\mathbf{X} \leftarrow (1 - \eta)\mathbf{X} + \eta(\alpha \mathbf{v}\mathbf{v}^*)$ 
9      $\mathbf{y} \leftarrow \mathbf{y} + \gamma(\mathcal{A}\mathbf{X} - \mathbf{b})$  ▷ Step size  $\gamma$  satisfies (3.8)

```

3.3. The augmented Lagrangian strategy. The form of the augmented Lagrangian suggests an algorithm. We generate a sequence $\{(\mathbf{X}_t; \mathbf{y}_t)\}$ of primal–dual pairs by alternately minimizing over the primal variable \mathbf{X} and taking a gradient step in the dual variable \mathbf{y} .

1. **Initialize:** Choose $\mathbf{X}_1 \in \alpha \Delta_n$ and $\mathbf{y}_1 \in \mathbb{R}^d$.
2. **Primal step:** $\mathbf{X}_{t+1} \in \arg \min \{L_\beta(\mathbf{X}; \mathbf{y}_t) : \mathbf{X} \in \alpha \Delta_n\}$.
3. **Dual step:** $\mathbf{y}_{t+1} = \mathbf{y}_t + \beta(\mathcal{A}\mathbf{X} - \mathbf{b})$.

As we proceed, we can also increase the smoothing parameter β to make violations of the affine constraint in (3.1) more and more intolerable.

The augmented Lagrangian strategy is powerful, but it is hard to apply in this setting because of the cost of implementing the primal step, even approximately.

3.4. The CGAL iteration. The CGAL iteration [102] is related to the augmented Lagrangian paradigm, but the primal steps are inspired by the conditional gradient method [39]. CGAL identifies an update direction for the primal variable by minimizing a linear proxy for the augmented Lagrangian. We take a small primal step in the update direction, and we improve the dual variable by taking a small gradient step. At each iteration, the smoothing parameter increases according to a fixed schedule.

This subsection outlines the steps in the CGAL iteration. Afterward, we give *a priori* guarantees on the convergence rate. Pseudocode for CGAL appears as Algorithm 3.1.

3.4.1. Initialization. Let $\beta_0 > 0$ be an initial smoothing parameter, and fix a (large) bound $K > 0$ on the maximum allowable size of the dual variable. We also assume that we have access to the norm $\|\mathcal{A}\|$ of the constraint matrix, or—failing that—a *lower* bound. Begin with an arbitrary choice $\mathbf{X}_1 \in \mathbb{S}_n$ for the primal variable; set the initial dual variable $\mathbf{y}_1 = \mathbf{0}$.

3.4.2. Primal updates via linear minimization. At iteration $t = 1, 2, 3, \dots$, we increase the smoothing parameter to $\beta_t := \beta_0 \sqrt{t+1}$, and we form the partial derivative of the augmented Lagrangian with respect to the primal variable at the current pair of iterates:

$$(3.4) \quad \mathbf{D}_t := \partial_{\mathbf{X}} L_{\beta_t}(\mathbf{X}_t; \mathbf{y}_t) = \mathbf{C} + \mathcal{A}^*(\mathbf{y}_t + \beta_t(\mathcal{A}\mathbf{X}_t - \mathbf{b})).$$

Then we compute an update $\mathbf{H}_t \in \alpha\Delta_t$ by finding the feasible point that is most correlated with the negative gradient $-\mathbf{D}_t$. This amounts to a linear minimization problem:

$$(3.5) \quad \begin{aligned} \mathbf{H}_t &\in \arg \min \{ \langle \mathbf{D}_t, \mathbf{H} \rangle : \mathbf{H} \in \alpha\Delta_n \} \\ &= \text{convex hull} \{ \alpha \mathbf{v} \mathbf{v}^* : \mathbf{v} \text{ is a unit-norm minimum eigenvector of } \mathbf{D}_t \}. \end{aligned}$$

In particular, we may take $\mathbf{H}_t = \alpha \mathbf{v}_t \mathbf{v}_t^*$ for a minimum eigenvector \mathbf{v}_t of the gradient \mathbf{D}_t . Next, update the matrix primal variable by taking a small step in the direction \mathbf{H}_t :

$$(3.6) \quad \mathbf{X}_{t+1} := (1 - \eta_t) \mathbf{X}_t + \eta_t \mathbf{H}_t \in \alpha\Delta_n \quad \text{where} \quad \eta_t := \frac{2}{t+1}.$$

The appeal of this approach is that it only requires a single eigenvector computation (3.5). Moreover, we need not compute the eigenvector accurately; see subsection 3.4.5 for details.

3.4.3. Dual updates via gradient ascent. Next, we update the dual variable by taking a small gradient step on the dual variable in the augmented Lagrangian:

$$(3.7) \quad \mathbf{y}_{t+1} = \mathbf{y}_t + \gamma_t \partial_{\mathbf{y}} L_{\beta_t}(\mathbf{X}_{t+1}; \mathbf{y}_t) = \mathbf{y}_t + \gamma_t (\mathcal{A} \mathbf{X}_{t+1} - \mathbf{b}).$$

The dual step size γ_t is the largest number that satisfies the conditions

$$(3.8) \quad \gamma_t \|\mathcal{A} \mathbf{X}_{t+1} - \mathbf{b}\|^2 \leq \frac{4\alpha^2 \beta_0}{(t+1)^{3/2}} \|\mathcal{A}\|^2 \quad \text{and} \quad 0 \leq \gamma_t \leq \beta_0.$$

We omit the dual step if it makes the dual variable too large. More precisely, if (3.7) and (3.8) result in $\|\mathbf{y}_{t+1}\| > K$, then we set $\mathbf{y}_{t+1} = \mathbf{y}_t$ instead. This is the CGAL iteration.

3.4.4. Assessing solution quality. Given a primal–dual pair $(\mathbf{X}_t; \mathbf{y}_t)$, we can assess the quality of the primal solution to the model problem (2.2). First, note that we can simply compute the infeasibility: $\mathcal{A} \mathbf{X}_t - \mathbf{b}$. Second, we can bound the suboptimality of the primal objective value. To do so, define the *surrogate duality gap*

$$(3.9) \quad g_t(\mathbf{X}) := \max_{\mathbf{H} \in \alpha\Delta_n} \langle \mathbf{D}_t, \mathbf{X} - \mathbf{H} \rangle = \langle \mathbf{C}, \mathbf{X}_t \rangle + \langle \mathbf{y}_t + \beta_t (\mathcal{A} \mathbf{X}_t - \mathbf{b}), \mathcal{A} \mathbf{X}_t \rangle - \lambda_{\min}(\mathbf{D}_t).$$

The latter expression follows from (3.4) and (3.5). As usual, λ_{\min} is the minimum eigenvalue. The suboptimality of \mathbf{X}_t is bounded as

$$(3.10) \quad \langle \mathbf{C}, \mathbf{X}_t \rangle - \langle \mathbf{C}, \mathbf{X}_\star \rangle \leq g_t(\mathbf{X}_t) - \langle \mathbf{y}_t, \mathcal{A} \mathbf{X}_t - \mathbf{b} \rangle - \frac{1}{2} \beta_t \|\mathcal{A} \mathbf{X}_t - \mathbf{b}\|^2,$$

where \mathbf{X}_\star is a primal optimal point. See subsection SM1.7 for the derivation of (3.10).

For CGAL, the surrogate duality gap (3.9) is analogous to the Frank–Wolfe duality gap [50]. It offers a practical tool for detecting early convergence. Unfortunately, we have not been able to prove that the CGAL error bound (3.10) converges to zero as the algorithm converges.

3.4.5. Approximate eigenvector computations. A crucial fact is that the CGAL strategy provably works if we replace (3.5) with an approximate minimum eigenvector computation. At step t , suppose that we find an update direction

$$(3.11) \quad \mathbf{H}_t := \alpha \mathbf{v}_t \mathbf{v}_t^* \quad \text{where} \quad \mathbf{v}_t^* \mathbf{D}_t \mathbf{v}_t \leq \lambda_{\min}(\mathbf{D}_t) + \frac{1}{\sqrt{t+1}} \|\mathbf{D}_t\|$$

for a unit vector $\mathbf{v}_t \in \mathbb{F}^n$. The matrix \mathbf{H}_t defined in (3.11) serves in place of a solution to (3.5) throughout the algorithm. We discuss solvers for the subproblem (3.11) in section 4.

3.5. Convergence guarantees for CGAL. The following result describes the convergence of the CGAL iteration. The analysis is adapted from [102, Thm. 3.1]; see section SM1 for details and a recapitulation of the proof.

Fact 3.1 (CGAL: Convergence). *Assume problem (2.2) satisfies strong duality. The CGAL iteration (subsection 3.4) with approximate eigenvector computations (3.11) yields a sequence $\{\mathbf{X}_t : t = 1, 2, 3, \dots\} \subset \alpha\Delta_n$ that satisfies*

$$(3.12) \quad \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\| \leq \frac{\text{Const}}{\sqrt{t}} \quad \text{and} \quad |\langle \mathbf{C}, \mathbf{X}_t \rangle - \langle \mathbf{C}, \mathbf{X}_\star \rangle| \leq \frac{\text{Const}}{\sqrt{t}}.$$

The matrix \mathbf{X}_\star solves (2.2). The constant depends on the problem data $(\mathbf{C}, \mathcal{A}, \mathbf{b}, \alpha)$, the minimum Euclidean norm of a dual solution, and the algorithm parameters β_0 and K .

In view of (3.12), CGAL produces a primal iterate \mathbf{X}_T that is ε -optimal after $T = O(\varepsilon^{-2})$ iterations. The numerical work in [102, Sec. 5] shows that CGAL finds an ε -optimal point after $T = O(\varepsilon^{-1})$ iterations for realistic problem instances. See also subsection SM5.3.

The analysis behind Fact 3.1 indicates that CGAL converges more quickly if we precondition the problem data by rescaling; see subsection 7.1.1. This step is critical in practice.

4. Approximate eigenvectors. Most of the computation in CGAL occurs in the approximate eigenvector step (3.11). This section describes our approach to this subproblem.

4.1. Krylov methods. The approximate minimum eigenvector computation (3.11) involves a matrix of the form $\mathbf{D}_t = \mathbf{C} + \mathcal{A}^* \mathbf{w}_t$. We can multiply the matrix \mathbf{D}_t by a vector using the primitives (2.4) 1–2. Krylov methods compute eigenvectors of \mathbf{D}_t by repeated matrix–vector multiplication with \mathbf{D}_t , so they are obvious tools for the subproblem (3.11).

Unfortunately, CGAL tends to generate matrices \mathbf{D}_t that have clustered eigenvalues. These matrices challenge standard eigenvector software, such as ARPACK [62], the engine behind the MATLAB command `eigs`. Instead, we retreat to a more classical technique.

4.2. A storage-optimal randomized Lanczos method. We use a nonstandard implementation of the randomized Lanczos method [57] to find an approximate eigenvector with minimal storage. This hinges on two ideas. First, we run the Lanczos iteration with a random initial vector, which supports convergence guarantees. Second, we do not store the Lanczos vectors, but rather regenerate them to construct the approximate eigenvector. See Algorithm 4.1 for pseudocode. Kuczyński & Woźniakowski [57, Thm. 4.2(a)] have obtained error bounds.

Fact 4.1 (Randomized Lanczos method). *Let $\mathbf{M} \in \mathbb{S}_n$. For $\varepsilon \in (0, 1]$ and $\delta \in (0, 0.5]$, the randomized Lanczos method, Algorithm 4.1, computes a unit vector $\mathbf{u} \in \mathbb{F}^n$ that satisfies*

$$\mathbf{u}^* \mathbf{M} \mathbf{u} \leq \lambda_{\min}(\mathbf{M}) + \frac{\varepsilon}{8} \|\mathbf{M}\| \quad \text{with probability at least } 1 - 2\delta$$

after $q \geq \frac{1}{2} + \varepsilon^{-1/2} \log(n/\delta^2)$ iterations. The arithmetic cost is at most q matrix–vector multiplies with \mathbf{M} and $O(qn)$ extra operations. The working storage is $O(n + q)$.

With constant probability, we solve the eigenvector problem (3.11) successfully in every iteration t of CGAL if we use $q_t = O(t^{1/4} \log(tn))$ iterations of Algorithm 4.1. In practice, we implement CGAL with $q_t = t^{1/4} \log n$. Although the Lanczos method has complicated numerical behavior in finite-precision arithmetic, it serves well as a subroutine within CGAL.

Algorithm 4.1 ApproxMinEvec via storage-optimal randomized Lanczos (subsection 4.2).

Input: Input matrix $M \in \mathbb{S}_n$ and maximum number q of iterations

Output: Approximate minimum eigenpair $(\xi, \mathbf{v}) \in \mathbb{R} \times \mathbb{F}^n$ of the matrix M

```

1 function ApproxMinEvec( $M$ ;  $q$ )
2    $\mathbf{v}_1 \leftarrow \text{randn}(n, 1)$ 
3    $\mathbf{v}_1 \leftarrow \mathbf{v}_1 / \|\mathbf{v}_1\|$  ▷ Store  $\mathbf{v}_1$  for reuse
4   for  $i \leftarrow 1, 2, 3, \dots, \min\{q, n - 1\}$  do ▷ During loop, store only  $\mathbf{v}_i$  and  $\mathbf{v}_{i+1}$ 
5      $\omega_i \leftarrow \text{Re}(\mathbf{v}_i^*(M\mathbf{v}_i))$ 
6      $\mathbf{v}_{i+1} \leftarrow M\mathbf{v}_i - \omega_i\mathbf{v}_i - \rho_{i-1}\mathbf{v}_{i-1}$  ▷ Three-term Lanczos recurrence;  $\rho_0\mathbf{v}_0 = \mathbf{0}$ 
7      $\rho_i \leftarrow \|\mathbf{v}_{i+1}\|$ 
8     if  $\rho_i = 0$  then break ▷ Found an invariant subspace!
9      $\mathbf{v}_{i+1} \leftarrow \mathbf{v}_{i+1} / \rho_i$ 
10     $T \leftarrow \text{tridiag}(\rho_{1:(i-1)}, \omega_{1:i}, \rho_{1:(i-1)})$  ▷ Form tridiagonal matrix
11     $[\xi, \mathbf{u}] \leftarrow \text{MinEvec}(T)$  ▷ Exploit tridiagonal form [77, Ch. 7], or just use eig
12     $\mathbf{v} \leftarrow \sum_{j=1}^i u_j \mathbf{v}_j$  ▷ Modify lines 4–9 to regenerate  $\mathbf{v}_2, \dots, \mathbf{v}_i$  and form sum

```

Remark 4.2 (Time–storage tradeoff). We can retain the vectors \mathbf{v}_i in Algorithm 4.1 to reduce the flop count by approximately a factor two, but the storage increases to $O(qn)$.

5. Sketching and reconstruction of a psd matrix. The CGAL iteration generates a psd matrix that solves the model problem (2.2) via a sequence (3.6) of rank-one linear updates. To control storage costs, SketchyCGAL retains only a compressed version of the psd matrix variable \mathbf{X}_t . This section outlines the *Nyström sketch*, an established method [45, 43, 64, 93] that can track the evolving psd matrix and then report a provably accurate low-rank approximation.

For more information about the implementation and behavior of low-rank matrix approximation from streaming data, see section SM2 and our papers [93, 94, 95].

5.1. Sketching and updates. Consider a psd input matrix $\mathbf{X} \in \mathbb{S}_n$. Let R be a parameter that modulates the storage cost of the sketch and the quality of the matrix approximation.

To construct the Nyström sketch, we draw and fix a standard normal¹ test matrix $\mathbf{\Omega} \in \mathbb{F}^{n \times R}$. Our summary, or *sketch*, of the matrix \mathbf{X} takes the form

$$(5.1) \quad \mathbf{S} = \mathbf{X}\mathbf{\Omega} \in \mathbb{F}^{n \times R}.$$

The sketch \mathbf{S} supports linear rank-one updates to \mathbf{X} . Indeed, we can track the evolution

$$(5.2) \quad \begin{aligned} \mathbf{X} &\leftarrow (1 - \eta)\mathbf{X} + \eta\mathbf{v}\mathbf{v}^* && \text{for } \eta \in [0, 1] \text{ and } \mathbf{v} \in \mathbb{F}^n \\ \text{via } \mathbf{S} &\leftarrow (1 - \eta)\mathbf{S} + \eta\mathbf{v}(\mathbf{v}^*\mathbf{\Omega}). \end{aligned}$$

The test matrix $\mathbf{\Omega}$ and the sketch \mathbf{S} require storage of $2Rn$ numbers in \mathbb{F} . The arithmetic cost of the linear update (5.2) to the sketch is $\Theta(Rn)$ numerical operations.

Remark 5.1 (Structured random matrices). We can reduce storage costs by a factor of two by using a structured random matrix in place of $\mathbf{\Omega}$. For example, see [95, Sec. 3] or [91].

¹Each entry of the matrix is an independent Gaussian random variable with mean zero and variance one. In the complex setting, the real and imaginary parts of each entry are independent standard normal variables.

Algorithm 5.1 NystromSketch implementation (see section 5)**Input:** Dimension n of input matrix, size R of sketch**Output:** Rank- R approximation $\widehat{\mathbf{X}}$ of sketched matrix in factored form $\widehat{\mathbf{X}} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^*$, where $\mathbf{U} \in \mathbb{F}^{n \times R}$ has orthonormal columns and $\mathbf{\Lambda} \in \mathbb{R}^{R \times R}$ is nonnegative diagonal

```

1 function NystromSketch.Init( $n, R$ )
2    $\mathbf{\Omega} \leftarrow \text{randn}(n, R)$  ▷ Draw and fix random test matrix
3    $\mathbf{S} \leftarrow \text{zeros}(n, R)$  ▷ Form sketch of zero matrix

4 function NystromSketch.RankOneUpdate( $\mathbf{v}, \eta$ ) ▷ Implements (5.2)
5    $\mathbf{S} \leftarrow (1 - \eta)\mathbf{S} + \eta\mathbf{v}(\mathbf{v}^*\mathbf{\Omega})$  ▷ Update sketch of matrix

6 function NystromSketch.Reconstruct()
7    $\sigma \leftarrow \sqrt{n} \text{eps}(\text{norm}(\mathbf{S}))$  ▷ Compute a shift parameter
8    $\mathbf{S}_\sigma \leftarrow \mathbf{S} + \sigma\mathbf{\Omega}$  ▷ Implicitly form sketch of  $\mathbf{X} + \sigma\mathbf{I}$ 
9    $\mathbf{L} \leftarrow \text{chol}(\mathbf{\Omega}^*\mathbf{S}_\sigma)$ 
10   $[\mathbf{U}, \mathbf{\Sigma}, \sim] \leftarrow \text{svd}(\mathbf{S}_\sigma/\mathbf{L})$  ▷ Dense SVD
11   $\mathbf{\Lambda} \leftarrow \max\{0, \mathbf{\Sigma}^2 - \sigma\mathbf{I}\}$  ▷ Remove shift

```

5.2. The reconstruction process. Given the test matrix $\mathbf{\Omega}$ and the sketch $\mathbf{S} = \mathbf{X}\mathbf{\Omega}$, we form a rank- R approximation $\widehat{\mathbf{X}}$ of the sketched matrix \mathbf{X} . The approximation is defined by

$$(5.3) \quad \widehat{\mathbf{X}} := \mathbf{S}(\mathbf{\Omega}^*\mathbf{S})^\dagger \mathbf{S}^* = (\mathbf{X}\mathbf{\Omega})(\mathbf{\Omega}^*\mathbf{X}\mathbf{\Omega})^\dagger (\mathbf{X}\mathbf{\Omega})^*,$$

where † is the pseudoinverse. This reconstruction is called a *Nyström approximation*. We often truncate $\widehat{\mathbf{X}}$ by replacing it with its best rank- r approximation $[\widehat{\mathbf{X}}]_r$ for some $r \leq R$.

See Algorithm 5.1 for a numerically stable implementation of the Nyström approximation (5.3). The algorithm takes $\Theta(R^2n)$ numerical operations and $\Theta(Rn)$ storage.

5.3. A priori error bounds. The Nyström approximation $\widehat{\mathbf{X}}$ yields a provably good estimate for the matrix \mathbf{X} contained in the sketch [93, Thm. 4.1].

Fact 5.2 (Nyström sketch: Error bound). Fix a psd matrix $\mathbf{X} \in \mathbb{S}_n$. Let $\mathbf{S} = \mathbf{X}\mathbf{\Omega}$ where $\mathbf{\Omega} \in \mathbb{F}^{n \times R}$ is standard normal. For each $r < R - 1$, the Nyström approximation (5.3) satisfies

$$(5.4) \quad \mathbb{E}_\Omega \|\mathbf{X} - \widehat{\mathbf{X}}\|_* \leq \left(1 + \frac{r}{R - r - 1}\right) \cdot \|\mathbf{X} - [\mathbf{X}]_r\|_*,$$

where \mathbb{E}_Ω is the expectation with respect to $\mathbf{\Omega}$. If we replace $\widehat{\mathbf{X}}$ with its rank- r truncation $[\widehat{\mathbf{X}}]_r$, the error bound (5.4) remains valid. Similar results hold with high probability.

6. Scalable semidefinite programming via SketchyCGAL. We may now present an extension of CGAL that solves the model problem (2.2) with controlled storage and computation. Our new algorithm, SketchyCGAL, enhances the CGAL iteration from subsection 3.4. Instead of storing the matrix \mathbf{X}_t in the CGAL iteration,

1. We drive the iteration with the d -dimensional primal state variable $\mathbf{z}_t := \mathcal{A}\mathbf{X}_t$.
2. We maintain a Nyström sketch of the primal iterate \mathbf{X}_t using storage $\Theta(Rn)$.
3. When the iteration is halted, say at time T , we use the sketch to construct a rank- R approximation $\widehat{\mathbf{X}}_T$ of the implicitly computed solution \mathbf{X}_T of the model problem.

As we will see, the resulting method exhibits almost the same convergence behavior as the CGAL algorithm, but it also enjoys a strong storage guarantee (when $d \ll n^2$ and $R \ll n$).

6.1. The SketchyCGAL iteration. To develop the SketchyCGAL iteration, we start with the CGAL iteration. Then we make the substitutions $\mathbf{z}_t = \mathcal{A}\mathbf{X}_t$ and $\mathbf{h}_t = \mathcal{A}\mathbf{H}_t$ to eliminate the matrix variables. Let us summarize what happens; see [subsection 6.2](#) for additional explanation. [Algorithm 6.1](#) contains pseudocode with implementation recommendations.

6.1.1. Initialization. First, we choose the size R of the Nyström sketch. Then draw and fix the random test matrix $\mathbf{\Omega} \in \mathbb{F}^{n \times R}$. Select an initial smoothing parameter β_0 and a (finite) bound K on the dual variable. Define the smoothing parameters $\beta_t := \beta_0 \sqrt{t+1}$ and the step size parameters $\eta_t := 2/(t+1)$. Initialize the primal state variable $\mathbf{z}_1 := \mathbf{0} \in \mathbb{R}^d$, the sketch $\mathbf{S}_1 := \mathbf{0} \in \mathbb{F}^{n \times R}$, and the dual variable $\mathbf{y}_1 := \mathbf{0} \in \mathbb{R}^d$. These choices correspond to the simplest initial iterate $\mathbf{X}_1 := \mathbf{0}$.

6.1.2. Primal updates. At iteration $t = 1, 2, 3, \dots$, we compute a unit-norm vector \mathbf{v}_t that is an approximate minimum eigenvector of the gradient \mathbf{D}_t of the smoothed objective:

$$(6.1) \quad \xi_t := \mathbf{v}_t^* \mathbf{D}_t \mathbf{v}_t \leq \lambda_{\min}(\mathbf{D}_t) + \frac{\|\mathbf{D}_t\|}{\sqrt{t+1}} \quad \text{where} \quad \mathbf{D}_t := \mathbf{C} + \mathcal{A}^*(\mathbf{y}_t + \beta_t(\mathbf{z}_t - \mathbf{b})).$$

This calculation corresponds with [\(3.11\)](#). Form the primal update direction $\mathbf{h}_t = \mathcal{A}(\alpha \mathbf{v}_t \mathbf{v}_t^*)$, and then update the primal state variable \mathbf{z}_t and the sketch \mathbf{S}_t :

$$(6.2) \quad \mathbf{z}_{t+1} := (1 - \eta_t)\mathbf{z}_t + \eta_t \mathbf{h}_t \quad \text{and} \quad \mathbf{S}_{t+1} := (1 - \eta_t)\mathbf{S}_t + \eta_t \alpha \mathbf{v}_t (\mathbf{v}_t^* \mathbf{\Omega}).$$

We obtain the update rule for the primal state variable \mathbf{z}_t by applying the linear map \mathcal{A} to the primal update rule [\(3.6\)](#). The update rule for the sketch \mathbf{S}_t follows from [\(5.2\)](#).

6.1.3. Dual updates. The update to the dual variable takes the form

$$(6.3) \quad \mathbf{y}_{t+1} = \mathbf{y}_t + \gamma_t (\mathbf{z}_{t+1} - \mathbf{b})$$

where we choose the largest γ_t that satisfies the conditions

$$(6.4) \quad \gamma_t \|\mathbf{z}_{t+1} - \mathbf{b}\|^2 \leq \frac{4\alpha^2 \beta_0}{(t+1)^{3/2}} \|\mathcal{A}\|^2 \quad \text{and} \quad 0 \leq \gamma_t \leq \beta_0.$$

If needed, we set $\gamma_t = 0$ to prevent $\|\mathbf{y}_{t+1}\| > K$. This is the SketchyCGAL iteration.

6.2. Connection with CGAL. There is a tight connection between the iterates of SketchyCGAL and CGAL. Let $\mathbf{X}_1 := \mathbf{0}$. Using the vectors \mathbf{v}_t computed in [\(6.1\)](#), define matrices

$$(6.5) \quad \mathbf{H}_t := \alpha \mathbf{v}_t \mathbf{v}_t^* \quad \text{and} \quad \mathbf{X}_{t+1} := (1 - \eta_t)\mathbf{X}_t + \eta_t \mathbf{H}_t.$$

With these definitions, the following loop invariants are in force:

$$(6.6) \quad \mathbf{h}_t = \mathcal{A}\mathbf{H}_t \quad \text{and} \quad \mathbf{z}_t = \mathcal{A}\mathbf{X}_t \quad \text{and} \quad \mathbf{S}_t = \mathbf{X}_t \mathbf{\Omega}.$$

By comparing [subsections 3.4](#) and [6.1](#), we see that the trajectory $\{(\mathbf{X}_t, \mathbf{H}_t, \mathbf{y}_t) : t = 1, 2, 3, \dots\}$ could also have been generated by running the CGAL iteration. In other words, the variables in SketchyCGAL track the variables of some invocation of CGAL and inherit their behavior. We refer to the matrices \mathbf{X}_t as the *implicit* CGAL iterates.

Algorithm 6.1 SketchyCGAL for the model problem (2.2)

Input: Problem data for (2.2) implemented via the primitives (2.4), sketch size R , number T of iterations
Output: Rank- R approximate solution to (2.2) in factored form $\widehat{\mathbf{X}}_T = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^*$, where $\mathbf{U} \in \mathbb{F}^{n \times R}$ has orthonormal columns and $\mathbf{\Lambda} \in \mathbb{R}^{R \times R}$ is nonnegative diagonal

```

1 function SketchyCGAL( $R; T$ )
2   Scale problem data (subsection 7.1.1) ▷ [opt] Recommended!
3    $\beta_0 \leftarrow 1$  and  $K \leftarrow +\infty$  ▷ Default parameters
4   NystromSketch.Init( $n, R$ )
5    $\mathbf{z} \leftarrow \mathbf{0}_d$  and  $\mathbf{y} \leftarrow \mathbf{0}_d$ 
6   for  $t \leftarrow 1, 2, 3, \dots, T$  do
7      $\beta \leftarrow \beta_0 \sqrt{t+1}$  and  $\eta \leftarrow 2/(t+1)$ 
8      $[\xi, \mathbf{v}] \leftarrow \text{ApproxMinEvec}(\mathbf{C} + \mathcal{A}^*(\mathbf{y} + \beta(\mathbf{z} - \mathbf{b})); q_t)$  ▷ Algorithm 4.1 with  $q_t = t^{1/4} \log n$ 
▷ Implement with primitives (2.4) 1 2!
9      $\mathbf{z} \leftarrow (1 - \eta)\mathbf{z} + \eta \mathcal{A}(\alpha \mathbf{v} \mathbf{v}^*)$  ▷ Use primitive (2.4) 3
10     $\mathbf{y} \leftarrow \mathbf{y} + \gamma(\mathbf{z} - \mathbf{b})$  ▷  $\gamma$  is the largest solution to (6.4)
11    NystromSketch.RankOneUpdate( $\sqrt{\alpha} \mathbf{v}, \eta$ )
12     $[\mathbf{U}, \mathbf{\Lambda}] \leftarrow \text{NystromSketch.Reconstruct}()$ 
13     $\mathbf{\Lambda} \leftarrow \mathbf{\Lambda} + (\alpha - \text{tr}(\mathbf{\Lambda}))\mathbf{I}_{R/R}$  ▷ [opt] Enforce trace constraint in (2.2)

```

6.2.1. Approximating the CGAL iterates. We do not have access to the implicit CGAL iterates described above. Nevertheless, the sketch permits us to approximate them! After iteration t of SketchyCGAL, we can form a rank- R approximation $\widehat{\mathbf{X}}_t$ of the implicit iterate \mathbf{X}_t by invoking the formula (5.3) with $\mathbf{S} = \mathbf{S}_t$. According to Fact 5.2, for each $r < R - 1$,

$$(6.7) \quad \mathbb{E}_\Omega \|\mathbf{X}_t - \widehat{\mathbf{X}}_t\|_* \leq \left(1 + \frac{r}{R - r - 1}\right) \cdot \|\mathbf{X}_t - \llbracket \mathbf{X}_t \rrbracket_r\|_*.$$

In other words, the computed approximation $\widehat{\mathbf{X}}_t$ is a good proxy for the implicit iterate \mathbf{X}_t whenever the latter matrix is well-approximated by a low-rank matrix. The same bound (6.7) holds if we replace $\widehat{\mathbf{X}}_t$ by the truncated rank- r matrix $\llbracket \widehat{\mathbf{X}}_t \rrbracket_r$.

Remark 6.1 (Trace correction). The model problem (2.2) requires the matrix variable to have trace α , but the computed solution $\widehat{\mathbf{X}}_t$ rarely satisfies this constraint. Algorithm 6.1 includes an optional projection step (line 13) that corrects the trace. This step never increases the error in the Nyström approximation by more than a factor of two (subsection SM2.4). Our analysis of SketchyCGAL *does not* include the projection, but it is valuable in practice.

6.2.2. Assessing solution quality. Given quantities computed by SketchyCGAL, we can assess how well the implicit iterate \mathbf{X}_t solves the model problem (2.2). The infeasibility is just $\mathbf{z}_t - \mathbf{b}$. To bound suboptimality, we track the objective $p_t := \langle \mathbf{C}, \mathbf{X}_t \rangle$ by setting $p_1 = 0$ and using the update $p_{t+1} = (1 - \eta_t)p_t + \eta_t \alpha \mathbf{v}_t^*(\mathbf{C} \mathbf{v}_t)$. The surrogate duality gap (3.9) takes the form

$$g_t(\mathbf{X}_t) = p_t + \langle \mathbf{y}_t + \beta_t(\mathbf{z}_t - \mathbf{b}), \mathbf{z}_t \rangle - \lambda_{\min}(\mathbf{D}_t).$$

In practice, we use (6.1) to approximate the minimum eigenvalue: $\lambda_{\min}(\mathbf{D}_t) \approx \xi_t$. In light of (3.10), we can bound the suboptimality of \mathbf{X}_t via the expression

$$(6.8) \quad \langle \mathbf{C}, \mathbf{X}_t \rangle - \langle \mathbf{C}, \mathbf{X}_\star \rangle \leq g_t(\mathbf{X}_t) - \langle \mathbf{y}_t, \mathbf{z}_t - \mathbf{b} \rangle - \frac{1}{2} \beta_t \|\mathbf{z}_t - \mathbf{b}\|^2,$$

where \mathbf{X}_\star is a primal optimal point. See [subsection SM3.1](#) for more details.

6.3. Convergence of SketchyCGAL. The implicit iterates \mathbf{X}_t converge to a solution of the model problem (2.2) at the same rate as the iterates of CGAL would. On average, the rank- R iterates $\widehat{\mathbf{X}}_t$ track the implicit iterates. The discrepancy between them depends on how well the implicit iterates are approximated by low-rank matrices. Here is a simple convergence result that reflects this intuition; see [subsection SM3.2](#) for the easy proof and further results.

Theorem 6.2 (SketchyCGAL: Convergence). *Assume problem (2.2) satisfies strong duality, and let Ψ_\star be its solution set. For each $r < R - 1$, the iterates $\widehat{\mathbf{X}}_t$ computed by SketchyCGAL (subsections 6.1 and 6.2) satisfy*

$$\limsup_{t \rightarrow \infty} \mathbb{E}_\Omega \text{dist}_*(\widehat{\mathbf{X}}_t, \Psi_\star) \leq \left(1 + \frac{r}{R - r - 1}\right) \cdot \max_{\mathbf{X} \in \Psi_\star} \|\mathbf{X} - \llbracket \mathbf{X} \rrbracket_r\|_*.$$

The same bound holds if we replace $\widehat{\mathbf{X}}_t$ with the truncated approximation $\llbracket \widehat{\mathbf{X}}_t \rrbracket_r$. Here, dist_* is the nuclear-norm distance between a matrix and a set of matrices.

6.4. Resource usage. The main working variables in the SketchyCGAL iteration are the primal state $\mathbf{z}_t \in \mathbb{R}^d$, the computed eigenvector $\mathbf{v}_t \in \mathbb{R}^n$, the update direction $\mathbf{h}_t \in \mathbb{R}^d$, the test matrix $\Omega \in \mathbb{R}^{n \times R}$, the sketch $\mathbf{S}_t \in \mathbb{R}^{n \times R}$, and the dual variable $\mathbf{y}_t \in \mathbb{R}^d$. The total cost for storing these variables is $\Theta(d + Rn)$.

The arithmetic bottleneck in SketchyCGAL comes from the approximate eigenvector computation (6.1); see [subsection 4.2](#) for resource requirements. The remaining computation takes place in the variable updates. To form the primal update direction \mathbf{h}_t , we invoke the primitive (2.4)ⓐ. To update the primal state variable \mathbf{z}_t and the sketch \mathbf{S}_t in (6.2), we need $O(d + Rn)$ arithmetic operations. No further storage is required at this stage.

Table 1 documents the cost of performing T iterations of the SketchyCGAL method. The first column summarizes the resources consumed in the outer iteration. The second column tabulates the total resources spent to solve the eigenvalue problem (6.1) via the randomized Lanczos method ([Algorithm 4.1](#)).

In light of [Fact 3.1](#) and [subsection 6.2](#), we can be confident that the implicit iterate \mathbf{X}_T is ε -optimal within $T = O(\varepsilon^{-2})$ iterations. The formula (6.7) gives *a priori* guarantees on the quality of the approximation $\widehat{\mathbf{X}}_T$.

6.5. Theoretical performance of SketchyCGAL. We can package up this discussion in a theorem that describes the theoretical performance of the SketchyCGAL method.

Theorem 6.3 (SketchyCGAL). *Assume the model problem (2.2) satisfies strong duality. Fix a rank r , and set the sketch size $R \geq \zeta^{-1}r + 1$. After $T = O(\varepsilon^{-2})$ iterations, SketchyCGAL (subsections 6.1 and 6.2) returns a rank- r approximation $\widehat{\mathbf{X}} = \llbracket \widehat{\mathbf{X}}_T \rrbracket_r$ to an ε -optimal point of (2.2) that satisfies (2.3) with constant probability.*

Suppose we use the storage-optimal Lanczos method ([Algorithm 4.1](#)) to solve (6.1). Then the total storage cost for SketchyCGAL is $O(d + Rn)$ floats. The arithmetic requirements are $O(\varepsilon^{-2}(d + Rn) + \varepsilon^{-5/2}n \log(n/\varepsilon))$ flops, $O(\varepsilon^{-5/2} \log(n/\varepsilon))$ invocations of the primitives (2.4)ⓐⓑ, and $O(\varepsilon^{-2})$ applications of the primitive (2.4)ⓐ.

The constants depend on the problem data $(\mathbf{C}, \mathcal{A}, \mathbf{b}, \alpha)$ and algorithm parameters (β_0, K) .

Table 1

Resource usage for T iterations of SketchyCGAL with sketch size R . The algorithm returns a rank- R approximation to an ε -optimal solution (subsection 2.2) to the model problem (2.2). In theory, $T \sim \varepsilon^{-2}$. In practice, $T \sim \varepsilon^{-1}$. Constants and logarithms are omitted. See subsection 6.5.

	SketchyCGAL Algorithm 6.1	with Lanczos Algorithm 4.1
Storage (floats)	$d + Rn$	n
Arithmetic (flops)	$T(d + Rn)$	$T^{5/4}n$
Primitives (calls)		
(2.4) ①②	—	$T^{5/4}$
(2.4) ③	T	—

In practice, SketchyCGAL performs better than Theorem 6.3 suggests: empirically, we find that $T = O(\varepsilon^{-1})$ for real-world problem instances.

7. Numerical examples. This section showcases computational experiments that establish that SketchyCGAL is a practical method for solving large SDPs. We show that the algorithm is flexible by applying it to several classes of SDPs, and we show it is reliable by solving a large number of instances of each type. We give empirical evidence that the (implicit) iterates converge to optimality much faster than Theorem 6.3 suggests. Comparisons with other general-purpose SDP solvers demonstrate that SketchyCGAL is competitive for small SDPs, while it scales to problems that standard methods cannot handle.

7.1. Setup. All experiments are performed in MATLAB_R2018a with double-precision arithmetic. Source code is included with the supplementary material. To simulate the processing power of a personal laptop computer, we use a single Intel Xeon CPU E5-2630 v3, clocked at 2.40 GHz, with RAM usage capped at 16 GB.

Arithmetic costs are measured in terms of actual run time. MATLAB does not currently offer a memory profiler, so we externally monitor the total memory allocated. We approximate the storage cost by reporting the peak value minus the storage at MATLAB’s idle state.

7.1.1. Problem scaling. The bounds in the convergence theorem Fact 3.1 for the implicit iterates of SketchyCGAL depend on problem scaling. Our analysis motivates us to set

$$(7.1) \quad \|C\|_F = \|\mathcal{A}\| = \alpha = 1 \quad \text{and} \quad \|\mathbf{A}_1\|_F = \|\mathbf{A}_2\|_F = \cdots = \|\mathbf{A}_d\|_F.$$

In our experiments, we enforce the scalings (7.1), except where noted.

7.1.2. Implementation. Our experiments require a variant of SketchyCGAL that can handle inequality constraints; see section SM4. We implement the algorithm with the default parameters ($\beta_0 = 1$ and $K = +\infty$). The sketch uses a Gaussian test matrix. The eigenvalue subproblem is solved via randomized Lanczos (Algorithm 4.1). We include the optional trace normalization (line 13 in Algorithm 6.1) whenever appropriate. **No other tuning is done.**

7.2. The MaxCut SDP. We begin with MaxCut SDPs (1.3). Our goal is to assess the storage and arithmetic costs of SketchyCGAL for a standard testbed. We compare with provable solvers for general SDPs: SEDUMI [90], SDPT3 [92], MOSEK [70], and SDPNAL+ [100].

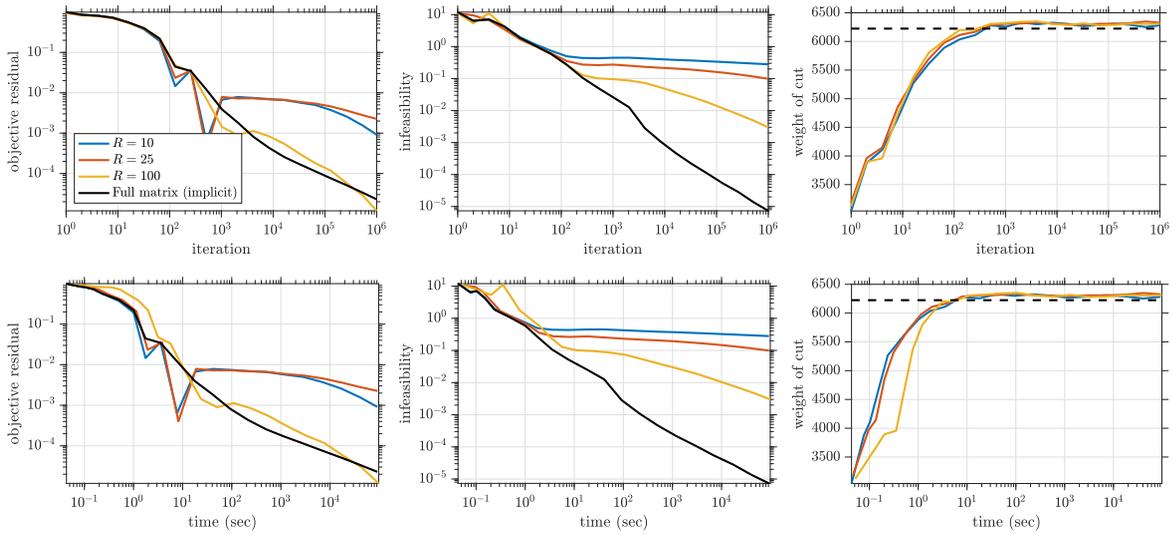


Figure 7.1. MaxCut SDP: Convergence. We solve the MaxCut SDP for the G67 dataset ($n = 10\,000$) with SketchyCGAL. The subplots show the suboptimality [left], infeasibility [center], and cut value [right] of the implicit iterates and low-rank approximations for sketch size $R \in \{10, 25, 100\}$ as a function of iteration [top] and run time [bottom]. The dashed line is the cut value of a high-accuracy SDPT3 solution. See subsection 7.2.4.

7.2.1. Rounding. To extract a cut from an approximate solution to (1.3), we apply a simple rounding procedure. SketchyCGAL returns a matrix $\widehat{\mathbf{X}} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^*$, where $\mathbf{U} \in \mathbb{R}^{n \times R}$ has orthonormal columns. The columns of the entrywise signum, $\text{sgn}(\mathbf{U})$, are the signed indicators of R cuts. We compute the weights of all R cuts and select the largest. The other solvers return a full-dimensional solution $\widehat{\mathbf{X}}$; we compute the top R eigenvectors of $\widehat{\mathbf{X}}$ using the MATLAB command `eigs` and invoke the same rounding procedure.

7.2.2. Datasets. We consider datasets from two different benchmark groups:

1. GSET: 67 binary-valued matrices generated by an autonomous random graph generator and published online [101]. The dimension n varies from 800 to 10 000.
2. DIMACS10: This benchmark consists of 150 symmetric matrices (with n varying from 39 to 50 912 018) chosen for the 10th DIMACS Implementation Challenge [12]. We consider 148 datasets with dimension $n \leq 24\,000\,000$. Two problems (`rgg_n_2_24_s0` and `Europe_osm`) exceeded the 16 GB storage limit.

7.2.3. Storage and arithmetic comparisons. We run each solver for each dataset and measure the storage cost and the runtime. We invoke SketchyCGAL with rank parameter $R = 10$, and we stop the algorithm when the error bound (6.8) guarantees that the implicit iterates have both relative suboptimality and infeasibility below 10^{-1} . For other solvers, we set the relative error tolerance to 10^{-1} . See section SM5 for implementation details.

Figure 1.1 displays the results of this experiment. The conventional convex solvers do not scale beyond $n = 10^4$ due to their high storage costs. In contrast, SketchyCGAL handles problems with $n \approx 2.2 \cdot 10^7$. For 9 datasets, SketchyCGAL did not reach the error tolerance within 5 days. Further details and results appear in subsection SM5.2.

7.2.4. Empirical convergence rates. Next, we investigate the empirical convergence of SketchyCGAL and the effect of the sketch size parameter R . We consider the MaxCut SDP (1.3) for the G67 dataset ($n = 10000$), the largest instance in GSET. We run 10^6 iterations of SketchyCGAL for each $R \in \{10, 25, 100\}$.

We use a high-accuracy solution from SDPT3 to approximate an optimal point \mathbf{X}_* of (2.2). Given a prospective solution \mathbf{X} , we compute its relative suboptimality and feasibility as

$$\text{objective residual} = \frac{|\langle \mathbf{C}, \mathbf{X} \rangle - \langle \mathbf{C}, \mathbf{X}_* \rangle|}{1 + |\langle \mathbf{C}, \mathbf{X}_* \rangle|} \quad \text{and} \quad \text{infeasibility} = \frac{\|\mathcal{A}\mathbf{X} - \mathbf{b}\|}{1 + \|\mathbf{b}\|}.$$

It is standard to increment the denominator by one to handle small values gracefully. These quantities are evaluated with respect to the original (not rescaled) problem data.

Figure 7.1 illustrates the convergence of SketchyCGAL for this problem. After t iterations, the residual and infeasibility decay like $O(t^{-1})$, which is far better than the $O(t^{-1/2})$ bound in Theorem 6.3. Similar behavior is manifest for other datasets and other problems.

Figure 7.1 also displays the weight of the cut obtained after rounding, compared with the weight of the cut obtained from the SDPT3 solution. Observe that sketch size $R = 10$ is sufficient, and the SketchyCGAL solutions yield excellent cuts after a few hundred iterations.

7.2.5. Primal–dual convergence. We have observed empirically that convergence occurs for the implicit primal sequence (\mathbf{X}_t) , the dual sequence (\mathbf{y}_t) , and the posterior error bound (6.8) generated by SketchyCGAL. See subsection SM5.3 for numerical evidence.

7.2.6. Hard MaxCut instances. Waldspurger & Waters [97] construct instances of the MaxCut SDP (1.3) that are challenging for algorithms based on the Burer–Monteiro [24] factorization heuristic (subsection 8.4). Each instance has a unique solution and the solution has rank 1, but Burer–Monteiro methods require factorization rank $R = \Theta(\sqrt{n})$, resulting in storage cost $\Theta(n^{3/2})$. In subsection SM5.4, we give numerical evidence that SketchyCGAL can solve these instances with sketch size $R = 2$, achieving the optimal storage $\Theta(n)$. We confirm that Burer–Monteiro usually fails in the optimal-storage regime.

7.3. Abstract phase retrieval. Phase retrieval is the problem of reconstructing a complex-valued signal from intensity-only measurements. It arises in interferometry [38], speech processing [15], array imaging [30], microscopy [48], and many other applications. We will outline a standard method [30, 48] for performing phase retrieval by means of an SDP.

This section uses synthetic instances of a phase retrieval SDP to compare the scaling behavior of SketchyCGAL and CGAL. We also consider a third algorithm ThinCGAL, inspired by [104], that maintains a thin SVD of the matrix variable via rank-one updates [22].

7.3.1. Phase retrieval SDPs. Let $\mathbf{x}_\natural \in \mathbb{C}^n$ be an unknown (discrete) signal. For known vectors $\mathbf{a}_i \in \mathbb{C}^n$, we acquire measurements of the form

$$(7.2) \quad b_i = |\langle \mathbf{a}_i, \mathbf{x}_\natural \rangle|^2 \quad \text{for } i = 1, 2, 3, \dots, d.$$

Abstract phase retrieval is the challenging problem of recovering \mathbf{x}_\natural from \mathbf{b} .

Let us summarize a lifting approach introduced by Balan et al. [14]. The key idea is to replace the signal vector \mathbf{x}_\natural by the matrix $\mathbf{X}_\natural = \mathbf{x}_\natural \mathbf{x}_\natural^*$. Then rewrite (7.2) as

$$(7.3) \quad b_i = \mathbf{a}_i^* \mathbf{X}_\natural \mathbf{a}_i = \langle \mathbf{A}_i, \mathbf{X}_\natural \rangle \quad \text{where} \quad \mathbf{A}_i = \mathbf{a}_i \mathbf{a}_i^* \quad \text{for } i = 1, 2, 3, \dots, d.$$

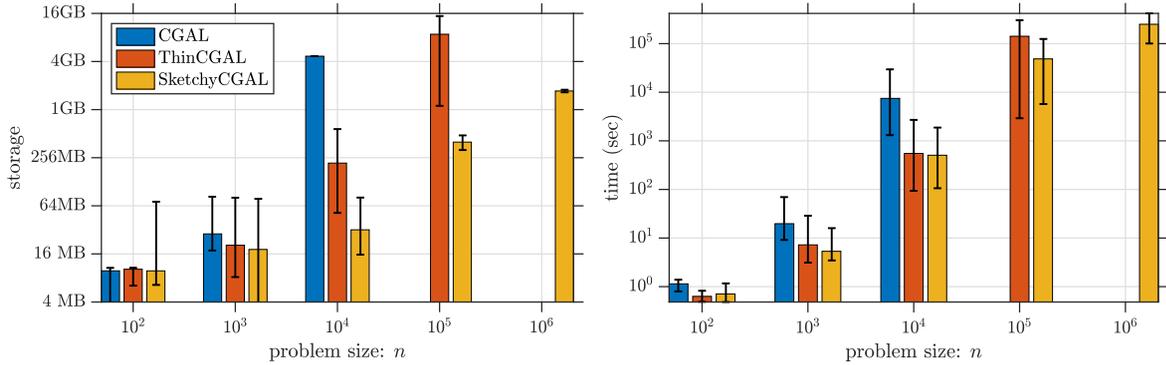


Figure 7.2. Phase retrieval SDP: Scalability. Storage cost [left] and runtime [right] to solve random instances with algorithms *CGAL*, *ThinCGAL*, and *SketchyCGAL*. The height of each bar is the mean; the interval marks the minimum and maximum over 20 trials. Missing bars indicate total failure. See [subsection 7.3.4](#).

Promoting the implicit constraints on \mathbf{X}_\dagger and forming the \mathbf{A}_i into a linear map \mathcal{A} , we can express the problem of finding \mathbf{X}_\dagger as a feasibility problem with a matrix variable:

$$\text{find } \mathbf{X} \in \mathbb{S}_n \text{ subject to } \mathcal{A}\mathbf{X} = \mathbf{b}, \quad \mathbf{X} \text{ is psd}, \quad \text{rank}(\mathbf{X}) = 1.$$

To reach a tractable convex formulation, we pass to a trace minimization SDP [37]:

$$(7.4) \quad \text{minimize } \text{tr } \mathbf{X} \quad \text{subject to } \mathcal{A}\mathbf{X} = \mathbf{b}, \quad \mathbf{X} \text{ is psd}, \quad \text{tr } \mathbf{X} \leq \alpha.$$

The parameter α is an upper bound on the signal energy $\|\chi_\dagger\|^2$, which can be estimated from the observed data \mathbf{b} ; see [104] for the details. We can solve the SDP (7.4) via *SketchyCGAL*.

7.3.2. Rounding. Suppose that we have obtained an approximate solution \mathbf{X} to (7.4). To estimate the signal χ_\dagger , we form the vector $\chi = \sqrt{\lambda}\mathbf{u}$ where (λ, \mathbf{u}) is a maximum eigenpair of \mathbf{X} . Both *SketchyCGAL* and *ThinCGAL* return eigenvalue decompositions, so this step is trivial. For *CGAL*, we use the MATLAB function `eigs` to perform this computation.

7.3.3. Dataset and evaluation. We consider synthetic phase retrieval instances. For each $n \in \{10^2, 10^3, \dots, 10^6\}$, we generate 20 independent datasets as follows. First, draw $\chi_\dagger \in \mathbb{C}^n$ from the complex standard normal distribution. Then acquire $d = 12n$ phaseless measurements (7.2) using the coded diffraction pattern model [28]; see [subsection SM6.1](#). The induced linear maps \mathcal{A} and \mathcal{A}^* can be applied via the fast Fourier transform (FFT).

The relative error in a signal reconstruction χ is given by $\min_{\phi \in \mathbb{R}} \|e^{i\phi}\chi - \chi_\dagger\| / \|\chi_\dagger\|$. In the SDP (7.4), we set $\alpha = 3n$ to demonstrate insensitivity of the algorithm to the choice of α .

7.3.4. Storage and arithmetic comparisons. For each algorithm, we report the storage cost and runtime required to produce a signal estimate χ with (exact) relative error below 10^{-2} . We invoke *SketchyCGAL* with sketch size parameter $R = 5$.

[Figure 7.2](#) displays the outcome. We witness the benefit of sketching for both storage and arithmetic. *CGAL* fails for all large-scale instances ($n = 10^5$ and 10^6) due to storage allocation. For the same reason, *ThinCGAL* fails for $n = 10^6$. *SketchyCGAL* successfully solves all problem instances to the target accuracy.

7.4. Phase retrieval in microscopy. Next, we study a more realistic phase retrieval problem that arises from a type of microscopy system [110] called Fourier ptychography (FP). Phase retrieval SDPs offer a potential approach to FP imaging [48]. This section shows that SketchyCGAL can successfully solve the difficult phase retrieval SDPs that arise from FP.

7.4.1. Fourier ptychography. FP microscopes circumvent the physical limits of a simple lens to achieve high-resolution and wide field-of-view simultaneously [110]. To do so, an FP microscope illuminates a sample from many angles and uses a simple lens to collect low-resolution intensity-only images. The measurements are low-pass filters, whose transfer functions depend on the lens and the angle of illumination [48]. From the data, we form a high-resolution image by solving a phase retrieval problem; e.g., via the SDP (7.4).

The high-resolution image of the sample is represented by a Fourier-domain vector $\chi_{\dagger} \in \mathbb{C}^n$. We acquire d intensity-only measurements of the form (7.2), where d is the total number of pixels in the low-resolution illuminations. The low-pass measurements are encoded in vectors \mathbf{a}_i . The operators \mathcal{A} and \mathcal{A}^* , built from the matrices $\mathbf{A}_i = \mathbf{a}_i \mathbf{a}_i^*$, can be applied via the FFT.

7.4.2. Dataset and evaluation. The authors of [48] provided transmission matrices \mathbf{A}_i of a working FP system. We simulate this system in the computer environment to acquire noiseless intensity-only measurements of a high-resolution target image [33]. In this setup, $\chi_{\dagger} \in \mathbb{C}^n$ corresponds to the Fourier transform of an $n = 320^2 = 102\,400$ pixel grayscale image. We normalize χ_{\dagger} so that $\|\chi_{\dagger}\| = 1$. We acquire 225 low-resolution illuminations of the original image, each with $64^2 = 4\,096$ pixels. The total number of measurements is $d = 921\,600$.

We evaluate the error of an estimate χ as in subsection 7.3.3. Although we know that the signal energy equals 1, it is more realistic to approximate the signal energy by setting $\alpha = 1.5$ in the SDP (7.4).

7.4.3. FP imaging. We solve the phase retrieval SDP (7.4) by performing 10 000 iterations of SketchyCGAL with rank parameter $R = 5$. The top eigenvector of the output gives an approximation $\chi \in \mathbb{C}^n$ of the signal. The inverse Fourier transform of χ is the desired image.

Figure 7.3 displays the image reconstruction after $t \in \{10, 10^2, 10^3, 10^4\}$ iterations. We obtain a good-quality result in 5 hours after 1 000 iterations; a sharper image emerges in 59 hours after 10 000 iterations are complete. We believe the computational time can be reduced substantially with a parallel or GPU implementation. Nevertheless, it is gratifying that we have solved a difficult SDP whose matrix variable has over 10^{10} entries. See subsection SM6.2 for a larger FP instance with $n = 640^2$ pixels.

7.5. The quadratic assignment problem. The quadratic assignment problem (QAP) is a very difficult combinatorial optimization problem that includes the traveling salesman, max-clique, bandwidth problems, and many others as special cases [66]. SDP relaxations offer a powerful approach for obtaining good solutions to large QAP problems [108]. In this section, we demonstrate that SketchyCGAL can solve these challenging SDPs.

7.5.1. QAP. We begin with the simplest form of the QAP. Fix symmetric $\mathbf{n} \times \mathbf{n}$ matrices $\mathbf{A}, \mathbf{B} \in \mathbb{S}_{\mathbf{n}}$ where \mathbf{n} is a natural number. We wish to “align” the matrices by solving

$$(7.5) \quad \text{minimize} \quad \text{tr}(\mathbf{A}\mathbf{\Pi}\mathbf{B}\mathbf{\Pi}^*) \quad \text{subject to} \quad \mathbf{\Pi} \text{ is an } \mathbf{n} \times \mathbf{n} \text{ permutation matrix.}$$

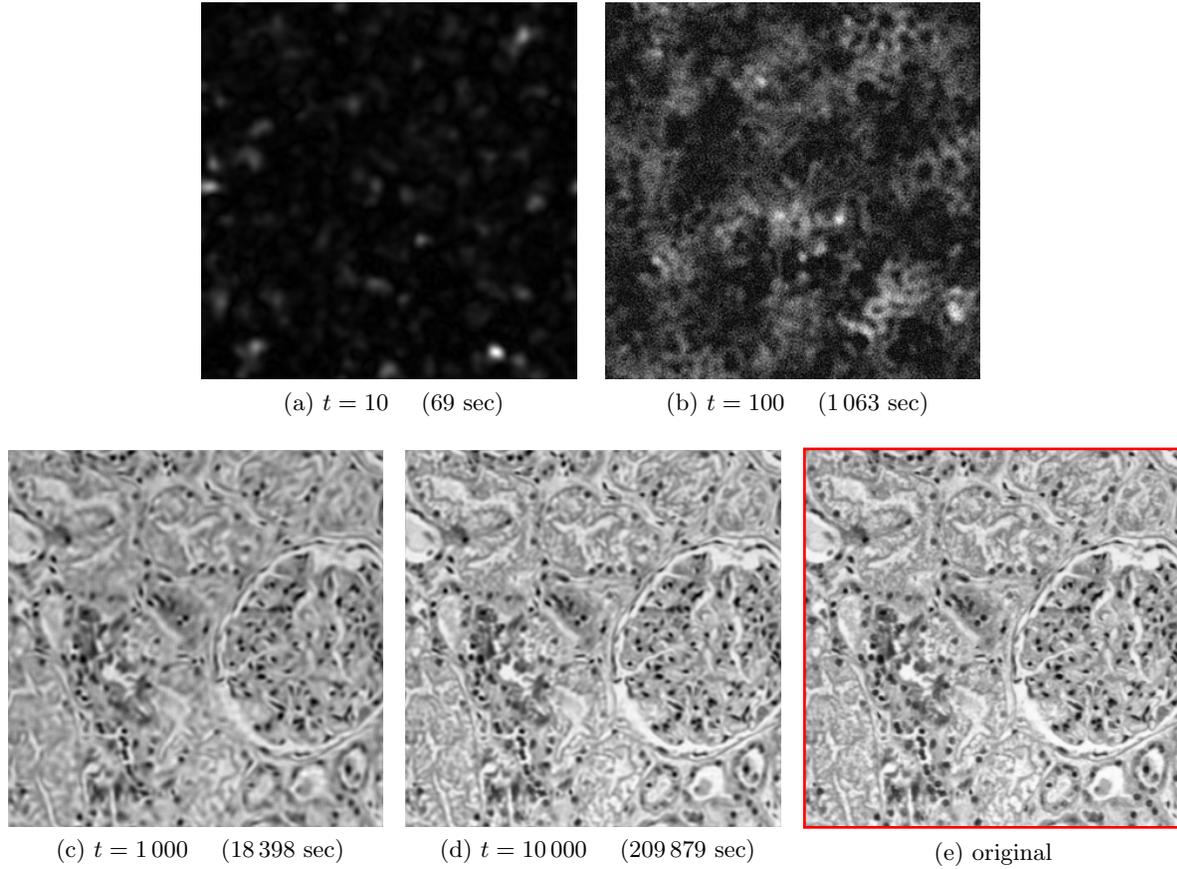


Figure 7.3. Phase retrieval SDP: Imaging. *Reconstruction of an $n = 320^2$ pixel image from Fourier ptychography data. We solve an $n \times n$ phase retrieval SDP via SketchyCGAL with rank parameter $R = 5$ and show the images obtained at iterations $t = 10^3, 10^4$. The last subfigure is the original. See [subsection 7.4.3](#).*

Recall that a permutation matrix $\mathbf{\Pi}$ has precisely one nonzero entry in each row and column, and that nonzero entry equals one. A brute force search over the $n!$ permutation matrices of size n quickly becomes intractable as n grows. Unsurprisingly, the QAP problem (7.5) is NP-hard [86]. Instances with $n > 30$ usually cannot be solved in reasonable time.

7.5.2. Relaxations. There is an extensive literature on SDP relaxations for QAP, beginning with the work [108] of Zhao et al. We consider a weaker relaxation inspired by [49, 23]:

$$\begin{aligned}
 & \text{minimize} && \text{tr}[(\mathbf{B} \otimes \mathbf{A})\mathbf{Y}] \\
 & \text{subject to} && \text{tr}_1(\mathbf{Y}) = \mathbf{I}, \quad \text{tr}_2(\mathbf{Y}) = \mathbf{I}, \quad \mathcal{G}(\mathbf{Y}) \geq 0, \\
 (7.6) \quad & && \text{vec}(\mathbf{P}) = \text{diag}(\mathbf{Y}), \quad \mathbf{P}\mathbf{1} = \mathbf{1}, \quad \mathbf{1}^*\mathbf{P} = \mathbf{1}^*, \quad \mathbf{P} \geq \mathbf{0}, \\
 & && \begin{bmatrix} 1 & \text{vec}(\mathbf{P})^* \\ \text{vec}(\mathbf{P}) & \mathbf{Y} \end{bmatrix} \succcurlyeq \mathbf{0}, \quad \text{tr} \mathbf{Y} = n.
 \end{aligned}$$

We have written \otimes for the Kronecker product.

The constraint $\mathcal{G}(\mathbf{Y}) \geq 0$ enforces nonnegativity of a subset of the entries in \mathbf{Y} . In the Zhao et al. relaxation, \mathcal{G} is the identity map, so it yields $O(\mathbf{n}^4)$ constraints. We reduce the complexity by choosing \mathcal{G} more carefully. In our formulation, \mathcal{G} extracts precisely the nonzero entries of the matrix $\mathbf{B} \otimes \mathbf{1}\mathbf{1}^*$. This is beneficial because \mathbf{B} is sparse in many applications. For example, in traveling salesman and bandwidth problems, \mathbf{B} has $O(\mathbf{n})$ nonzero entries, so the map \mathcal{G} produces only $O(\mathbf{n}^3)$ constraints.

The main variable \mathbf{Y} in (7.6) has dimension $\mathbf{n}^2 \times \mathbf{n}^2$. As a consequence, the problem has $O(\mathbf{n}^4)$ degrees of freedom, together with $O(\mathbf{n}^3)$ to $O(\mathbf{n}^4)$ constraints (depending on \mathcal{G}). The explosive growth of this relaxation scuttles most algorithms by the time $\mathbf{n} > 50$. To solve larger instances, many researchers resort to even weaker relaxations.

In contrast, we can solve the relaxation (7.6) directly using SketchyCGAL, up to $\mathbf{n} = 150$. By limiting the number of inequality constraints, via the operator \mathcal{G} , we achieve substantial reductions in resource usage. We validate our algorithm on QAPs where the exact solution is known, and we compare the performance with algorithms [107, 23] for other relaxations.

7.5.3. Rounding. Given an approximate solution to (7.6), we use a rounding method to construct a permutation. SketchyCGAL returns a matrix $\widehat{\mathbf{X}} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^*$ where \mathbf{U} has dimension $(\mathbf{n}^2 + 1) \times R$. We extract the first column of \mathbf{U} , discard its first entry and reshape the remaining part into an $\mathbf{n} \times \mathbf{n}$ matrix. Then we project this matrix onto the set of permutation matrices via the Hungarian method [58, 71, 52]. This yields a feasible point $\mathbf{\Pi}$ for the problem (7.5). We repeat this procedure for all R columns of \mathbf{U} , and we pick the one that minimizes $\text{tr}(\mathbf{A}\mathbf{\Pi}\mathbf{B}\mathbf{\Pi}^*)$. This permutation gives an upper bound on the optimal value of (7.5).

7.5.4. Datasets and evaluation. We consider instances from QAPLIB [27] and TSPLIB [83] that are used in [23]. The optimal values are known, and the permutation size \mathbf{n} varies between 12 and 150. (Recall that the SDP matrix dimension $n = \mathbf{n}^2 + 1$.) We report

$$(7.7) \quad \text{relative gap \%} = \frac{\text{upper bound obtained} - \text{optimum}}{\text{optimum}} \times 100$$

7.5.5. Solving QAPs. To solve (7.6), we cannot use the scaling (7.1) because $\|\mathcal{A}\|$ is not available; see the source code for our approach. We apply SketchyCGAL with sketch size $R = \mathbf{n}$, so the sketch uses storage $\Theta(\mathbf{n}^3)$. After rounding, a low- or medium-accuracy solution of (7.6) often provides a better permutation than a high-accuracy solution. Therefore, we applied the rounding step at iterations 2, 4, 8, 16, \dots and tracked the quality of the best permutation attained on the solution path. We stopped after (the first of) 10^6 iterations or 72 hours.

The results of this experiment appear in Figure 7.4. We compare against the best value reported by Bravo Ferreira et al. in [23, Tables 4 and 6] for their CSDP method with clique size $k \in \{2, 3, 4\}$. We also include the results that [23] lists for the PATH method [107].

SketchyCGAL allows us to solve a tighter SDP relaxation of QAP than the other methods (CSDP, PATH). As a consequence, we obtain significantly smaller gaps for most instances.

8. Related work. There is a large body of literature on numerical methods for solving SDPs; see [67] for a recent survey. Here we describe the SDP literature just enough to contextualize our contributions.

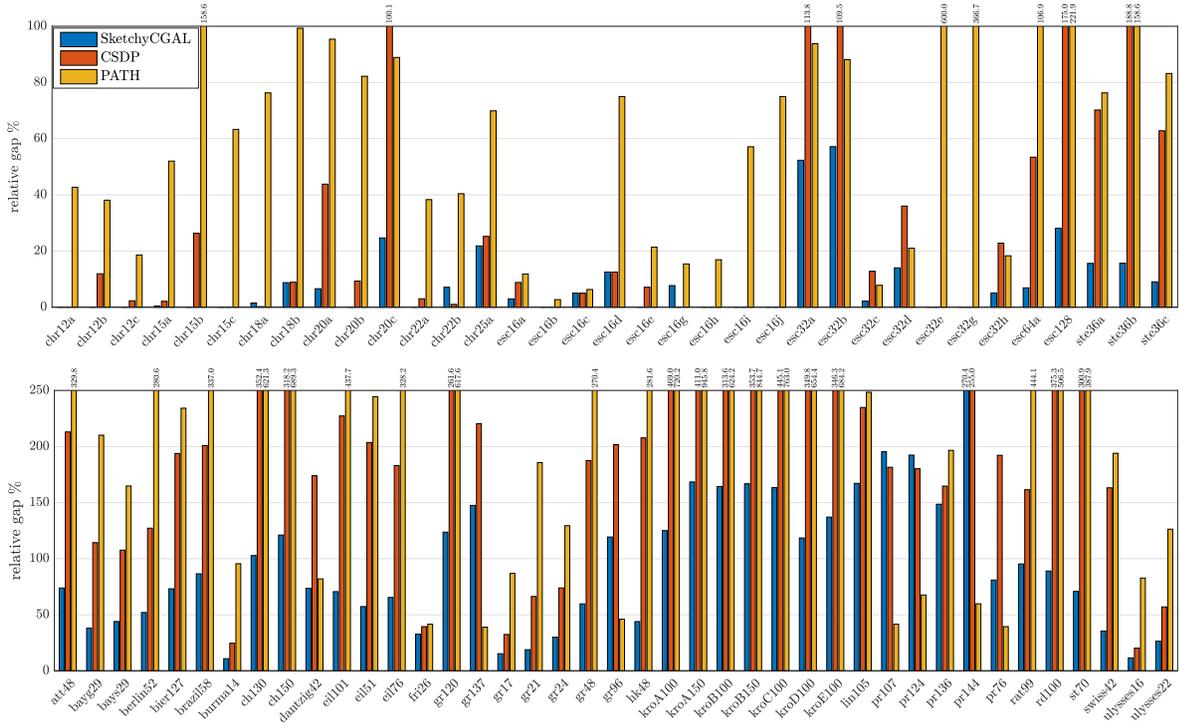


Figure 7.4. QAP relaxation: Solution quality. Using *SketchyCGAL*, the *CSDP* method [23], and the *PATH* method [107], we solve SDP relaxations of QAP instances from *QAPLIB* [top] and *TSPLIB* [bottom]. The bars compare the cost of the computed solution, against the optimal value; shorter is better. See subsection 7.5.5.

8.1. Standard methodologies. First, we outline the approaches that drive most of the reliable, general-purpose SDP software packages that are currently available.

Interior-point methods (IPMs) [75, 76, 3, 4, 53, 54] reformulate the SDP as an unconstrained problem and take an (approximate) Newton step at each iteration. In exchange, they deliver quadratic convergence. Hence IPMs are widely used to solve SDPs to high precision [56, 9, 73, 17]. Software packages include *SeDuMi* [90], *MoSeK* [70], and *SDPT3* [92]. Alas, IPMs do not scale to large problems: to solve the Newton system we must store and factor large, dense matrices. A typical IPM for the SDP (2.2) requires $\Theta(n^3 + d^2n^2 + d^3)$ arithmetic operations per iteration and $\Theta(n^2 + dn + d^2)$ memory [6].

Several effective SDP solvers are based on the augmented Lagrangian (AL) paradigm [26, 98, 109, 99]. In particular, Zhao et al. [109] employ a semi-smooth Newton method and the conjugate gradient method to solve the subproblems. Their method is enhanced and implemented in the software package *SDPNAL+* [100]. AL methods for SDPs typically require storage $\Omega(n^2)$.

8.2. First-order methods. First-order methods use only gradient information to solve the SDP to reduce runtime and storage requirements. We focus on *projection-free* algorithms, suitable large SDPs, that do not require full SVD computations. Major first-order methods for SDP include approaches based on the conditional gradient method (CGM) [39, 63, 51]

and extensions that handle more complex constraints [32, 46, 50], primal–dual subgradient algorithms [74, 105], and the matrix multiplicative weight (MMW) method [96, 10], or equivalently, the mirror-prox algorithm with the quantum entropy mirror map [72].

The standard CGM algorithm does not apply to the model problem (2.2) because of the affine constraint $\mathcal{A}\mathbf{X} = \mathbf{b}$. Several variants [42, 65, 87, 103] of CGM can handle affine constraints. In particular, CGAL [102] does so by applying CGM to an AL formulation. We have chosen to extend CGAL because of its strong empirical performance and its robustness to inexact eigenvector computations; see [102, Sec. 4–5].

Primal–dual subgradient methods perform subgradient ascent on the dual problem; the cost of each iteration is dominated by an eigenvector computation. Nesterov [74] constructs primal iterates by proximal mapping. The algorithm in Yurtsever et al. [105] constructs primal iterates by an averaging technique, similar to the updates in CGM; this method involves a line search, so it requires very accurate eigenvector calculations.

The MMW method is derived by reducing the SDP to a sequence of feasibility problems. These are reformulated as a primal–dual game whose dual is an eigenvalue optimization problem. The resulting MMW algorithm can be interpreted as performing gradient descent in a dual space and using the matrix exponential map to transfer information back to the primal space. To scale this approach to larger problems, researchers have proposed linearization, random projection, sparsification techniques, and stochastic Lanczos quadrature to approximate the matrix exponential [11, 7, 80, 40, 41, 8, 13, 61, 29]. Even so, the reduction to a sequence of feasibility problems makes this technique impractical for general SDPs. We are aware of only one computational evaluation of the MMW idea [13].

8.3. Storage considerations. Almost all provably correct SDP algorithms store and operate on a full-dimensional matrix variable, so they are not suitable for very large SDPs.

Some primal–dual subgradient methods and CGM variants build an approximate solution as a convex combination of rank-one updates, so the rank of the solution does not exceed the number of iterations. This fact has led researchers to call these methods “storage-efficient,” but this claim is misleading because the algorithms require many iterations to converge.

In the conference paper [106], written by a subset of the authors, we observed that certain types of optimization algorithms can be combined with sketching to control storage costs. As a first example, we augmented CGM with sketching to obtain a new algorithm called *SketchyCGM*. This method solves a special class of low-rank matrix optimization problems that arise in statistics and machine learning. We believe that *SketchyCGM* is the first algorithm for this class of problems that provably succeeds with optimal storage.

To develop *SketchyCGAL*, we changed the base optimization algorithm (to CGAL) so that we can solve standard-form SDPs. We switched to a simpler sketching technique (the Nyström sketch) that has better empirical performance. We also analyzed how accurately to solve the eigenvalue problems to ensure that *SketchyCGAL* succeeds (section SM1), and we deployed an approximate eigenvalue computation method (randomized Lanczos) that meets our needs. Altogether, this effort leads to a storage-optimal algorithm that works for all standard-form SDPs with a minimum of tuning. Reducing the storage has the ancillary benefit of reducing arithmetic and communication costs, which also improves scalability.

In concurrent work with Ding [36], a subset of the authors developed a new *approximate*

mate complementarity principle that also yields a storage-optimal algorithm for standard-form SDPs. This approach uses a suboptimal dual point to approximate the range of the primal solution to the SDP. By compressing the primal problem to this subspace, we can solve the primal SDP with limited storage. This method, however, has more limited guarantees than SketchyCGAL. A numerical evaluation is in progress.

8.4. Nonconvex Burer–Monteiro methods. The most famous approach to low-storage semidefinite programming is the factorization heuristic proposed by Homer and Peinado [47] and refined by Burer and Monteiro (BM) [24]. The main idea is to reformulate the model problem (2.2) by expressing the psd matrix variable $\mathbf{X} = \mathbf{F}\mathbf{F}^*$ in terms of a factor $\mathbf{F} \in \mathbb{R}^{n \times R}$, where the rank parameter $R \ll n$. That is,

$$(8.1) \quad \text{minimize } \langle \mathbf{C}, \mathbf{F}\mathbf{F}^* \rangle \quad \text{subject to } \mathcal{A}(\mathbf{F}\mathbf{F}^*) = \mathbf{b}, \quad \mathbf{F}\mathbf{F}^* \in \alpha \Delta_n.$$

This approach controls storage by sacrificing convexity and the associated guarantees.

Many nonlinear programming methods have been applied to optimize (8.1). AL methods are commonly used [24, 25, 59, 85, 88]. The most popular research software based on BM factorization is Manopt [20], which implements manifold optimization algorithms including Riemannian gradient and Riemannian trust region methods [1]. Consequently, Manopt is limited to problems where the factorized formulation (8.1) defines a smooth manifold.

There has been an intense effort to establish theoretical results for the BM factorization approach. It is clear that every solution to the SDP (2.2) of rank R or less is also a solution to the factorized problem (8.1). On the other hand, (8.1) may admit local minima that are not global minima of (2.2). Some guarantees are available. For example, if \mathbf{C} is generic and the constraint set of (8.1) is a smooth manifold and $R \geq \sqrt{2(d+1)}$, then each second-order critical point of (8.1) is a global optimum [21]. A second-order critical point can be located using a Riemannian trust region method. See [19, 81, 31] for additional theoretical analysis.

The storage and arithmetic costs of solving (8.1) depend on the factorization rank R . Unfortunately, the BM method may fail when $R = o(\sqrt{d})$. Below this threshold, the BM formulation (8.1) can have spurious solutions (second-order critical points that are not globally optimal), and the bad problem instances can form a set of positive measure [97]. Hence the Burer–Monteiro approach cannot support provably correct algorithms with storage costs better than $\Omega(n\sqrt{d})$. See subsection SM5.4 for numerical evidence.

9. Conclusion. We have presented a practical, new approach for solving SDPs at scale. Our algorithm, SketchyCGAL, combines a primal–dual optimization method with randomized linear algebra techniques to achieve unprecedented guarantees when the problem is weakly constrained and the solution is approximately low rank. We hope that our ideas lead to further algorithmic advances and support new applications of semidefinite programming.

SketchyCGAL is currently limited by the arithmetic cost of solving large eigenvalue problems to increasing accuracy. It also falters for SDPs with a large number of constraints because it depends on a primal–dual approach. Moreover, our analysis does not fully explain the observed behavior of the algorithm, including the rate of convergence of the primal variable or the convergence of the dual variable and the surrogate duality gap. These topics merit further research.

Acknowledgments. Nicolas Boumal encouraged us to reconsider the storage-optimal randomized Lanczos method ([Algorithm 4.1](#)); this algorithm simplifies the presentation while improving our theoretical and numerical results. Richard Kueng allowed us to include his results on trace normalization (Step 13 in [Algorithm 6.1](#)). Irène Waldspurger provided code that generates instances of the MaxCut problem that are difficult for the Burer–Monteiro heuristic. We would also like to thank the editor and reviewers for their feedback.

REFERENCES

- [1] P.-A. ABSIL, C. BAKER, AND K. GALLIVAN, *Trust-region methods on Riemannian manifolds*, *Found. Comput. Math.*, 7 (2007), pp. 303–330.
- [2] A. Y. ALFAKIH, A. KHANDANI, AND H. WOLKOWICZ, *Solving Euclidean distance matrix completion problems via semidefinite programming*, *Comput. Optim. Appl.*, 12 (1999), pp. 13–30.
- [3] F. ALIZADEH, *Combinatorial optimization with interior point methods and semi-definite matrices*, PhD thesis, Univ. Minnesota, 1991.
- [4] F. ALIZADEH, *Interior point methods in semidefinite programming with applications to combinatorial optimization*, *SIAM J. Optim.*, 5 (1995), pp. 13–51.
- [5] F. ALIZADEH, J.-P. A. HAEBERLY, AND M. L. OVERTON, *Complementarity and nondegeneracy in semidefinite programming*, *Math. Programming*, 77 (1997), pp. 111–128.
- [6] F. ALIZADEH, J.-P. A. HAEBERLY, AND M. L. OVERTON, *Primal-dual interior point methods for semidefinite programming: Convergence rates, stability and numerical results*, *SIAM J. Optim.*, 8 (1998), pp. 746–768.
- [7] Z. ALLEN-ZHU, Y. T. LEE, AND L. ORECCHIA, *Using optimization to obtain a width-independent, parallel, simpler, and faster positive SDP solver*, 2015, <https://arxiv.org/abs/1507.02259>.
- [8] Z. ALLEN-ZHU AND Y. LI, *Follow the compressed leader: Faster online learning of eigenvectors and faster MMWU*, June 2017, <https://arxiv.org/abs/1701.01722>.
- [9] M. F. ANJOS AND J. B. LASSERRE, *Handbook on Semidefinite, Conic and Polynomial Optimization*, Springer USA, 2012.
- [10] S. ARORA, E. HAZAN, AND S. KALE, *Fast algorithms for approximate semidefinite programming using the multiplicative weights update method*, in *Proc. 46th Ann. IEEE Symp. Foundations of Computer Science*, FOCS '05, Washington, DC, USA, 2005, pp. 339–348.
- [11] S. ARORA AND S. KALE, *A combinatorial, primal-dual approach to semidefinite programs*, *J. ACM*, 63 (2016), pp. 12:1–12:35.
- [12] D. BADER, H. MEYERHENKE, P. SANDERS, AND D. WAGNER, *Graph partitioning and graph clustering. 10th DIMACS Implementation Challenge Workshop*, February 2012, <https://www.cise.ufl.edu/research/sparse/matrices/DIMACS10/index.html> (accessed October 2019).
- [13] M. BAES, M. BÜRGISSER, AND A. NEMIROVSKI, *A randomized mirror-prox method for solving structured large-scale matrix saddle-point problems*, *SIAM J. Optim.*, 23 (2013), pp. 934–962.
- [14] R. BALAN, B. G. BODMANN, P. G. CASAZZA, AND D. EDIDIN, *Painless reconstruction from magnitudes of frame coefficients*, *J. Fourier Anal. Appl.*, 15 (2009), pp. 488–501.
- [15] R. BALAN, P. CASAZZA, AND D. EDIDIN, *On signal reconstruction without phase*, *Appl. Comp. Harmonic Anal.*, 20 (2006), pp. 345–356.
- [16] A. BARVINOK, *A course in convexity*, AMS, Providence, RI, 2002.
- [17] A. BEN-TAL AND A. NEMIROVSKI, *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*, SIAM, Philadelphia, PA, 2001.
- [18] D. P. BERTSEKAS, *Constrained optimization and Lagrange multiplier methods*, *Computer Science and Applied Mathematics*, Academic Press, New York, NY, 1982.
- [19] S. BHOJANAPALLI, N. BOUMAL, P. JAIN, AND P. NETRAPALLI, *Smoothed analysis for low-rank solutions to semidefinite programs in quadratic penalty form*, in *Proc. 31st Conf. Learning Theory*, vol. 75, 2018, pp. 3243–3270.
- [20] N. BOUMAL, B. MISHRA, P.-A. ABSIL, AND R. SEPULCHRE, *Manopt, a Matlab toolbox for optimization on manifolds*, *J. Mach. Learn. Res.*, 15 (2014), pp. 1455–1459.

- [21] N. BOUMAL, V. VORONINSKI, AND A. BANDEIRA, *The non-convex Burer–Monteiro approach works on smooth semidefinite programs*, in Adv. Neural Information Processing Systems 29, 2016, pp. 2757–2765.
- [22] M. BRAND, *Fast low-rank modifications of the thin singular value decomposition*, Linear Alg. Appl., 415 (2006), pp. 20 – 30.
- [23] J. F. S. BRAVO FERREIRA, Y. KHOO, AND A. SINGER, *Semidefinite programming approach for the quadratic assignment problem with a sparse graph*, Comput. Optim. Appl., 69 (2018), pp. 677–712.
- [24] S. BURER AND R. D. MONTEIRO, *A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization*, Math. Prog., 95 (2003), pp. 329–357.
- [25] S. BURER AND R. D. MONTEIRO, *Local minima and convergence in low-rank semidefinite programming*, Math. Prog., 103 (2005), pp. 427–444.
- [26] S. BURER AND D. VANDENBUSSCHE, *Solving lift-and-project relaxations of binary integer programs*, SIAM J. Optim., 16 (2006), pp. 726–750.
- [27] R. E. BURKARD, S. E. KARISCH, AND F. RENDL, *QAPLIB – a quadratic assignment problem library*, J. Global Opt., 10 (1997), pp. 391–403.
- [28] E. J. CANDÈS, X. LI, AND M. SOLTANOLKOTABI, *Phase retrieval from coded diffraction patterns*, Appl. Comp. Harmonic Anal., 39 (2015), pp. 277 – 299.
- [29] Y. CARMON, J. C. DUCHI, S. AARON, AND T. KEVIN, *A rank-1 sketch for matrix multiplicative weights*, in Proc. 32nd Conf. Learning Theory, vol. 99, 25–28 Jun 2019, pp. 589–623.
- [30] A. CHAI, M. MOSCOSO, AND G. PAPANICOLAOU, *Array imaging using intensity-only measurements*, Inv. Prob., 27 (2010), p. 015005.
- [31] D. CIFUENTES, *Burer–Monteiro guarantees for general semidefinite programs*, 2019, <https://arxiv.org/abs/1904.07147>.
- [32] K. L. CLARKSON, *Coresets, sparse greedy approximation, and the Frank–Wolfe algorithm*, ACM Trans. Algorithms, 6 (2010), pp. 63:1–63:30.
- [33] J. L. CONRAD, *Marburg virus hemorrhagic fever: Cytoarchitecture and histopathology*, <https://pixnio.com/de/wissenschaft/mikroskopische-aufnahmen/hamorrhagisches-fieber-marburg-virus/cytoarchitectural-histopathologischen-erkannt-niere-probe-marburg-patient> (accessed November 2019). Public domain (CC0 license).
- [34] C. DELORME AND S. POLJAK, *Laplacian eigenvalues and the maximum cut problem*, Math. Prog., Ser. A, 62 (1993), pp. 557–574.
- [35] C. DELORME AND S. POLJAK, *The performance of an eigenvalue bound on the max-cut problem in some classes of graphs*, Discrete Math., 111 (1993), pp. 145–156.
- [36] L. DING, A. YURTSEVER, V. CEVHER, J. A. TROPP, AND M. UDELL, *An optimal-storage approach to semidefinite programming using approximate complementarity*, 2019, <https://arxiv.org/abs/1902.03373>.
- [37] M. FAZEL, *Matrix Rank Minimization with Applications*, PhD thesis, Stanford University, Palo Alto, CA, USA, 2002.
- [38] J. R. FIENUP, *Phase retrieval algorithms: a comparison*, Appl. Optics, 21 (1982), pp. 2758–2769.
- [39] M. FRANK AND P. WOLFE, *An algorithm for quadratic programming*, Naval Res. Logistics Quart., 3 (1956), pp. 95–110.
- [40] D. GARBER AND E. HAZAN, *Approximating semidefinite programs in sublinear time*, in Adv. Neural Information Processing Systems 25, 2011.
- [41] D. GARBER AND E. HAZAN, *Sublinear time algorithms for approximate semidefinite programming*, Math. Prog., 158 (2016), pp. 329–361.
- [42] G. GIDEL, F. PEDREGOSA, AND S. LACOSTE-JULIEN, *Frank–Wolfe splitting via augmented Lagrangian method*, in Proc. 21st Int. Conf. Artificial Intelligence and Statistics, vol. 84, 2018, pp. 1456–1465.
- [43] A. GITTENS, *Topics in Randomized Numerical Linear Algebra*, PhD thesis, California Institute of Technology, Pasadena, CA, USA, 2013.
- [44] M. X. GOEMANS AND D. P. WILLIAMSON, *Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming*, J. ACM, 42 (1995), pp. 1115–1145.
- [45] N. HALKO, P. G. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Rev., 53 (2011), pp. 217–288.
- [46] E. HAZAN, *Sparse approximate solutions to semidefinite programs*, in LATIN 2008: Theoretical Infor-

- mathematics, 2008, pp. 306–316.
- [47] S. HOMER AND M. PEINADO, *Design and performance of parallel and distributed approximation algorithms for maxcut*, J. Parallel Distrib. Comput., 46 (1997), pp. 48 – 61.
- [48] R. HORSTMAYER, R. Y. CHEN, X. OU, B. AMES, J. A. TROPP, AND C. YANG, *Solving ptychography with a convex relaxation*, New J. Phys., 17 (2015), p. 053044.
- [49] Q. HUANG, Y. CHEN, AND L. GUIBAS, *Scalable semidefinite relaxation for maximum a posterior estimation*, in Proc. 31st Int. Conf. Machine Learning, vol. 32(2), 2014, pp. 64–72.
- [50] M. JAGGI, *Revisiting Frank–Wolfe: Projection-free sparse convex optimization*, in Proc. 30th Int. Conf. Machine Learning, vol. 28(1), 2013, pp. 427–435.
- [51] L. K. JONES, *A simple lemma on greedy approximation in Hilbert space and convergence rates for projection pursuit regression and neural network training*, Ann. Statist., 20 (1992), pp. 608–613.
- [52] R. JONKER AND A. VOLGENANT, *A shortest augmenting path algorithm for dense and sparse linear assignment problems*, Computing, 38 (1987), pp. 325–340.
- [53] A. P. KAMATH AND N. K. KARMARKAR, *A continuous method for computing bounds in integer quadratic optimization problems*, J. Global Optim., 2 (1991), pp. 229–241.
- [54] A. P. KAMATH AND N. K. KARMARKAR, *An $O(nL)$ iteration algorithm for computing bounds in quadratic optimization problems*, in Complexity in Numerical Optimization, World Scientific Publishing, July 1993, pp. 254–268.
- [55] R. M. KARP, *Reducibility among combinatorial problems*, in Complexity of computer computations, 1972, pp. 85–103.
- [56] K. KRISHNAN AND T. TERLAKY, *Interior point and semidefinite approaches in combinatorial optimization*, in Graph Theory and Combinatorial Optimization, Springer, Boston, MA, 2005, ch. 1.
- [57] J. KUCZYŃSKI AND H. WOŹNIAKOWSKI, *Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 1094–1122.
- [58] H. KUHN, *The Hungarian Method for the assignment problem*, Naval Res. Logistics Quart., 2 (1955), pp. 83–97.
- [59] B. KULIS, A. C. SURENDRAN, AND J. C. PLATT, *Fast low-rank semidefinite programming for embedding and clustering*, in Proc. 11th Int. Conf. Artificial Intelligence and Statistics, vol. 2, San Juan, Puerto Rico, 21–24 Mar 2007, pp. 235–242.
- [60] J. LAVAEI AND S. H. LOW, *Zero duality gap in optimal power flow problem*, IEEE Trans. Power Sys., 27 (2012), pp. 92–107.
- [61] Y. T. LEE AND S. PADMANABHAN, *An $\tilde{O}(m/\varepsilon^{3.5})$ -cost algorithm for semidefinite programs with diagonal constraints*, in Conference on Learning Theory, 2020, pp. 3069–3119.
- [62] R. B. LEHOUCQ, D. C. SORENSSEN, AND C. YANG, *ARPACK users’ guide*, vol. 6 of Software, Environments, and Tools, SIAM, Philadelphia, PA, 1998. Solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods.
- [63] E. S. LEVITIN AND B. T. POLJAK, *Minimization methods in the presence of constraints*, Ž. Vyčisl. Mat i Mat. Fiz., 6 (1966), pp. 787–823.
- [64] H. LI, G. C. LINDERMAN, A. SZLAM, K. P. STANTON, Y. KLUGER, AND M. TYGERT, *Algorithm 971: an implementation of a randomized algorithm for principal component analysis*, ACM Trans. Math. Software, 43 (2017), pp. Art. 28, 14.
- [65] Y.-F. LIU, X. LIU, AND S. MA, *On the nonergodic convergence rate of an inexact augmented lagrangian framework for composite convex programming*, Math. Oper. Res., 44 (2019), pp. 632–650.
- [66] E. M. LOIOLA, N. M. M. DE ABREU, P. O. BOAVENTURA-NETTO, P. HAHN, AND T. QUERIDO, *A survey for the quadratic assignment problem*, Europ. J. Oper. Res., 176 (2007), pp. 657–690.
- [67] A. MAJUMDAR, G. HALL, AND A. A. AHMADI, *A survey of recent scalability improvements for semidefinite programming with applications to machine learning, control, and robotics*. Available at <http://arXiv.org/abs/1908.05209>, Sep. 2019.
- [68] M. MESBAHI AND G. P. PAPAVALLOPOULOS, *On the rank minimization problem over a positive semidefinite linear matrix inequality*, IEEE Trans. Automatic Control, 42 (1997), pp. 239–243.
- [69] L. MIRSKY, *Symmetric gauge functions and unitarily invariant norms*, Quart. J. Math. Oxford Ser. (2), 11 (1960), pp. 50–59.
- [70] MOSEK APS, *The MOSEK optimization toolbox for MATLAB manual. Version 9.0.*, 2019.
- [71] J. MUNKRES, *Algorithms for the assignment and transportation problems*, J. SIAM, 5 (1957), pp. 32–38.

- [72] A. NEMIROVSKI, *Prox-method with rate of convergence $O(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems*, SIAM J. Optim., 15 (2004), pp. 229–251.
- [73] Y. NESTEROV, *Introductory Lectures on Convex Optimization: A Basic Course*, Springer US, 2004.
- [74] Y. NESTEROV, *Primal-dual subgradient methods for convex problems*, Math. Program., 120 (2009), pp. 221–259.
- [75] Y. NESTEROV AND A. NEMIROVSKI, *Self-concordant functions and polynomial time methods in convex programming*, Central Economic & Mathematic Institute report, USSR Acad. Sci., April 1990.
- [76] Y. NESTEROV AND A. NEMIROVSKI, *Interior-Point Polynomial Algorithms in Convex Programming*, SIAM, Philadelphia, PA, 1994.
- [77] B. N. PARLETT, *The symmetric eigenvalue problem*, vol. 20 of Classics in Applied Mathematics, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998, <https://doi.org/10.1137/1.9781611971163>, <https://doi-org.clsproxy.library.caltech.edu/10.1137/1.9781611971163>. Corrected reprint of the 1980 original.
- [78] P. PARRILO, *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*, PhD thesis, California Institute of Technology, Pasadena, CA, 2000.
- [79] G. PATAKI, *On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues*, Math. Oper. Res., 23 (1998), pp. 339–358.
- [80] R. PENG AND K. TANGWONGSAN, *Faster and simpler width-independent parallel algorithms for positive semidefinite programming*, in Proc. 24th Ann. ACM Symp. Parallelism in Algorithms and Architectures, 2012, pp. 101–108.
- [81] T. PUMIR, S. JELASSI, AND N. BOUMAL, *Smoothed analysis of the low-rank approach for smooth semidefinite programs*, in Adv. Neural Information Processing Systems 31, 2018, pp. 2281–2290.
- [82] A. RAGHUNATHAN, J. STEINHARDT, AND P. LIANG, *Semidefinite relaxations for certifying robustness to adversarial examples*, in Adv. Neural Information Processing Systems 31, 2018, pp. 10900–10910.
- [83] G. REINELT, *TSPLIB*, <http://comopt.ifi.uni-heidelberg.de/software/TSPLIB95/>.
- [84] D. M. ROSEN, L. CARLONE, A. S. BANDEIRA, AND J. J. LEONARD, *SE-sync: A certifiably correct algorithm for synchronization over the special euclidean group*, Int. J. Robot. Res., 38 (2019), pp. 95–125.
- [85] M. F. SAHIN, A. EFTEKHARI, A. ALACAOGLU, F. LATORRE, AND V. CEVHER, *An inexact augmented Lagrangian framework for nonconvex optimization with nonlinear constraints*, in Adv. Neural Information Processing Systems 32, 2019, pp. 13965–13977.
- [86] S. SAHNI AND T. GONZALEZ, *P-complete approximation problems*, J. ACM, 23 (1976), pp. 555–565.
- [87] A. SILVETI-FALLS, C. MOLINARI, AND J. FADILI, *Generalized conditional gradient with augmented Lagrangian for composite minimization*, 2019, <https://arxiv.org/abs/1901.01287>.
- [88] M. SOUTO, J. D. GARCIA, AND A. VEIGA, *Exploiting low-rank structure in semidefinite programming by approximate operator splitting*, 2018, <https://arxiv.org/abs/1810.05231>.
- [89] N. SREBRO, *Learning with matrix factorizations*, PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 2004.
- [90] J. STURM, *Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones*, Optimization Methods and Software, 11–12 (1999), pp. 625–653.
- [91] Y. SUN, Y. GUO, J. A. TROPP, AND M. UDELL, *Tensor random projection for low memory dimension reduction*, in NeurIPS Workshop on Relational Representation Learning, 2018.
- [92] K. A. TOH, M. J. TODD, AND R. H. TÜTÜNCÜ, *SDPT3 — a Matlab software package for semidefinite programming, optimization methods and software*, Optim. Methods Software, 11 (1999), pp. 545–581.
- [93] J. A. TROPP, A. YURTSEVER, M. UDELL, AND V. CEVHER, *Fixed-rank approximation of a positive-semidefinite matrix from streaming data*, in Adv. Neural Information Processing Systems 30, 2017, pp. 1225–1234.
- [94] J. A. TROPP, A. YURTSEVER, M. UDELL, AND V. CEVHER, *Practical sketching algorithms for low-rank matrix approximation*, SIAM J. Matrix Anal. Appl., 38 (2017), pp. 1454–1485.
- [95] J. A. TROPP, A. YURTSEVER, M. UDELL, AND V. CEVHER, *Streaming low-rank matrix approximation with an application to scientific simulation*, SIAM J. Sci. Comput., 41 (2019), pp. A2430–A2463.
- [96] K. TSUDA, G. RÄTSCH, AND M. K. WARMUTH, *Matrix exponentiated gradient updates for on-line learning and Bregman projections*, J. Mach. Learn. Res., 6 (2005), pp. 995–1018.

- [97] I. WALDSPURGER AND A. WATERS, *Rank optimality for the Burer–Monteiro factorization*, 2018, <https://arxiv.org/abs/1812.03046>.
- [98] Z. WEN, D. GOLDFARB, S. MA, AND K. SCHEINBERG, *Row by row methods for semidefinite programming*, IEOR report, Columbia University, 2009.
- [99] Z. WEN, D. GOLDFARB, AND W. YIN, *Alternating direction augmented Lagrangian methods for semidefinite programming*, *Math. Program. Comp.*, 2 (2010), pp. 203–230.
- [100] L. YANG, D. SUN, AND K.-C. TOH, *SDPNAL+: A majorized semismooth Newton-CG augmented Lagrangian method for semidefinite programming with nonnegative constraints*, *Math. Prog. Comput.*, 7 (2015), pp. 331–366.
- [101] Y. YE, *GSet random graphs*, <https://www.cise.ufl.edu/research/sparse/matrices/Gset/> (accessed October 2019).
- [102] A. YURTSEVER, O. FERCOQ, AND V. CEVHER, *A conditional-gradient-based augmented Lagrangian framework*, in *Proc. 36th Int. Conf. Machine Learning*, vol. 97, 09–15 Jun 2019, pp. 7272–7281.
- [103] A. YURTSEVER, O. FERCOQ, F. LOCATELLO, AND V. CEVHER, *A conditional gradient framework for composite convex minimization with applications to semidefinite programming*, in *Proc. 35th Intl. Conf. Machine Learning*, vol. 80, 2018, pp. 5727–5736.
- [104] A. YURTSEVER, Y.-P. HSIEH, AND V. CEVHER, *Scalable convex methods for phase retrieval*, in *6th Int. IEEE Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, 2015, pp. 381–384.
- [105] A. YURTSEVER, Q. TRAN DINH, AND V. CEVHER, *A universal primal-dual convex optimization framework*, in *Adv. Neural Information Processing Systems* 28, 2015, pp. 3150–3158.
- [106] A. YURTSEVER, M. UDELL, J. TROPP, AND V. CEVHER, *Sketchy Decisions: Convex Low-Rank Matrix Optimization with Optimal Storage*, in *Proc. 20th Int. Conf. Artificial Intelligence and Statistics*, vol. 54, 20–22 Apr 2017, pp. 1188–1196.
- [107] M. ZASLAVSKIY, F. BACH, AND J. VERT, *A path following algorithm for the graph matching problem*, *IEEE Trans. Pattern Anal. Mach. Intel.*, 31 (2009), pp. 2227–2242.
- [108] Q. ZHAO, S. E. KARISCH, F. RENDL, AND H. WOLKOWICZ, *Semidefinite programming relaxations for the quadratic assignment problem*, *J. Comb. Optim.*, 2 (1998), pp. 71–109.
- [109] X.-Y. ZHAO, D. SUN, AND K.-C. TOH, *A Newton-CG augmented Lagrangian method for semidefinite programming*, *SIAM J. Optim.*, 20 (2010), pp. 1737–1765.
- [110] G. ZHENG, R. HORSTMAYER, AND C. YANG, *Wide-field, high-resolution Fourier ptychography microscopy*, *Nat. Photonics*, 7 (2013), pp. 739–745.

SUPPLEMENTARY MATERIALS: Scalable Semidefinite Programming*

Alp Yurtsever[†], Joel A. Tropp[‡], Olivier Fercoq[§], Madeleine Udell[¶], and Volkan Cevher[†]

SM1. Analysis of the CGAL Algorithm. This section contains a complete analysis of the convergence of the CGAL algorithm and its arithmetic costs. For simplicity, we have specialized this presentation to the model problem that is the focus of this paper; many extensions are possible. The convergence results here are adapted from the initial paper [SM27] on the CGAL algorithm. The analysis of the approximate eigenvector computation and the detailed results for the model problem are new. Empirical work suggests that the analysis is still qualitatively suboptimal, which is a direction for future research.

SM1.1. The model problem. We focus on solving the optimization template

$$(SM1.1) \quad \text{minimize } \langle \mathbf{C}, \mathbf{X} \rangle \quad \text{subject to } \mathcal{A}\mathbf{X} = \mathbf{b}, \quad \mathbf{X} \in \alpha\Delta_n.$$

The constraint set $\alpha\Delta_n$ consists of $n \times n$ psd matrices with trace α . The objective function depends on a matrix $\mathbf{C} \in \mathbb{S}_n$. The linear constraints are determined by the linear map $\mathcal{A} : \mathbb{S}_n \rightarrow \mathbb{R}^d$ and the right-hand side vector $\mathbf{b} \in \mathbb{R}^d$.

SM1.2. Elements of Lagrangian duality. Introduce the Lagrangian function

$$(SM1.2) \quad L(\mathbf{X}; \mathbf{y}) := \langle \mathbf{C}, \mathbf{X} \rangle + \langle \mathbf{y}, \mathcal{A}\mathbf{X} - \mathbf{b} \rangle \quad \text{for } \mathbf{X} \in \alpha\Delta_n \text{ and } \mathbf{y} \in \mathbb{R}^d.$$

We assume that the Lagrangian admits at least one saddle point $(\mathbf{X}_*, \mathbf{y}_*) \in \alpha\Delta_n \times \mathbb{R}^d$:

$$(SM1.3) \quad L(\mathbf{X}_*, \mathbf{y}) \leq L(\mathbf{X}_*, \mathbf{y}_*) \leq L(\mathbf{X}; \mathbf{y}_*) \quad \text{for all } \mathbf{X} \in \alpha\Delta_n \text{ and } \mathbf{y} \in \mathbb{R}^d.$$

This hypothesis is guaranteed under Slater’s condition. Write p_* for the shared extremal value of the dual and primal problems:

$$(SM1.4) \quad \max_{\mathbf{y} \in \mathbb{R}^d} \min_{\mathbf{X} \in \alpha\Delta_n} L(\mathbf{X}, \mathbf{y}) = p_* = \min_{\mathbf{X} \in \alpha\Delta_n} \sup_{\mathbf{y} \in \mathbb{R}^d} L(\mathbf{X}, \mathbf{y}).$$

In particular, note that $p_* = \langle \mathbf{C}, \mathbf{X}_* \rangle$.

*Submitted to the editors 6 December 2019. Revised on 18 January 2020, 24 July 2020 and 12 November 2020.

Funding: VC and AY have received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program under the grant agreement number 725594 (time-data) and the Swiss National Science Foundation (SNSF) under the grant number 200021_178865/1. JAT gratefully acknowledges ONR Awards N00014-11-1-0025, N00014-17-1-2146, and N00014-18-1-2363. MU gratefully acknowledges DARPA Award FA8750-17-2-0101. Part of this research is conducted while AY is at Massachusetts Institute of Technology, Cambridge, MA, USA. AY acknowledges the Early Postdoc.Mobility Fellowship P2ELP2_187955 from the Swiss National Science Foundation and partial postdoctoral support from the NSF-CAREER grant IIS-1846088.

[†]Laboratory for Information and Inference Systems, Department of Electrical Engineering, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland (alp.yurtsever@epfl.ch, volkan.cevher@epfl.ch, <https://lions.epfl.ch/>).

[‡]Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA, USA (jtropp@cms.caltech.edu, <http://users.cms.caltech.edu/~jtropp>).

[§]Laboratoire Traitement et Communication d’Information, Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France (olivier.fercoq@telecom-paris.fr, <https://perso.telecom-paristech.fr/ofercq/>).

[¶]Department of Operations Research and Information Engineering, Cornell University, Ithaca, NY, USA (udell@cornell.edu, <https://people.orie.cornell.edu/mru8/>).

SM1.3. The CGAL iteration. The CGAL iteration solves the model problem (SM1.1) using an augmented Lagrangian method where the primal step is inspired by the conditional gradient method. See Algorithm SM1.1 for pseudocode.

Let $\beta_0 > 0$ be an initial smoothing parameter. Fix a schedule for the step size parameter η_t and the smoothing parameter β_t :

$$(SM1.5) \quad \eta_t := \frac{2}{t+1} \quad \text{and} \quad \beta_t := \beta_0 \sqrt{t+1} \quad \text{for } t = 1, 2, 3, \dots$$

Define the augmented Lagrangian L_t with smoothing parameter β_t

$$(SM1.6) \quad L_t(\mathbf{X}; \mathbf{y}) := \langle \mathbf{C}, \mathbf{X} \rangle + \langle \mathbf{y}, \mathcal{A}\mathbf{X} - \mathbf{b} \rangle + \frac{1}{2}\beta_t \|\mathcal{A}\mathbf{X} - \mathbf{b}\|^2.$$

The CGAL algorithm solves the model problem (SM1.1) by alternating between primal and dual update steps on (SM1.6), while increasing the smoothing parameter.

Fix the initial iterates

$$(SM1.7) \quad \mathbf{X}_1 = \mathbf{0} \in \mathbb{S}_n \quad \text{and} \quad \mathbf{y}_1 = \mathbf{0} \in \mathbb{R}^d.$$

At each iteration $t = 1, 2, 3, \dots$, we implicitly compute the partial derivative

$$(SM1.8) \quad \mathbf{D}_t := \partial_{\mathbf{X}} L_t(\mathbf{X}_t; \mathbf{y}_t) = \mathbf{C} + \mathcal{A}^*(\mathbf{y}_t + \beta_t(\mathcal{A}\mathbf{X}_t - \mathbf{b})).$$

Then we identify a primal update direction $\mathbf{H}_t \in \alpha\Delta_n$ that satisfies

$$(SM1.9) \quad \langle \mathbf{D}_t, \mathbf{H}_t \rangle \leq \min_{\mathbf{H} \in \alpha\Delta_n} \langle \mathbf{D}_t, \mathbf{H} \rangle + \frac{\alpha\beta_0}{\beta_t} \|\mathbf{D}_t\|.$$

In other words, the smoothing parameter β_t also controls the amount of inexactness we are willing to tolerate in the linear minimization at iteration t . We construct the next primal iterate via the rule

$$(SM1.10) \quad \mathbf{X}_{t+1} = \mathbf{X}_t + \eta_t(\mathbf{H}_t - \mathbf{X}_t) \in \alpha\Delta_n.$$

The dual update takes the form

$$(SM1.11) \quad \mathbf{y}_{t+1} = \mathbf{y}_t + \gamma_t(\mathcal{A}\mathbf{X}_{t+1} - \mathbf{b}).$$

We select the largest dual step size parameter $0 \leq \gamma_t \leq \beta_0$ that satisfies the condition

$$(SM1.12) \quad \gamma_t \|\mathcal{A}\mathbf{X}_{t+1} - \mathbf{b}\|^2 \leq \beta_t \eta_t^2 \alpha^2 \|\mathcal{A}\|^2.$$

The latter inequality always holds when $\gamma_t = 0$. We will also choose the dual step size to maintain a bounded travel condition:

$$(SM1.13) \quad \|\mathbf{y}_t\| \leq K.$$

If the bounded travel condition holds at iteration $t-1$, then the choice $\gamma_t = 0$ ensures that it holds at iteration t . In practice, it is not necessary to enforce (SM1.13). The iteration continues until it reaches a maximum iteration count T_{\max} .

Algorithm SM1.1 CGAL for the model problem (SM1.1)

Input: Problem data for (SM1.1) implemented via the primitives (2.4); number T of iterations

Output: Approximate solution \mathbf{X}_T to (2.2)

Recommendation: Set $T \approx \varepsilon^{-1}$ to achieve ε -optimal solution (SM1.17)

```

1 function CGAL( $T$ )
2   Scale problem data (subsection 7.1.1)                                ▷ [opt] Recommended!
3    $\beta_0 \leftarrow 1$  and  $K \leftarrow +\infty$                             ▷ Default parameters
4    $\mathbf{X} \leftarrow \mathbf{0}_{n \times n}$  and  $\mathbf{y} \leftarrow \mathbf{0}_d$ 
5   for  $t \leftarrow 1, 2, 3, \dots, T$  do
6      $\beta \leftarrow \beta_0 \sqrt{t+1}$  and  $\eta \leftarrow 2/(t+1)$ 
7      $(\xi, \mathbf{v}) \leftarrow \text{ApproxMinEvec}(\mathbf{C} + \mathcal{A}^*(\mathbf{y} + \beta(\mathcal{A}\mathbf{X} - \mathbf{b})); q_t)$ 
                                                    ▷ Algorithm 4.1 with  $q_t = t^{1/4} \log n$ 
                                                    ▷ Implement with primitives (2.4) 12!
8      $\mathbf{X} \leftarrow (1 - \eta)\mathbf{X} + \eta(\alpha \mathbf{v}\mathbf{v}^*)$ 
9      $\mathbf{y} \leftarrow \mathbf{y} + \gamma(\mathcal{A}\mathbf{X} - \mathbf{b})$                                 ▷ Step size  $\gamma$  satisfies (SM1.12) and (SM1.13)
    
```

SM1.4. Distributed computation. Since CGAL builds on the augmented Lagrangian framework, we can apply it even when the problem is too large to solve on one computational node. In particular, when d is large, it may be advantageous to partition the constraint matrices \mathbf{A}_i and the associated dual variables y_i among several workers. Distributed CGAL has a structure similar to the alternating directions method of multipliers (ADMM) [SM5].

SM1.5. Theoretical analysis of CGAL. We develop two results on the behavior of CGAL. The first concerns its convergence to optimality, and the second concerns the computational resource usage.

SM1.5.1. Convergence. The first result demonstrates that the CGAL algorithm always converges to a primal optimal solution of the model problem (SM1.1). This result is adapted from [SM27]; a complete proof appears below in subsection SM1.6.

Theorem SM1.1 (CGAL: Convergence). *Define*

$$E := 6\alpha^2 \|\mathcal{A}\|^2 + \alpha(\|\mathbf{C}\| + K\|\mathcal{A}\|).$$

The primal iterates $\{\mathbf{X}_t : t = 2, 3, 4, \dots\}$ generated by the CGAL iteration satisfy the a priori bounds

$$(SM1.14) \quad \langle \mathbf{C}, \mathbf{X}_t \rangle - p_\star \leq \frac{2\beta_0 E}{\sqrt{t}} + K \cdot \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|;$$

$$(SM1.15) \quad -\|\mathbf{y}_\star\| \cdot \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\| \leq \langle \mathbf{C}, \mathbf{X}_t \rangle - p_\star;$$

$$(SM1.16) \quad \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\| \leq \frac{2\beta_0^{-1}(K + \|\mathbf{y}_\star\|) + 2\sqrt{E}}{\sqrt{t}}.$$

The a priori bounds ensure that

$$(SM1.17) \quad \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\| \leq \varepsilon \quad \text{and} \quad |\langle \mathbf{C}, \mathbf{X}_t \rangle - p_\star| \leq \varepsilon$$

within $O(\varepsilon^{-2})$ iterations. The big- O suppresses constants that depend on the problem data $(\alpha, \|\mathcal{A}\|, \|\mathbf{C}\|)$ and the algorithm parameters β_0 and K .

SM1.5.2. Problem scaling. Theorem SM1.1 indicates that it is valuable to scale the model problem (SM1.1) so that $\alpha = \|\mathbf{C}\| = \|\mathcal{A}\| = 1$. In this case, a good choice for the smoothing parameter is $\beta_0 = 1$. Nevertheless, the algorithm converges, regardless of the scaling and regardless of the parameter choices β_0 and K . We use a slightly different scaling in practice; see subsection 7.1.1.

SM1.5.3. Bounded travel? Theorem SM1.1 suggests that the optimal choice for the travel bound is $K = 0$. In other words, the dual vector $\mathbf{y}_t = \mathbf{0}$, and it does not evolve. The algorithm that results from this choice is called HomotopyCGM [SM28]. The numerical work on CGAL, reported in [SM27], makes it clear that updating the dual variable, as CGAL does, allows for substantial performance improvements over HomotopyCGM. Unfortunately, the theory developed in [SM27], and echoed here, does not comprehend the reason for this phenomenon. This is an obvious direction for further research.

SM1.6. Proof of Theorem SM1.1. In this section, we establish the convergence guarantee stated in Theorem SM1.1.

SM1.6.1. Analysis of the primal update. The first steps in the proof address the role of the primal update rule (SM1.10). The arguments parallels the standard convergence analysis [SM12] of CGM, applied to the function

$$(SM1.18) \quad f_t(\mathbf{X}) := L_t(\mathbf{X}; \mathbf{y}_t) = \langle \mathbf{C}, \mathbf{X} \rangle + \langle \mathbf{y}_t, \mathcal{A}\mathbf{X} - \mathbf{b} \rangle + \frac{1}{2}\beta_t\|\mathcal{A}\mathbf{X} - \mathbf{b}\|^2.$$

Observe that the gradient $\nabla f_t(\mathbf{X}_t)$ coincides with the partial derivative (SM1.8).

To begin, we exploit the smoothness of f_t to control the change in its value at adjacent primal iterates. The function f_t is convex on \mathbb{S}_n , and its gradient has Lipschitz constant $\beta_t\|\mathcal{A}\|^2$, so

$$f_t(\mathbf{X}_+) - f_t(\mathbf{X}) \leq \langle \nabla f_t(\mathbf{X}), \mathbf{X}_+ - \mathbf{X} \rangle + \frac{1}{2}\beta_t\|\mathcal{A}\|^2\|\mathbf{X}_+ - \mathbf{X}\|_{\mathbb{F}}^2 \quad \text{for } \mathbf{X}, \mathbf{X}_+ \in \mathbb{S}_n.$$

In particular, with $\mathbf{X}_+ = \mathbf{X}_{t+1}$ and $\mathbf{X} = \mathbf{X}_t$, we obtain

$$(SM1.19) \quad \begin{aligned} f_t(\mathbf{X}_{t+1}) &\leq f_t(\mathbf{X}_t) + \langle \mathbf{D}_t, \mathbf{X}_{t+1} - \mathbf{X}_t \rangle + \frac{1}{2}\beta_t\|\mathcal{A}\|^2\|\mathbf{X}_{t+1} - \mathbf{X}_t\|_{\mathbb{F}}^2 \\ &= f_t(\mathbf{X}_t) + \eta_t \langle \mathbf{D}_t, \mathbf{H}_t - \mathbf{X}_t \rangle + \frac{1}{2}\beta_t\eta_t^2\|\mathcal{A}\|^2\|\mathbf{H}_t - \mathbf{X}_t\|_{\mathbb{F}}^2 \\ &\leq f_t(\mathbf{X}_t) + \eta_t \langle \mathbf{D}_t, \mathbf{H}_t - \mathbf{X}_t \rangle + \beta_t\eta_t^2\alpha^2\|\mathcal{A}\|^2. \end{aligned}$$

The second identity follows from the update rule (SM1.10). The bound on the last term depends on the fact that the constraint set $\alpha\mathbf{\Delta}_n$ has Frobenius-norm diameter $\alpha\sqrt{2}$.

Next, we use the construction of the primal update to control the linear term in the last display. The update direction \mathbf{H}_t satisfies the inequality (SM1.9), so

$$(SM1.20) \quad \begin{aligned} \langle \mathbf{D}_t, \mathbf{H}_t - \mathbf{X}_t \rangle &\leq \min_{\mathbf{H} \in \alpha\mathbf{\Delta}_n} \langle \mathbf{D}_t, \mathbf{H} - \mathbf{X}_t \rangle + \frac{\alpha\beta_0}{\beta_t}\|\mathbf{D}_t\| \\ &\leq \langle \mathbf{D}_t, \mathbf{X}_\star - \mathbf{X}_t \rangle + \frac{\alpha\beta_0}{\beta_t}\|\mathbf{D}_t\|. \end{aligned}$$

The second inequality depends on the fact that $\mathbf{X}_\star \in \alpha \Delta_n$.

We can use the explicit formula (SM1.8) for the derivative \mathbf{D}_t to control the two terms in (SM1.20). First,

$$\begin{aligned}
 \langle \mathbf{D}_t, \mathbf{X}_\star - \mathbf{X}_t \rangle &= \langle \mathbf{C} + \mathcal{A}^*(\mathbf{y}_t + \beta_t(\mathcal{A}\mathbf{X}_t - \mathbf{b})), \mathbf{X}_\star - \mathbf{X}_t \rangle \\
 &= \langle \mathbf{C}, \mathbf{X}_\star - \mathbf{X}_t \rangle - \langle \mathbf{y}_t, \mathcal{A}\mathbf{X}_t - \mathbf{b} \rangle - \beta_t \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|^2 \\
 &= \langle \mathbf{C}, \mathbf{X}_\star \rangle - f_t(\mathbf{X}_t) - \frac{1}{2}\beta_t \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|^2.
 \end{aligned}
 \tag{SM1.21}$$

We have invoked the definition of the adjoint \mathcal{A}^* and the fact that $\mathcal{A}\mathbf{X}_\star = \mathbf{b}$. Last, we used the definition (SM1.18) to identify the quantity $f_t(\mathbf{X}_t)$. Second,

$$\begin{aligned}
 \|\mathbf{D}_t\| &= \|\mathbf{C} + \mathcal{A}^*(\mathbf{y}_t + \beta_t(\mathcal{A}\mathbf{X}_t - \mathbf{b}))\| \\
 &\leq \|\mathbf{C}\| + K\|\mathcal{A}\| + \beta_t \|\mathcal{A}\| \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|.
 \end{aligned}
 \tag{SM1.22}$$

We have used the assumption that \mathbf{y}_t satisfies the bounded travel condition (SM1.13).

Combine the last three displays to obtain the estimate

$$\begin{aligned}
 \langle \mathbf{D}_t, \mathbf{H}_t - \mathbf{X}_t \rangle &\leq \langle \mathbf{C}, \mathbf{X}_\star \rangle - f_t(\mathbf{X}_t) \\
 &\quad + \left(\alpha\beta_0 \|\mathcal{A}\| \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\| - \frac{1}{2}\beta_t \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|^2 \right) + \frac{\alpha\beta_0}{\beta_t} (\|\mathbf{C}\| + K\|\mathcal{A}\|).
 \end{aligned}
 \tag{SM1.23}$$

We can now control the decrease in the function f_t between adjacent primal iterates. Combine the displays (SM1.19) and (SM1.23) to arrive at the bound

$$\begin{aligned}
 f_t(\mathbf{X}_{t+1}) &\leq (1 - \eta_t) f_t(\mathbf{X}_t) + \eta_t \langle \mathbf{C}, \mathbf{X}_\star \rangle \\
 &\quad + \left(\eta_t \alpha \beta_0 \|\mathcal{A}\| \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\| - \frac{1}{2} \beta_t \eta_t \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|^2 \right) + \left(\beta_t \eta_t^2 \alpha^2 \|\mathcal{A}\|^2 + \frac{\eta_t \alpha \beta_0}{\beta_t} (\|\mathbf{C}\| + K\|\mathcal{A}\|) \right).
 \end{aligned}$$

Subtract $p_\star = \langle \mathbf{C}, \mathbf{X}_\star \rangle$ from both sides to arrive at

$$\begin{aligned}
 f_t(\mathbf{X}_{t+1}) - p_\star &\leq (1 - \eta_t) (f_t(\mathbf{X}_t) - p_\star) \\
 &\quad + \left(\eta_t \alpha \beta_0 \|\mathcal{A}\| \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\| - \frac{1}{2} \beta_t \eta_t \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|^2 \right) + \left(\beta_t \eta_t^2 \alpha^2 \|\mathcal{A}\|^2 + \frac{\eta_t \alpha \beta_0}{\beta_t} (\|\mathbf{C}\| + K\|\mathcal{A}\|) \right).
 \end{aligned}$$

Finally, use the definition (SM1.18) again to pass back to the augmented Lagrangian:

$$\begin{aligned}
 L_t(\mathbf{X}_{t+1}; \mathbf{y}_t) - p_\star &\leq (1 - \eta_t) (L_t(\mathbf{X}_t; \mathbf{y}_t) - p_\star) \\
 &\quad + \left(\eta_t \alpha \beta_0 \|\mathcal{A}\| \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\| - \frac{1}{2} \beta_t \eta_t \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|^2 \right) \\
 &\quad + \left(\beta_t \eta_t^2 \alpha^2 \|\mathcal{A}\|^2 + \frac{\eta_t \alpha \beta_0}{\beta_t} (\|\mathbf{C}\| + K\|\mathcal{A}\|) \right).
 \end{aligned}
 \tag{SM1.24}$$

This bound describes the evolution of the augmented Lagrangian as the primal iterate advances. But we still need to include the effects of increasing the smoothing parameter and updating the dual variable.

SM1.6.2. Analysis of the smoothing update. To continue, observe that updates to the smoothing parameter have a controlled impact on the augmented Lagrangian (SM1.6):

$$L_t(\mathbf{X}_t; \mathbf{y}_t) - L_{t-1}(\mathbf{X}_t; \mathbf{y}_t) = \frac{1}{2}(\beta_t - \beta_{t-1})\|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|^2.$$

Add and subtract $L_{t-1}(\mathbf{X}_t; \mathbf{y}_t)$ in the large parenthesis in the first line of (SM1.24) and invoke the last identity to obtain

$$\begin{aligned} \text{(SM1.25)} \quad L_t(\mathbf{X}_{t+1}; \mathbf{y}_t) - p_\star &\leq (1 - \eta_t)\left(L_{t-1}(\mathbf{X}_t; \mathbf{y}_t) - p_\star\right) \\ &+ \left(\eta_t\alpha\beta_0\|\mathcal{A}\|\|\mathcal{A}\mathbf{X}_t - \mathbf{b}\| + \frac{1}{2}[(1 - \eta_t)(\beta_t - \beta_{t-1}) - \beta_t\eta_t]\|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|^2\right) \\ &\quad + \left(\beta_t\eta_t^2\alpha^2\|\mathcal{A}\|^2 + \frac{\eta_t\alpha\beta_0}{\beta_t}(\|\mathbf{C}\| + K\|\mathcal{A}\|)\right). \end{aligned}$$

The next step is to develop a uniform bound on the terms in the second line so that we can ignore the role of the feasibility gap $\|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|$ in the subsequent calculations. The choice (SM1.5) of the parameters ensures that

$$(1 - \eta_t)(\beta_t - \beta_{t-1}) - \beta_t\eta_t < \frac{-\beta_0^2}{\beta_t}.$$

Introduce this bound into the second line of (SM1.25) and maximize the resulting concave quadratic function to reach

$$\begin{aligned} \eta_t\alpha\beta_0\|\mathcal{A}\|\|\mathcal{A}\mathbf{X}_t - \mathbf{b}\| + \frac{1}{2}[(1 - \eta_t)(\beta_t - \beta_{t-1}) - \beta_t\eta_t]\|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|^2 \\ \leq \eta_t\alpha\beta_0\|\mathcal{A}\|\|\mathcal{A}\mathbf{X}_t - \mathbf{b}\| - \frac{\beta_0^2}{2\beta_t}\|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|^2 \leq \beta_t\eta_t^2\alpha^2\|\mathcal{A}\|^2. \end{aligned}$$

Substitute the last display into (SM1.25) to determine that

$$\begin{aligned} \text{(SM1.26)} \quad L_t(\mathbf{X}_{t+1}; \mathbf{y}_t) - p_\star &\leq (1 - \eta_t)\left(L_{t-1}(\mathbf{X}_t; \mathbf{y}_t) - p_\star\right) \\ &\quad + \left(2\beta_t\eta_t^2\alpha^2\|\mathcal{A}\|^2 + \frac{\eta_t\alpha\beta_0}{\beta_t}(\|\mathbf{C}\| + K\|\mathcal{A}\|)\right). \end{aligned}$$

To develop a recursion, we need to assess how the left-hand side changes when we update the dual variable.

SM1.6.3. Analysis of the dual update. To incorporate the dual update, observe that

$$\begin{aligned} L_t(\mathbf{X}_{t+1}; \mathbf{y}_{t+1}) &= L_t(\mathbf{X}_{t+1}; \mathbf{y}_t) + \langle \mathbf{y}_{t+1} - \mathbf{y}_t, \mathcal{A}\mathbf{X}_{t+1} - \mathbf{b} \rangle \\ &= L_t(\mathbf{X}_{t+1}; \mathbf{y}_t) + \gamma_t\|\mathcal{A}\mathbf{X}_{t+1} - \mathbf{b}\|^2 \\ &\leq L_t(\mathbf{X}_{t+1}; \mathbf{y}_t) + \beta_t\eta_t^2\alpha^2\|\mathcal{A}\|^2. \end{aligned}$$

The first relation is simply the definition (SM1.6) of the augmented Lagrangian, while the second relation depends on the dual update rule (SM1.11). The last step follows from the selection rule (SM1.12) for the dual step size parameter.

Introduce the latter display into (SM1.25) to discover that

$$(SM1.27) \quad L_t(\mathbf{X}_{t+1}; \mathbf{y}_{t+1}) - p_\star \leq (1 - \eta_t) \left(L_{t-1}(\mathbf{X}_t; \mathbf{y}_t) - p_\star \right) \\ + \left(3\beta_t \eta_t^2 \alpha^2 \|\mathcal{A}\|^2 + \frac{\eta_t \alpha \beta_0}{\beta_t} (\|\mathbf{C}\| + K \|\mathcal{A}\|) \right).$$

We have developed a recursion for the value of the augmented Lagrangian as the iterates and the smoothing parameter evolve.

SM1.6.4. Solving the recursion. Next, we solve the recursion (SM1.27). We assert that

$$(SM1.28) \quad L_t(\mathbf{X}_{t+1}; \mathbf{y}_{t+1}) - p_\star < \frac{2\beta_0}{\sqrt{t+1}} [6\alpha^2 \|\mathcal{A}\|^2 + \alpha(\|\mathbf{C}\| + K \|\mathcal{A}\|)] \\ =: \frac{2\beta_0 E}{\sqrt{t+1}} \quad \text{for } t = 1, 2, 3, \dots$$

For the case $t = 1$, the definition (SM1.5) ensures that $\eta_1 = 1$ and $\beta_1 = \beta_0 \sqrt{2}$, so the bound (SM1.28) follows instantly from (SM1.27). When $t > 1$, an inductive argument using the recursion (SM1.27) and the bound (SM1.28) for $t - 1$ ensures that

$$L_t(\mathbf{X}_{t+1}; \mathbf{y}_{t+1}) - p_\star \leq \left[\frac{t-1}{t+1} \cdot \frac{2\beta_0}{\sqrt{t}} + \frac{2\beta_0}{(t+1)^{3/2}} \right] [6\alpha^2 \|\mathcal{A}\|^2 + \alpha(\|\mathbf{C}\| + K \|\mathcal{A}\|)] \\ < \frac{2\beta_0}{\sqrt{t+1}} [6\alpha^2 \|\mathcal{A}\|^2 + \alpha(\|\mathbf{C}\| + K \|\mathcal{A}\|)].$$

We have introduced the stated values (SM1.5) of the step size and smoothing parameters. The induction proceeds, and we conclude that (SM1.28) is valid.

SM1.6.5. Bound for the suboptimality of the objective. We are prepared to develop an upper bound on the suboptimality of the objective of the model problem (SM1.1). The definition (SM1.6) of the augmented Lagrangian directly implies that

$$(SM1.29) \quad \langle \mathbf{C}, \mathbf{X}_{t+1} \rangle - p_\star = L_t(\mathbf{X}_{t+1}; \mathbf{y}_{t+1}) - p_\star - \langle \mathbf{y}_{t+1}, \mathcal{A} \mathbf{X}_{t+1} - \mathbf{b} \rangle \\ - \frac{1}{2} \beta_t \|\mathcal{A} \mathbf{X}_{t+1} - \mathbf{b}\|^2.$$

Continuing from here,

$$\langle \mathbf{C}, \mathbf{X}_{t+1} \rangle - p_\star \leq \frac{2\beta_0 E}{\sqrt{t+1}} + K \cdot \|\mathcal{A} \mathbf{X}_{t+1} - \mathbf{b}\|.$$

The first identity follows from definition (SM1.6) of the augmented Lagrangian. The bound relies on (SM1.28) and the Cauchy–Schwarz inequality. We have also used the bounded travel condition (SM1.13). This establishes (SM1.14).

SM1.6.6. Bound for the superoptimality of the objective. The CGAL iterates \mathbf{X}_t are generally infeasible for (SM1.1), so they can be superoptimal. Nevertheless, we can easily control the superoptimality. By the saddle-point properties (SM1.3) and (SM1.4), the Lagrangian (SM1.2) satisfies

$$(SM1.30) \quad p_\star = \max_{\mathbf{y}} \min_{\mathbf{X} \in \alpha \Delta_n} L(\mathbf{X}; \mathbf{y}) \leq L(\mathbf{X}_{t+1}; \mathbf{y}_\star) = \langle \mathbf{C}, \mathbf{X}_{t+1} \rangle + \langle \mathbf{y}_\star, \mathcal{A} \mathbf{X}_{t+1} - \mathbf{b} \rangle.$$

Invoke the Cauchy–Schwarz inequality and rearrange to determine that

$$-\|\mathbf{y}_\star\| \cdot \|\mathcal{A} \mathbf{X}_{t+1} - \mathbf{b}\| \leq \langle \mathbf{C}, \mathbf{X}_{t+1} \rangle - p_\star.$$

This establishes (SM1.15).

SM1.6.7. Bound for the infeasibility of the iterates. Next, we demonstrate that the iterates converge toward the feasible set of (SM1.1). Combine (SM1.29) and (SM1.30) and rearrange to see that

$$\langle \mathbf{y}_{t+1} - \mathbf{y}_\star, \mathcal{A} \mathbf{X}_{t+1} - \mathbf{b} \rangle \leq L_t(\mathbf{X}_{t+1}; \mathbf{y}_{t+1}) - p_\star - \frac{1}{2} \beta_t \|\mathcal{A} \mathbf{X}_{t+1} - \mathbf{b}\|^2.$$

Bound the left-hand side below using Cauchy–Schwarz and the right-hand side above using (SM1.28):

$$-\|\mathbf{y}_{t+1} - \mathbf{y}_\star\| \cdot \|\mathcal{A} \mathbf{X}_{t+1} - \mathbf{b}\| \leq \frac{2\beta_0 E}{\sqrt{t+1}} - \frac{1}{2} \beta_t \|\mathcal{A} \mathbf{X}_{t+1} - \mathbf{b}\|^2.$$

Solve this quadratic inequality to obtain the bound

$$\begin{aligned} \|\mathcal{A} \mathbf{X}_{t+1} - \mathbf{b}\| &\leq \beta_t^{-1} \left(\|\mathbf{y}_{t+1} - \mathbf{y}_\star\| + \sqrt{\|\mathbf{y}_{t+1} - \mathbf{y}_\star\|^2 + \frac{4\beta_t \beta_0 E}{\sqrt{t+1}}} \right) \\ &\leq \beta_t^{-1} \left(2\|\mathbf{y}_{t+1} - \mathbf{y}_\star\| + \sqrt{4\beta_0^2 E} \right) \\ &= \frac{2\beta_0^{-1} \|\mathbf{y}_{t+1} - \mathbf{y}_\star\| + 2\sqrt{E}}{\sqrt{t+1}}. \end{aligned}$$

The second inequality depends on the definition (SM1.5) of the smoothing parameter β_t and the subadditivity of the square root.

Finally, we control the dual error using the bounded travel condition (SM1.13):

$$\|\mathbf{y}_{t+1} - \mathbf{y}_\star\| \leq \|\mathbf{y}_{t+1}\| + \|\mathbf{y}_\star\| \leq K + \|\mathbf{y}_\star\|.$$

The last two displays yield

$$\|\mathcal{A} \mathbf{X}_{t+1} - \mathbf{b}\| \leq \frac{2\beta_0^{-1}(K + \|\mathbf{y}_\star\|) + 2\sqrt{E}}{\sqrt{t+1}}.$$

This confirms (SM1.16).

SM1.7. Computable bounds for suboptimality. In this section, we assume that the linear minimization (SM1.9) is performed exactly at iteration t . That is, there is no error depending on $\|\mathbf{D}_t\|$. Introduce the duality gap surrogate

$$(SM1.31) \quad g_t(\mathbf{X}) := \max_{\mathbf{H} \in \alpha \Delta_n} \langle \nabla f_t(\mathbf{X}), \mathbf{X} - \mathbf{H} \rangle.$$

The function f_t is defined in (SM1.18). Note that the gap $g(\mathbf{X}_t)$ can be evaluated with the information we have at hand:

$$(SM1.32) \quad \begin{aligned} g_t(\mathbf{X}_t) &= \langle \mathbf{D}_t, \mathbf{X}_t \rangle - \langle \mathbf{D}_t, \mathbf{H}_t \rangle \\ &= \langle \mathbf{C}, \mathbf{X}_t \rangle + \langle \mathbf{y}_t + \beta_t(\mathcal{A}\mathbf{X}_t - \mathbf{b}), \mathcal{A}\mathbf{X}_t \rangle - \langle \mathbf{D}_t, \mathbf{H}_t \rangle. \end{aligned}$$

The last term is just the value of the linear minimization (SM1.9)

The gap gives us computable bounds on the suboptimality of the current iterate \mathbf{X}_t . Indeed, the convexity of f_t implies that

$$f_t(\mathbf{X}_t) - f_t(\mathbf{X}_\star) \leq \langle \nabla f_t(\mathbf{X}_t), \mathbf{X}_t - \mathbf{X}_\star \rangle \leq g_t(\mathbf{X}_t).$$

Using the definition (SM1.18) and the fact that \mathbf{X}_\star is feasible for (SM1.1), we find that

$$f_t(\mathbf{X}_t) - p_\star = f_t(\mathbf{X}_t) - f_t(\mathbf{X}_\star) \leq g_t(\mathbf{X}_t).$$

Invoking the definition (SM1.18) again and rearranging, we find that

$$(SM1.33) \quad \langle \mathbf{C}, \mathbf{X}_t \rangle - p_\star \leq g_t(\mathbf{X}_t) - \langle \mathbf{y}_t, \mathcal{A}\mathbf{X}_t - \mathbf{b} \rangle - \frac{1}{2}\beta_t \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|^2.$$

In other words, the suboptimality of the primal iterate \mathbf{X}_t is bounded in terms of the dual gap $g_t(\mathbf{X}_t)$, the feasibility gap $\mathcal{A}\mathbf{X}_t - \mathbf{b}$, and the dual variable \mathbf{y}_t .

Remark SM1.2 (Bounds with approximate linear minimization). We can extend the bound (SM1.33) for the approximate linear minimization in (SM1.9) by taking the error term into account. Based on the definition (SM1.31) of $g_t(\mathbf{X}_t)$ and the oracle (SM1.9),

$$g_t(\mathbf{X}_t) \leq \langle \mathbf{D}_t, \mathbf{X}_t \rangle - \langle \mathbf{D}_t, \mathbf{H}_t \rangle + \frac{\alpha\beta_0}{\beta_t} \|\mathbf{D}_t\|.$$

This leads to an extended version of (SM1.33):

$$(SM1.34) \quad \langle \mathbf{C}, \mathbf{X}_t \rangle - p_\star \leq g_t(\mathbf{X}_t) - \langle \mathbf{y}_t, \mathcal{A}\mathbf{X}_t - \mathbf{b} \rangle - \frac{1}{2}\beta_t \|\mathcal{A}\mathbf{X}_t - \mathbf{b}\|^2 + \frac{\alpha\beta_0}{\beta_t} \|\mathbf{D}_t\|.$$

Note that the additional error term converges to zero. We can simplify this bound further by using (SM1.22).

In practice, we have observed that the bound (SM1.34) is less informative than simply using the bound (SM1.33), which assumes that the eigenvalue computation is exact.

Remark SM1.3 (Superoptimality). Note that the suboptimality $\langle \mathbf{C}, \mathbf{X}_t \rangle - p_\star$ can attain negative values because the CGAL iterates \mathbf{X}_t are generally infeasible. Similarly, the upper bound in (SM1.33) can be negative. In other words, the iterates \mathbf{X}_t can be superoptimal. Nevertheless, the superoptimality is controlled by the feasibility gap, as shown in [subsection SM1.6.6](#).

SM2. Sketching and reconstruction of a positive-semidefinite matrix. This section reviews and gives additional details about the Nyström sketch [SM11, SM9, SM13, SM22]. This sketch tracks an evolving psd matrix and then reports a provably accurate low-rank approximation. The material on error estimation is new.

SM2.1. Sketching and updates. Consider a psd input matrix $\mathbf{X} \in \mathbb{S}_n$. Let R be a parameter that modulates the storage cost of the sketch and the quality of the matrix approximation.

To construct the sketch, we draw and fix a standard normal test matrix $\mathbf{\Omega} \in \mathbb{F}^{n \times R}$. The sketch of the matrix \mathbf{X} takes the form

$$(SM2.1) \quad \mathbf{S} = \mathbf{X}\mathbf{\Omega} \in \mathbb{F}^{n \times R} \quad \text{and} \quad \tau = \text{tr}(\mathbf{X}) \in \mathbb{R}.$$

The sketch supports linear rank-one updates to \mathbf{X} . Indeed, we can track the evolution

$$(SM2.2) \quad \begin{aligned} & \mathbf{X} \leftarrow (1 - \eta) \mathbf{X} + \eta \mathbf{v}\mathbf{v}^* && \text{for } \eta \in [0, 1] \text{ and } \mathbf{v} \in \mathbb{F}^n \\ \text{via } & \mathbf{S} \leftarrow (1 - \eta) \mathbf{S} + \eta \mathbf{v}(\mathbf{v}^*\mathbf{\Omega}) \\ & \text{and } \tau \leftarrow (1 - \eta)\tau + \eta \|\mathbf{v}\|^2. \end{aligned}$$

The test matrix $\mathbf{\Omega}$ and the sketch (\mathbf{S}, τ) require storage of $2Rn + 1$ numbers in \mathbb{F} . The arithmetic cost of the linear update (5.2) to the sketch is $\Theta(Rn)$ numerical operations.

Remark SM2.1 (Structured random matrices). We can reduce storage costs by a factor of two by using a structured random matrix in place of $\mathbf{\Omega}$. For example, see [SM24, Sec. 3] or [SM20]. This modification requires additional care with implementation (e.g., use of sparse arithmetic or fast trigonometric transforms), but the improvement can be significant for very large problems.

SM2.2. The reconstruction process. Given the test matrix $\mathbf{\Omega}$ and the sketch $\mathbf{S} = \mathbf{X}\mathbf{\Omega}$, we can form a rank- R approximation $\widehat{\mathbf{X}}$ of the matrix \mathbf{X} contained in the sketch. This approximation is defined by the formula

$$(SM2.3) \quad \widehat{\mathbf{X}} := \mathbf{S}(\mathbf{\Omega}^*\mathbf{S})^\dagger \mathbf{S}^* = (\mathbf{X}\mathbf{\Omega})(\mathbf{\Omega}^*\mathbf{X}\mathbf{\Omega})^\dagger (\mathbf{X}\mathbf{\Omega})^*.$$

This reconstruction is called a *Nyström approximation*. We often truncate $\widehat{\mathbf{X}}$ by replacing it with its best rank- r approximation $\llbracket \widehat{\mathbf{X}} \rrbracket_r$ for some $r \leq R$.

See [Algorithm SM2.1](#) for a numerically stable implementation of the Nyström reconstruction process (5.3), including error estimation steps. The algorithm takes $O(R^2n)$ numerical operations and $\Theta(Rn)$ storage.

SM2.3. The approximation error. We can easily determine the exact Schatten 1-norm error in the truncated Nyström approximation:

$$(SM2.4) \quad \text{err}(\llbracket \widehat{\mathbf{X}} \rrbracket_r) := \|\mathbf{X} - \llbracket \widehat{\mathbf{X}} \rrbracket_r\|_* = \tau - \text{tr}(\llbracket \widehat{\mathbf{X}} \rrbracket_r) \quad \text{for each } r \leq R.$$

Furthermore, we can use (SM2.4) to ascertain whether the unknown input matrix \mathbf{X} is (almost) low-rank. Indeed, the best rank- r approximation of \mathbf{X} satisfies

$$(SM2.5) \quad \|\mathbf{X} - \llbracket \mathbf{X} \rrbracket_r\|_* \leq \text{err}(\llbracket \widehat{\mathbf{X}} \rrbracket_r) \quad \text{for each } r \leq R.$$

Thus, large drops in the function $r \mapsto \text{err}(\llbracket \widehat{\mathbf{X}} \rrbracket_r)$ signal large drops in the eigenvalues of \mathbf{X} . See [subsection SM2.9](#) for the details.

Algorithm SM2.1 NystromSketch implementation (see [section SM2](#))

Input: Dimension n of input matrix, size R of sketch

Output: Rank- R approximation $\widehat{\mathbf{X}}$ of sketched matrix in factored form $\widehat{\mathbf{X}} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^*$, where $\mathbf{U} \in \mathbb{R}^{n \times R}$ has orthonormal columns and $\mathbf{\Lambda} \in \mathbb{R}^{R \times R}$ is nonnegative diagonal, and the Schatten 1-norm approximation errors $\text{err}(\llbracket \widehat{\mathbf{X}} \rrbracket_r)$ for $1 \leq r \leq R$, as defined in [\(SM2.4\)](#)

Recommendation: Choose R as large as possible

```

1 function NystromSketch.Init( $n, R$ )
2    $\mathbf{\Omega} \leftarrow \text{randn}(n, R)$  ▷ Draw and fix random test matrix
3    $\mathbf{S} \leftarrow \text{zeros}(n, R)$  and  $\tau \leftarrow 0$  ▷ Form sketch of zero matrix

4 function NystromSketch.RankOneUpdate( $\mathbf{v}, \eta$ ) ▷ Implements (5.2)
5    $\mathbf{S} \leftarrow (1 - \eta)\mathbf{S} + \eta\mathbf{v}(\mathbf{v}^*\mathbf{\Omega})$  ▷ Update sketch of matrix
6    $\tau \leftarrow (1 - \eta)\tau + \eta\|\mathbf{v}\|^2$  ▷ Update the trace

7 function NystromSketch.Reconstruct()
8    $\sigma \leftarrow \sqrt{n}\text{eps}(\text{norm}(\mathbf{S}))$  ▷ Compute a shift parameter
9    $\mathbf{S} \leftarrow \mathbf{S} + \sigma\mathbf{\Omega}$  ▷ Implicitly form sketch of  $\mathbf{X} + \sigma\mathbf{I}$ 
10   $\mathbf{L} \leftarrow \text{chol}(\mathbf{\Omega}^*\mathbf{S})$ 
11   $[\mathbf{U}, \mathbf{\Sigma}, \sim] \leftarrow \text{svd}(\mathbf{S}/\mathbf{L})$  ▷ Dense SVD
12   $\mathbf{\Lambda} \leftarrow \max\{0, \mathbf{\Sigma}^2 - \sigma\mathbf{I}\}$  ▷ Remove shift
13   $\text{err} \leftarrow \tau - \text{cumsum}(\text{diag}(\mathbf{\Lambda}))$  ▷ Compute approximation errors
    
```

SM2.4. Trace correction. The trace of the Nyström approximation $\widehat{\mathbf{X}}$ does not exceed the trace of the input matrix \mathbf{X} . It can sometimes be helpful to replace the Nyström approximation by another matrix whose trace matches the trace of the input matrix.

Suppose that the Nyström approximation has the form $\widehat{\mathbf{X}} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^*$, where $\mathbf{U} \in \mathbb{F}^{n \times R}$ is orthonormal and $\mathbf{\Lambda} \in \mathbb{S}_R^+$ is diagonal. Assume that $\text{tr}(\mathbf{X}) = \alpha$ and $\text{tr}(\widehat{\mathbf{X}}) = \widehat{\alpha}$. Since $\widehat{\mathbf{X}}$ is a Nyström approximation of \mathbf{X} , it holds that $\widehat{\alpha} \leq \alpha$. Now, we can construct a second approximation

$$(SM2.6) \quad \widetilde{\mathbf{X}} := \mathbf{U}(\mathbf{\Lambda} + (\alpha - \widehat{\alpha})\mathbf{I}_{R/R})\mathbf{U}^*.$$

It is evident that $\widehat{\mathbf{X}} \preceq \widetilde{\mathbf{X}}$ and that $\text{tr}\widetilde{\mathbf{X}} = \alpha$.

One might hope that the trace-corrected approximation $\widetilde{\mathbf{X}}$ is better than the original Nyström approximation $\widehat{\mathbf{X}}$, but this is not necessarily the case. Fortunately, we always have the relation

$$\|\mathbf{X} - \widetilde{\mathbf{X}}\|_* \leq 2\|\mathbf{X} - \widehat{\mathbf{X}}\|_*.$$

In other words, correcting the trace doubles the error, at worst. See [subsection SM2.10](#) for the proof.

SM2.5. Statistical properties of the Nyström sketch. The truncated Nyström approximation has a number of attractive statistical properties. For a fixed input matrix \mathbf{X} , the expected approximation error $\mathbb{E}_{\mathbf{\Omega}} \|\mathbf{X} - \llbracket \widehat{\mathbf{X}} \rrbracket_r\|_*$ is monotone decreasing in both the sketch size

R and the truncation rank r . Furthermore, if $\text{rank}(\mathbf{X}) = r$ for $r \leq R$, then $\|\mathbf{X} - \llbracket \widehat{\mathbf{X}} \rrbracket_r\|_* = 0$ with probability one. We establish these results below in [subsection SM2.11](#).

SM2.6. A priori error bounds. The Nyström approximation $\widehat{\mathbf{X}}$ yields a provably good estimate for the matrix \mathbf{X} contained in the sketch [[SM22](#), Thms. 4.1].

Fact SM2.2 (Nyström sketch: Error bound). Fix a psd matrix $\mathbf{X} \in \mathbb{S}_n$. Let $\mathbf{S} = \mathbf{X}\mathbf{\Omega}$ where $\mathbf{\Omega} \in \mathbb{F}^{n \times R}$ is standard normal. For each $r < R-1$, the Nyström approximation ([SM2.3](#)) satisfies

$$(SM2.7) \quad \mathbb{E}_{\mathbf{\Omega}} \|\mathbf{X} - \widehat{\mathbf{X}}\|_* \leq \left(1 + \frac{r}{R-r-1}\right) \|\mathbf{X} - \llbracket \mathbf{X} \rrbracket_r\|_*.$$

If we replace $\widehat{\mathbf{X}}$ with the rank- r truncation $\llbracket \widehat{\mathbf{X}} \rrbracket_r$, the error bound ([SM2.7](#)) remains valid. Similar results hold with high probability.

The truncated Nyström approximations satisfy a stronger error bound when the input matrix \mathbf{X} exhibits spectral decay.

Fact SM2.3 (Nyström sketch: Error bound II). Fix a psd matrix $\mathbf{X} \in \mathbb{S}_n$. Let $\mathbf{S} = \mathbf{X}\mathbf{\Omega}$ where $\mathbf{\Omega} \in \mathbb{F}^{n \times R}$ is standard normal. For each $r < R-1$, the Nyström approximation ([SM2.3](#)) satisfies

$$(SM2.8) \quad \mathbb{E}_{\mathbf{\Omega}} \|\mathbf{X} - \llbracket \widehat{\mathbf{X}} \rrbracket_r\|_* \leq \|\mathbf{X} - \llbracket \mathbf{X} \rrbracket_r\|_* + \left(1 + \frac{r}{R-r-1}\right) \|\mathbf{X} - \llbracket \mathbf{X} \rrbracket_r\|_*.$$

If we replace $\widehat{\mathbf{X}}$ with the rank- r truncation $\llbracket \widehat{\mathbf{X}} \rrbracket_r$, the error bound ([SM2.8](#)) remains valid. Similar results hold with high probability.

Proof. Combine the proofs of [[SM22](#), Thm. 4.2] and [[SM11](#), Thm. 9.3]. ■

SM2.7. Discussion. In practice, it is best to minimize the error attributable to sketching. To that end, we recommend choosing the sketch size parameter R as large as possible, given resource constraints, so that we can obtain the highest-quality Nyström approximation.

In some problems, e.g., MaxCut with eigenvector rounding, the desired rank r of the truncation is known in advance. In this case, [Fact 5.2](#) offers guidance about how to select R to achieve a specific error tolerance $(1 + \zeta)$ in ([2.3](#)). For example, when $5r + 1 \leq R$, the expected Schatten 1-norm error in the rank- r approximation $\llbracket \widehat{\mathbf{X}} \rrbracket_r$ is at most $1.25 \times$ the error in the best rank- r approximation of \mathbf{X} .

When the input matrix \mathbf{X} has decaying eigenvalues, the error in the truncated approximation may be far smaller than [Fact 5.2](#) predicts; see [[SM22](#), Thm. 4.2]. This happy situation is typical when \mathbf{X} is generated by the CGAL iteration.

SM2.8. Representation of the truncated Nyström approximation. The key tool in the analysis is a simple representation for the truncated approximation. These facts are extracted from [[SM22](#), Supp.].

Let $\mathbf{X} \in \mathbb{S}_n$ be a fixed psd matrix, and let $\mathbf{\Omega} \in \mathbb{R}^{n \times R}$ be an arbitrary test matrix. Let $\mathbf{P} \in \mathbb{S}_n$ be the orthoprojector onto the range of $\mathbf{X}^{1/2}\mathbf{\Omega}$. Then we can write the Nyström approximation ([SM2.3](#)) as

$$\widehat{\mathbf{X}} = \mathbf{X}^{1/2}\mathbf{P}\mathbf{X}^{1/2}.$$

For each $r \leq R$, define \mathbf{P}_r to be the orthoprojector onto the co-range of the matrix $\llbracket \mathbf{X}^{1/2} \mathbf{P} \rrbracket_r$. By construction, $\llbracket \mathbf{X}^{1/2} \mathbf{P} \rrbracket_r = \mathbf{X}^{1/2} \mathbf{P}_r$. As a consequence,

$$\llbracket \widehat{\mathbf{X}} \rrbracket_r = (\llbracket \mathbf{X}^{1/2} \mathbf{P} \rrbracket_r)(\llbracket \mathbf{X}^{1/2} \mathbf{P} \rrbracket_r)^* = \mathbf{X}^{1/2} \mathbf{P}_r \mathbf{X}^{1/2}.$$

These results allow us to relate the truncated approximations to each other:

SM2.9. Approximation errors: Analysis. We may now obtain explicit formulas for the error in each Nyström approximation. For $r \leq R$, note that

$$\mathbf{X} - \llbracket \widehat{\mathbf{X}} \rrbracket_r = \mathbf{X}^{1/2}(\mathbf{I} - \mathbf{P}_r)\mathbf{X}^{1/2} \succcurlyeq \mathbf{0}.$$

It follows immediately that

$$\text{err}(\llbracket \widehat{\mathbf{X}} \rrbracket_r) = \|\mathbf{X} - \llbracket \widehat{\mathbf{X}} \rrbracket_r\|_* = \text{tr}(\mathbf{X} - \llbracket \widehat{\mathbf{X}} \rrbracket_r) = \text{tr}(\mathbf{X}) - \text{tr}(\llbracket \widehat{\mathbf{X}} \rrbracket_r).$$

This is the relation (SM2.4).

Assume that $r \leq r'$. Since $\text{tr}(\llbracket \widehat{\mathbf{X}} \rrbracket_r) \leq \text{tr}(\llbracket \widehat{\mathbf{X}} \rrbracket_{r'})$, we have the bound

$$\text{err}(\llbracket \widehat{\mathbf{X}} \rrbracket_r) \geq \text{err}(\llbracket \widehat{\mathbf{X}} \rrbracket_{r'}) \quad \text{for } r \leq r'.$$

In other words, for fixed sketch size R , the error in the truncated Nyström approximation is monotone decreasing in the approximation rank.

To obtain (SM2.5), observe that

$$\|\mathbf{X} - \llbracket \mathbf{X} \rrbracket_r\|_* \leq \|\mathbf{X} - \llbracket \widehat{\mathbf{X}} \rrbracket_r\|_* = \text{err}(\llbracket \widehat{\mathbf{X}} \rrbracket_r).$$

The inequality holds because $\llbracket \mathbf{X} \rrbracket_r$ is a best rank- r approximation of \mathbf{X} in Schatten 1-norm, while $\llbracket \widehat{\mathbf{X}} \rrbracket_r$ is another rank- r matrix.

SM2.10. Trace correction: Analysis. Next, we study the trace-corrected Nyström approximation, introduced in subsection SM2.4. Recall that \mathbf{X} is a psd matrix with trace α and $\widehat{\mathbf{X}}$ is a Nyström approximation of \mathbf{X} with $\text{tr}(\widehat{\mathbf{X}}) = \widehat{\alpha}$. Owing to the projection representation of $\widehat{\mathbf{X}}$ in subsection SM2.8, it must be the case that $\widehat{\alpha} \leq \alpha$. The trace-corrected approximation $\widetilde{\mathbf{X}}$ is defined in (SM2.6). By construction, this matrix satisfies $\widetilde{\mathbf{X}} \succcurlyeq \widehat{\mathbf{X}}$ and that $\text{tr}(\widetilde{\mathbf{X}}) = \alpha$.

First, we develop a variational interpretation of the trace-corrected approximation:

$$(SM2.9) \quad \widetilde{\mathbf{X}} \in \arg \min \{ \|\mathbf{Y} - \widehat{\mathbf{X}}\|_* : \text{tr}(\mathbf{Y}) = \alpha \text{ and } \mathbf{Y} \succcurlyeq \mathbf{0} \}.$$

Indeed, for any feasible point \mathbf{Y} ,

$$\|\mathbf{Y} - \widehat{\mathbf{X}}\|_* \geq \text{tr}(\mathbf{Y} - \widehat{\mathbf{X}}) = \alpha - \widehat{\alpha}.$$

In the first relation, equality holds if and only if $\mathbf{Y} \succcurlyeq \widehat{\mathbf{X}}$. We have already seen that the matrix $\widetilde{\mathbf{X}}$ is a feasible point that satisfies the equality condition. There are many solutions to the optimization problem (SM2.9); we have singled out $\widetilde{\mathbf{X}}$ as the one that simultaneously minimizes the Frobenius-norm error $\|\mathbf{Y} - \widehat{\mathbf{X}}\|_F$ over the same feasible set.

Second, we need to argue that trace correction has a controlled impact on the error in the Nyström approximation. This point follows from a standard calculation:

$$\|\widetilde{\mathbf{X}} - \mathbf{X}\|_* \leq \|\widetilde{\mathbf{X}} - \widehat{\mathbf{X}}\|_* + \|\mathbf{X} - \widehat{\mathbf{X}}\|_* \leq 2\|\mathbf{X} - \widehat{\mathbf{X}}\|_*.$$

The first relation is the triangle inequality. The second relation holds because $\widetilde{\mathbf{X}}$ solves the variational problem (SM2.9), while \mathbf{X} is also feasible for this optimization problem.

SM2.11. Statistical properties: Analysis. Next, let us verify the statistical properties of the error. In this section, the test matrix $\mathbf{\Omega} \in \mathbb{F}^{n \times R}$ is standard normal.

Assuming that $\text{rank}(\mathbf{X}) = r \leq R$, let us prove that $\|\mathbf{X} - \llbracket \widehat{\mathbf{X}} \rrbracket_r\|_* = 0$ with probability one. To that end, we observe

$$\text{range}(\mathbf{P}) = \text{range}(\mathbf{X}^{1/2}\mathbf{\Omega}) = \text{range}(\mathbf{X}^{1/2}) \quad \text{with probability one.}$$

It follows that

$$\widehat{\mathbf{X}} = \mathbf{X}^{1/2}\mathbf{P}\mathbf{X}^{1/2} = \mathbf{X} \quad \text{with probability one.}$$

Moreover, $\text{rank}(\mathbf{X}^{1/2}\mathbf{P}) = r$ with probability one. Conditional on this event,

$$\llbracket \widehat{\mathbf{X}} \rrbracket_r = \mathbf{X}^{1/2}\mathbf{P}_r\mathbf{X}^{1/2} = (\llbracket \mathbf{X}^{1/2}\mathbf{P} \rrbracket_r)(\llbracket \mathbf{X}^{1/2}\mathbf{P} \rrbracket_r)^* = (\mathbf{X}^{1/2}\mathbf{P})(\mathbf{X}^{1/2}\mathbf{P})^* = \mathbf{X}.$$

This is the stated result.

Next, we show that the expected error in the truncated Nyström approximation is monotone decreasing with respect to the sketch size R . Fix the truncation rank r . Let $\mathbf{\Omega}_+ = [\mathbf{\Omega} \ \boldsymbol{\omega}]$, where $\boldsymbol{\omega} \in \mathbb{F}^n$ is a standard normal vector independent from $\mathbf{\Omega}$. Define $\mathbf{P}_+ \in \mathbb{S}_n$ to be the orthoprojector onto $\text{range}(\mathbf{X}^{1/2}\mathbf{\Omega}_+)$. It is clear that $\text{range}(\mathbf{P}) \subseteq \text{range}(\mathbf{P}_+)$, and so

$$\mathbf{X}^{1/2}\mathbf{P}\mathbf{X}^{1/2} \preceq \mathbf{X}^{1/2}\mathbf{P}_+\mathbf{X}^{1/2}.$$

As a consequence,

$$\text{tr}(\llbracket \mathbf{X}^{1/2}\mathbf{P}\mathbf{X}^{1/2} \rrbracket_r) \leq \text{tr}(\llbracket \mathbf{X}^{1/2}\mathbf{P}_+\mathbf{X}^{1/2} \rrbracket_r).$$

Equivalently,

$$\|\mathbf{X} - \llbracket \mathbf{X}^{1/2}\mathbf{P}\mathbf{X}^{1/2} \rrbracket_r\|_* \geq \|\mathbf{X} - \llbracket \mathbf{X}^{1/2}\mathbf{P}_+\mathbf{X}^{1/2} \rrbracket_r\|_*.$$

Take the expectation with respect to $\mathbf{\Omega}_+$. The left-hand side is the average error in the r -truncated Nyström approximation with a standard normal sketch of size R . The right-hand side is the same thing, except the sketch has size $R + 1$. This is the required result.

SM3. SketchyCGAL: Additional results. This section contains some additional material about the SketchyCGAL algorithm.

SM3.1. Assessing solution quality. We can develop estimates for the quality of the SketchyCGAL solution by adapted the approach that we used for CGAL.

To do so, we need to track the primal objective value at the current iterate:

$$p_t = \langle \mathbf{C}, \mathbf{X}_t \rangle.$$

At each iteration, we can easily update this estimate using the computed approximate eigenvector \mathbf{v}_t :

$$p_{t+1} = (1 - \eta_t)p_t + \eta_t \alpha \langle \mathbf{v}_t, \mathbf{C}\mathbf{v}_t \rangle.$$

This update rule is applied with the help of the primitive (2.4) \bullet .

When we wish to estimate the error, say in iteration t , we can solve the eigenvalue subproblem to very high accuracy:

$$\xi_t = \mathbf{v}_t^* \mathbf{D}_t \mathbf{v}_t = \min_{\|\mathbf{v}\|=1} \mathbf{v}^* \mathbf{D}_t \mathbf{v}.$$

Then, we can compute the surrogate duality gap:

$$g_t(\mathbf{X}_t) = p_t + \langle \mathbf{y}_t + \beta_t(\mathbf{z}_t - \mathbf{b}), \mathbf{z}_t \rangle - \xi_t.$$

This expression follows directly from the formula (SM1.32) using the loop invariant that $\mathbf{z}_t = \mathcal{A}\mathbf{X}_t$. We arrive at a computable error estimate:

$$p_t - p_\star \leq g_t(\mathbf{X}_t) - \langle \mathbf{y}_t, \mathbf{z}_t - \mathbf{b} \rangle - \frac{1}{2}\beta_t\|\mathbf{z}_t - \mathbf{b}\|^2.$$

This bound follows directly from (SM1.33).

SM3.2. Convergence theory. In this section, we establish two simple results on the convergence properties of the SketchyCGAL algorithm.

Theorem SM3.1 (SketchyCGAL: Convergence I). *Let Ψ_\star be the solution set of the model problem (2.2). For each $r < R - 1$, the iterates $\widehat{\mathbf{X}}_t$ computed by SketchyCGAL (subsections 6.1 and 6.2) satisfy*

$$\limsup_{t \rightarrow \infty} \mathbb{E}_\Omega \text{dist}_*(\widehat{\mathbf{X}}_t, \Psi_\star) \leq \left(1 + \frac{r}{R - r - 1}\right) \cdot \max_{\mathbf{X} \in \Psi_\star} \|\mathbf{X} - \llbracket \mathbf{X} \rrbracket_r\|_*.$$

The same bound holds for the truncated approximations $\llbracket \widehat{\mathbf{X}}_t \rrbracket_r$. Here, dist_* is the nuclear-norm distance between a matrix and a set of matrices.

Proof. The implicit iterates \mathbf{X}_t satisfy the conclusions of Fact 3.1, so they converge toward the compact set Ψ_\star . Therefore, we can choose a sequence $\{\mathbf{X}_{t_\star}\} \subset \Psi_\star$ with the property that $\|\mathbf{X}_t - \mathbf{X}_{t_\star}\|_* \rightarrow 0$. By the triangle inequality and (6.7),

$$\begin{aligned} \mathbb{E}_\Omega \text{dist}_*(\widehat{\mathbf{X}}_t, \Psi_\star) &\leq \mathbb{E}_\Omega \|\widehat{\mathbf{X}}_t - \mathbf{X}_t\|_* + \text{dist}_*(\mathbf{X}_t, \Psi_\star) \\ &\leq \left(1 + \frac{r}{R - r - 1}\right) \cdot \|\mathbf{X}_t - \llbracket \mathbf{X}_t \rrbracket_r\|_* + \|\mathbf{X}_t - \mathbf{X}_{t_\star}\|_*. \end{aligned}$$

The rank- r approximation error in Schatten 1-norm is 1-Lipschitz with respect to the Schatten 1-norm (cf. [SM24, Sec. SM2.2]), so

$$\left| \|\mathbf{X}_t - \llbracket \mathbf{X}_t \rrbracket_r\|_* - \|\mathbf{X}_{t_\star} - \llbracket \mathbf{X}_{t_\star} \rrbracket_r\|_* \right| \leq \|\mathbf{X}_t - \mathbf{X}_{t_\star}\|_*.$$

Therefore,

$$\begin{aligned} \|\mathbf{X}_t - \llbracket \mathbf{X}_t \rrbracket_r\|_* &\leq \|\mathbf{X}_{t_\star} - \llbracket \mathbf{X}_{t_\star} \rrbracket_r\|_* + \|\mathbf{X}_t - \mathbf{X}_{t_\star}\|_* \\ &\leq \max_{\mathbf{X} \in \Psi_\star} \|\mathbf{X} - \llbracket \mathbf{X} \rrbracket_r\|_* + \|\mathbf{X}_t - \mathbf{X}_{t_\star}\|_*. \end{aligned}$$

Combine the last two displays, and extract the superior limit. ■

If the implicit iterates generated by SketchyCGAL happen to converge to a limit, we have a more precise result.

Theorem SM3.2 (SketchyCGAL: Convergence II). *Assume the implicit iterates \mathbf{X}_t induced by SketchyCGAL (subsections 6.1 and 6.2) converge to a matrix \mathbf{X}_{cgal} . For each $r < R - 1$, the computed iterates $\widehat{\mathbf{X}}_t$ satisfy*

$$\begin{aligned} \limsup_{t \rightarrow \infty} \mathbb{E}_{\Omega} \|\mathcal{A}\widehat{\mathbf{X}}_t - \mathbf{b}\| &\leq \left(1 + \frac{r}{R - r - 1}\right) \cdot \|\mathcal{A}\| \cdot \|\mathbf{X}_{\text{cgal}} - \llbracket \mathbf{X}_{\text{cgal}} \rrbracket_r\|_*; \\ \limsup_{t \rightarrow \infty} \mathbb{E}_{\Omega} |\langle \mathbf{C}, \widehat{\mathbf{X}}_t \rangle - \langle \mathbf{C}, \mathbf{X}_* \rangle| &\leq \left(1 + \frac{r}{R - r - 1}\right) \cdot \|\mathbf{C}\| \cdot \|\mathbf{X}_{\text{cgal}} - \llbracket \mathbf{X}_{\text{cgal}} \rrbracket_r\|_*. \end{aligned}$$

The same bound holds for the truncated approximations $\llbracket \widehat{\mathbf{X}}_t \rrbracket_r$. If $\text{rank}(\mathbf{X}_{\text{cgal}}) \leq R$, then the computed iterates $\widehat{\mathbf{X}}_t$ converge to the solution set of (2.2).

Proof. The implicit iterates \mathbf{X}_t satisfy the conclusions of Fact 3.1, so the limit \mathbf{X}_{cgal} solves (2.2). Using the triangle inequality, the operator norm bound, and (6.7), we obtain nonasymptotic error bounds

$$\begin{aligned} \mathbb{E}_{\Omega} \|\mathcal{A}\widehat{\mathbf{X}}_t - \mathbf{b}\| &\leq \frac{\text{Const}}{\sqrt{t}} + \left(1 + \frac{r}{R - r - 1}\right) \cdot \|\mathcal{A}\| \cdot \|\mathbf{X}_t - \llbracket \mathbf{X}_t \rrbracket_r\|_*; \\ \mathbb{E}_{\Omega} |\langle \mathbf{C}, \widehat{\mathbf{X}}_t \rangle - \langle \mathbf{C}, \mathbf{X}_* \rangle| &\leq \frac{\text{Const}}{\sqrt{t}} + \left(1 + \frac{r}{R - r - 1}\right) \cdot \|\mathbf{C}\| \cdot \|\mathbf{X}_t - \llbracket \mathbf{X}_t \rrbracket_r\|_*. \end{aligned}$$

Extract the limit as $t \rightarrow \infty$. The last conclusion follows from the facts outlined in subsection SM2.6. ■

SM4. Beyond the model problem. The CGAL algorithm [SM27] applies to a more general problem template than (2.2). Likewise, the SketchyCGAL algorithm can solve a wider class of problems in a scalable fashion. This section outlines some of the opportunities.

SM4.1. A more general template. Consider the optimization problem

$$(SM4.1) \quad \text{minimize } \langle \mathbf{C}, \mathbf{X} \rangle \quad \text{subject to } \mathcal{A}\mathbf{X} \in \mathbf{K} \quad \text{and } \mathbf{X} \in \mathbf{X}, \quad \mathbf{X} \text{ is psd.}$$

In this expression, $\mathbf{K} \subset \mathbb{R}^d$ is a closed, convex set and $\mathbf{X} \subset \mathbb{S}_n(\mathbb{F})$ is a compact, convex set of matrices. The rest of this section describes some problems that fall within the compass of (SM4.1), as well as new computational challenges that appear. Algorithm SM4.1 contains pseudocode for a version of SketchyCGAL tailored to (SM4.1).

Remark SM4.1 (Other matrix optimization problems). We can also extend SketchyCGAL to optimization problems involving matrices that are symmetric (but not psd) or that are rectangular. For example, matrix completion via nuclear-norm minimization [SM17] falls in this framework. In this case, we need to replace the Nyström sketch with a more general technique, such as [SM23, SM24]. Further extensions are also possible; see [SM27]. We omit these developments.

SM4.2. The convex constraint set. To handle the convex constraint \mathbf{X} that appears in (SM4.1), we must develop a subroutine for the linear minimization problem

$$(SM4.2) \quad \text{minimize}_{\mathbf{H} \in \mathbb{S}_n} \langle \mathbf{D}_t, \mathbf{H} \rangle \quad \text{subject to } \mathbf{H} \in \mathbf{X}, \quad \mathbf{H} \text{ is psd.}$$

To implement SketchyCGAL efficiently, we need the problem (SM4.2) to admit a structured (e.g., low-rank) approximate solution. Here are some situations where this is possible.

1. **Trace-bounded psd matrices.** $\mathcal{X} := \{\mathbf{X} \in \mathbb{S}_n : \text{tr } \mathbf{X} \leq \alpha \text{ and } \mathbf{X} \text{ is psd}\}$. For solving a standard-form SDP, this constraint is more natural than $\mathbf{X} = \alpha \mathbf{\Delta}_n$. Given an (approximate) minimum eigenpair (ξ_t, \mathbf{v}_t) of \mathbf{D}_t , the solution of (SM4.2) is

$$\mathbf{H}_t = \begin{cases} \alpha \mathbf{v}_t \mathbf{v}_t^*, & \xi_t < 0, \\ \mathbf{0}, & \xi_t \geq 0. \end{cases}$$

As before, we can solve the eigenvector problem with Algorithm 4.1.

2. **Relaxed orthoprojectors.** $\mathcal{X} := \{\mathbf{X} \in \mathbb{S}_n : \text{tr } \mathbf{X} = \alpha \text{ and } \mathbf{0} \preceq \mathbf{X} \preceq \mathbf{I}\}$. This is the best convex relaxation of the set of orthogonal projectors with rank α ; see [SM16]. When α is small, we can provably solve the linear minimization with randomized subspace iteration [SM11] or randomized block Lanczos methods [SM10, Sec. 10.3.6].

Remark SM4.2 (Standard-form SDP). It is often easy to find an upper bound for the trace of a solution for an SDP. In many applications, the constraint $\mathcal{A}\mathbf{X} = \mathbf{b}$ already enforces a constant trace (e.g., MaxCut or QAP). In many other applications (e.g., phase retrieval), \mathbf{C} is identity matrix, so the objective is to minimize the trace. In this setting, the trace of an arbitrary feasible point is an upper bound for α . If neither situation is in force, we can still solve a standard-form SDP by solving a small number of trace-bounded SDPs:

1. Start with an arbitrary bound $\alpha_0 > 0$, say, twice the trace of an arbitrary feasible point.
2. Solve the trace-bounded SDP with $\text{tr } \mathbf{X} \leq \alpha_i$ to obtain $\mathbf{X}_*(\alpha_i)$.
3. If the primal objective $\langle \mathbf{C}, \mathbf{X}_*(\alpha_i) \rangle = \langle \mathbf{C}, \mathbf{X}_*(\alpha_{i-1}) \rangle$, then terminate and return $\mathbf{X}_*(\alpha_i)$.
4. Otherwise, set $\alpha_{i+1} = 2\alpha_i$ and return to Step 2.

This procedure terminates in $\tilde{O}(1)$ iterations if a bounded solution exists. It is easy to show by contradiction that there exists no finite $\alpha > \alpha_i$ such that $\langle \mathbf{C}, \mathbf{X}_*(\alpha) \rangle < \langle \mathbf{C}, \mathbf{X}_*(\alpha_i) \rangle$. Unfortunately, it is nontrivial to extend this claim to the approximate solutions because they are infeasible.

SM4.3. Convex inclusions. To handle the inclusion in \mathbf{K} that appears in (SM4.1), we need an efficient algorithm to perform the Euclidean projection onto \mathbf{K} . That is,

$$\text{proj}_{\mathbf{K}}(\mathbf{w}) := \arg \min \{\|\mathbf{w} - \mathbf{u}\| : \mathbf{u} \in \mathbf{K}\} \quad \text{for } \mathbf{w} \in \mathbb{R}^d.$$

Here are some important examples:

1. **Inequality constraints.** $\mathbf{K} := \{\mathbf{u} \in \mathbb{R}^d : \mathbf{u} \leq \mathbf{b}\}$. In this case, the projection takes the form $\text{proj}_{\mathbf{K}}(\mathbf{w}) = (\mathbf{w} - \mathbf{b})_-$, where $(\cdot)_-$ reports the negative part of a vector.
2. **Norm constraints.** $\mathbf{K} := \{\mathbf{u} \in \mathbb{R}^d : \|\mathbf{u} - \mathbf{b}\| \leq \delta\}$, where $\|\cdot\|$ is a norm. The projector can be computed easily for many norms, including the ℓ_p norm for $p \in \{1, 2, \infty\}$.

SM4.4. The CGAL iteration for the general template. To extend the description of the CGAL iteration in subsection SM1.3 for the general template (SM4.1), we consider the following

augmented Lagrangian formulation with the slack variable $\mathbf{w} \in \mathbb{K}$ instead of (SM1.6):

$$L_t(\mathbf{X}; \mathbf{y}) := \langle \mathbf{C}, \mathbf{X} \rangle + \min_{\mathbf{w} \in \mathbb{K}} \left\{ \langle \mathbf{y}, \mathcal{A}\mathbf{X} - \mathbf{w} \rangle + \frac{1}{2} \beta_t \|\mathcal{A}\mathbf{X} - \mathbf{w}\|^2 \right\}.$$

Accordingly, the partial derivative (SM1.8) becomes

$$\mathbf{D}_t := \partial_{\mathbf{X}} L_t(\mathbf{X}_t; \mathbf{y}_t) = \mathbf{C} + \mathcal{A}^*(\mathbf{y}_t + \beta_t(\mathcal{A}\mathbf{X}_t - \mathbf{w}_t)) \quad \text{where} \quad \mathbf{w}_t := \text{proj}_{\mathbb{K}}(\mathcal{A}\mathbf{X}_t + \beta_t^{-1}\mathbf{y}_t).$$

We replace the linear minimization subroutine (SM1.9) with (SM4.2). We can still use an inexact variant of (SM4.2) with additive error. We also revise the dual update scheme by modifying (SM1.11) as

$$\mathbf{y}_{t+1} = \mathbf{y}_t + \gamma_t(\mathcal{A}\mathbf{X}_{t+1} - \bar{\mathbf{w}}_t) \quad \text{where} \quad \bar{\mathbf{w}}_t := \text{proj}_{\mathbb{K}}(\mathcal{A}\mathbf{X}_{t+1} + \beta_{t+1}^{-1}\mathbf{y}_t).$$

Finally, we replace the dual step size parameter selection rule (SM1.12) with

$$(SM4.3) \quad \gamma_t \|\mathcal{A}\mathbf{X}_{t+1} - \bar{\mathbf{w}}_t\|^2 \leq \beta_t \eta_t^2 \alpha^2 \|\mathcal{A}\|^2.$$

The bounded travel condition (SM1.13) remains the same.

To obtain the extension of SketchyCGAL to the general template, we simply pursue the same program outlined in section 6 to augment CGAL with sketching.

SM5. Details of MaxCut experiments. This section presents further details about the termination criteria that we used in the MaxCut experiments presented in section 7.

SM5.1. Comparison of SDP solvers. We begin with a discussion of how we compared the performance of different SDP solvers for the MaxCut problem (1.3).

SM5.1.1. Convention for sign of the dual. We used the sign convention from the Lagrangian formulation (SM1.2) when describing the quality guarantees of each solver, and also in the definition of DIMACS errors and the dual problem in subsection SM5.3. This convention can be different in other works, so these definitions might have the sign of the dual variable inverted in some references.

SM5.1.2. Details for SketchyCGAL. We implement the stopping criteria based on the discussion in subsections SM1.7 and SM3.1. We stop the algorithm when both

$$(SM5.1) \quad \frac{p_t + \langle \mathbf{y}_t, \mathbf{b} \rangle + \frac{1}{2} \beta_t \langle \mathbf{z}_t - \mathbf{b}, \mathbf{z}_t + \mathbf{b} \rangle - \lambda_{\min}(\mathbf{D}_t)}{1 + |p_t|} \leq \varepsilon \quad \text{and} \quad \frac{\|\mathbf{z}_t - \mathbf{b}\|}{1 + \|\mathbf{b}\|} \leq \varepsilon.$$

In theory, this bound requires computing $\lambda_{\min}(\mathbf{D}_t)$ to high accuracy. In practice, we did not observe a significant difference when using a high accuracy approximation or the inexact computation from step (SM1.9). We fixed the code to use the latter to avoid additional cost.

If SketchyCGAL terminates by (SM5.1), then the implicit variable \mathbf{X} provably satisfies

$$\frac{\langle \mathbf{C}, \mathbf{X} \rangle - \langle \mathbf{C}, \mathbf{X}_* \rangle}{1 + |\langle \mathbf{C}, \mathbf{X} \rangle|} \leq \varepsilon \quad \text{and} \quad \frac{\|\mathcal{A}\mathbf{X} - \mathbf{b}\|}{1 + \|\mathbf{b}\|} \leq \varepsilon.$$

This is the guarantee that we seek.

Algorithm SM4.1 SketchyCGAL for the general template (SM4.1)

Input: Problem data for (SM4.1) implemented via the primitives (2.4), sketch size R , number T of iterations

Output: Rank- R approximate solution to (SM4.1) in factored form $\widehat{\mathbf{X}}_T = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^*$ where $\mathbf{U} \in \mathbb{R}^{n \times R}$ has orthonormal columns and $\mathbf{\Lambda} \in \mathbb{R}^{R \times R}$ is nonnegative diagonal, and the Schatten 1-norm approximation errors $\text{err}(\|\widehat{\mathbf{X}}\|_r)$ for $1 \leq r \leq R$, as defined in (SM2.4)

Recommendation: To achieve (2.3), set R as large as possible, and set $T \approx \varepsilon^{-1}$

```

1 function SketchyCGAL( $R; T$ )
2   Scale problem data (subsection 7.1.1) ▷ [opt] Recommended!
3    $\beta_0 \leftarrow 1$  and  $K \leftarrow +\infty$  ▷ Default parameters
4   NystromSketch.Init( $n, R$ )
5    $\mathbf{z} \leftarrow \mathbf{0}_d$  and  $\mathbf{y} \leftarrow \mathbf{0}_d$ 
6   for  $t \leftarrow 1, 2, 3, \dots, T$  do
7      $\beta \leftarrow \beta_0 \sqrt{t+1}$  and  $\eta \leftarrow 2/(t+1)$ 
8      $\mathbf{w} \leftarrow \text{proj}_K(\mathbf{z} + \beta^{-1}\mathbf{y})$ 
9      $\mathbf{D} \leftarrow \mathbf{C} + \mathcal{A}^*(\mathbf{y} + \beta(\mathbf{z} - \mathbf{w}))$  ▷ Represent via primitives (2.4) 1 2
10     $\mathbf{H}$  is a (low-rank) matrix that solves (SM4.2)
11     $\mathbf{z} \leftarrow (1 - \eta)\mathbf{z} + \eta\mathcal{A}(\mathbf{H})$  ▷ Use primitive (2.4) 3
12     $\beta_+ \leftarrow \beta_0 \sqrt{t+2}$ 
13     $\mathbf{w} \leftarrow \text{proj}_K(\mathbf{z} + \beta_+^{-1}\mathbf{y})$ 
14     $\mathbf{y} \leftarrow \mathbf{y} + \gamma(\mathbf{z} - \mathbf{w})$  ▷ Step size  $\gamma$  satisfies (SM1.13) and (SM4.3)
15    NystromSketch.RankOneUpdate( $\sqrt{\alpha}\mathbf{v}, \eta$ )
16    [ $\mathbf{U}, \mathbf{\Lambda}, \text{err}$ ]  $\leftarrow$  NystromSketch.Reconstruct()
    
```

SM5.1.3. Details for SDPT3. We used SDPT3 version 4.0 [SM21] in the experiments. This software is designed for solving conic optimization problems by using a primal-dual interior-point algorithm. The algorithm iterates three variables: \mathbf{X} , \mathbf{y} , and \mathbf{Z} . For the MaxCut problem, \mathbf{X} and \mathbf{Z} are $n \times n$ symmetric positive semidefinite matrices, and $\mathbf{y} \in \mathbb{R}^n$.

We can control the desired accuracy by changing the parameter `OPTIONS.gaptol`. When we set this parameter to ε , the outputs ensure

$$\frac{\langle \mathbf{X}, \mathbf{Z} \rangle}{1 + |\langle \mathbf{C}, \mathbf{X} \rangle| + |\mathbf{b}^\top \mathbf{y}|} \leq \varepsilon, \quad \frac{\|\mathcal{A}\mathbf{X} - \mathbf{b}\|}{1 + \|\mathbf{b}\|} \leq \varepsilon, \quad \text{and} \quad \frac{\|\mathcal{A}^*\mathbf{y} - \mathbf{Z} + \mathbf{C}\|_F}{1 + \|\mathbf{C}\|_F} \leq \varepsilon.$$

These are the required bounds.

SM5.1.4. Details for SeDuMi. We used SeDuMi 1.3 [SM18] in the experiments. This software implements a self-dual embedding technique [SM26]. The algorithm has two outputs. For the MaxCut problem, they are the $n \times n$ symmetric positive semidefinite primal solution \mathbf{X} and n dimensional dual solution \mathbf{y} . We can control the desired accuracy by changing the parameter `pars.eps`.

For some instances, when we set `pars.eps` very small, even though the algorithm achieves the desired accuracy, we observed that SeDuMi 1.3 overwrites the variables as zeros after

solving the problem before returning them, based on a quality control procedure. To prevent this issue, we commented out the section between the lines 638 and 642 in `sedumi.m`.

SM5.1.5. Details for SDPNAL+. We used SDPNAL version 1.0 [SM19] which implements an augmented Lagrangian based method for solving semidefinite programs. When applied to the MAXCUT SDP, the algorithm outputs three variables \mathbf{X} , \mathbf{y} and \mathbf{Z} similar to SDPT3. We can control the desired accuracy by changing the parameter `OPTIONS.tol`. When we set this parameter to ε , the outputs ensure

$$\frac{\|\mathcal{A}\mathbf{X} - \mathbf{b}\|}{1 + \|\mathbf{b}\|} \leq \varepsilon, \quad \frac{\|\mathcal{A}^*\mathbf{y} - \mathbf{Z} + \mathbf{C}\|_F}{1 + \|\mathbf{C}\|_F} \leq \varepsilon, \quad \text{and} \quad \frac{\|\mathbf{X} - \text{proj}_{\text{psd}}(\mathbf{X} - \mathbf{Z})\|_F}{1 + \|\mathbf{X}\|_F + \|\mathbf{Z}\|_F} \leq 5\varepsilon,$$

where $\text{proj}_{\text{psd}} : \mathbb{R}^{n \times n} \rightarrow \mathbb{S}_n$ is the projection operator onto positive semidefinite cone.

SM5.1.6. Details for Mosek. We used the interior point optimizer for conic optimization from the Mosek Optimization Suite Release 8.1.0.64 [SM15]. This optimizer implements of the so-called homogeneous and self-dual algorithm [SM2]. When applied to the MaxCut SDP, the algorithm returns three variables \mathbf{X} , \mathbf{y} , and \mathbf{Z} . We can control the target accuracy by changing three parameters: `MSK_DPAR_INTPNT_CO_TOL_PFEAS`, `MSK_DPAR_INTPNT_CO_TOL_DFEAS`, and `MSK_DPAR_INTPNT_CO_TOL_RELGAP`. For simplicity, and as suggested in the manual, we relax these parameters together and set each one to ε . This ensures

$$\frac{\|\mathcal{A}\mathbf{X} - \mathbf{b}\|_\infty}{1 + \|\mathbf{b}\|_\infty} \leq \varepsilon, \quad \frac{\|\mathcal{A}^*\mathbf{y} - \mathbf{Z} + \mathbf{C}\|_\infty}{1 + \|\mathbf{C}\|_\infty} \leq \varepsilon,$$

$$\frac{\langle \mathbf{X}, \mathbf{Z} \rangle}{\max\{1, \min\{|\langle \mathbf{C}, \mathbf{X} \rangle|, |\mathbf{b}^\top \mathbf{y}|\}\}} \leq \varepsilon, \quad \text{and} \quad \frac{|\langle \mathbf{C}, \mathbf{X} \rangle + \mathbf{b}^\top \mathbf{y}|}{\max\{1, \min\{|\langle \mathbf{C}, \mathbf{X} \rangle|, |\mathbf{b}^\top \mathbf{y}|\}\}} \leq \varepsilon.$$

These are the quality guarantees that we need.

SM5.2. Convergence of SketchyCGAL for the MaxCut SDP. This section gives further information about our evaluation of the convergence behavior of the SketchyCGAL algorithm.

Table SM2 presents numerical data from the MaxCut SDP experiment with GSET benchmark. We compare the methods in terms of the

$$\text{suboptimality} = \frac{\langle \mathbf{C}, \mathbf{X} \rangle - \langle \mathbf{C}, \mathbf{X}_\star \rangle}{1 + |\langle \mathbf{C}, \mathbf{X}_\star \rangle|} \quad \text{and} \quad \text{infeasibility} = \frac{\|\mathcal{A}\mathbf{X} - \mathbf{b}\|}{1 + \|\mathbf{b}\|},$$

as well as the storage cost and computation time. We use a high-accuracy solution from SDPT3 with default parameters to approximate the optimal point. The storage cost is approximated by monitoring the virtual memory size of the process, hence it includes the memory that is swapped out and it can go beyond 16 GB.

We also compare the weight of cut evaluated after rounding (see subsection 7.2.1 for details of the rounding procedure), relative to the weight obtained by rounding \mathbf{X}_\star :

$$\text{relative cut weight} = \frac{\text{cut weight from } \mathbf{X} - \text{cut weight from } \mathbf{X}_\star}{\text{cut weight from } \mathbf{X}_\star}.$$

Positive values indicate better cuts with a higher weight. The average of this quantity over all datasets in the benchmark indicates a discrepancy of less than 1.5%, a small price for huge scalability benefits.

SM5.3. Primal–dual convergence. This section presents empirical evidence for the convergence of the primal variables, the dual variables, and the surrogate duality gap generated by SketchyCGAL. We use the MaxCut SDP (1.3) for these tests.

SM5.3.1. DIMACS errors. We evaluate the DIMACS errors [SM1] (with Euclidean scaling) to measure the suboptimality and infeasibilities for primal and dual problems. These measures are commonly used for benchmarking SDP solvers [SM14]. For a given approximate solution triplet $(\mathbf{X}, \mathbf{y}, \mathbf{Z})$, we compute

$$\begin{aligned} \text{err}_1 &= \frac{\|\mathcal{A}\mathbf{X} - \mathbf{b}\|}{1 + \|\mathbf{b}\|}, & \text{err}_2 &= \frac{\max\{-\lambda_{\min}(\mathbf{X}), 0\}}{1 + \|\mathbf{b}\|}, & \text{err}_3 &= \frac{\|\mathcal{A}^*\mathbf{y} - \mathbf{Z} + \mathbf{C}\|_F}{1 + \|\mathbf{C}\|_F}, \\ \text{err}_4 &= \frac{\max\{-\lambda_{\min}(\mathbf{Z}), 0\}}{1 + \|\mathbf{C}\|_F}, & \text{and} & & \text{err}_5 &= \frac{\langle \mathbf{C}, \mathbf{X} \rangle + \mathbf{b}^\top \mathbf{y}}{1 + |\langle \mathbf{C}, \mathbf{X} \rangle| + |\mathbf{b}^\top \mathbf{y}|}. \end{aligned}$$

We also compute the 6th error defined in [SM14] to measure the violation in complementary slackness

$$\text{err}_6 = \frac{\langle \mathbf{X}, \mathbf{Z} \rangle}{1 + |\langle \mathbf{C}, \mathbf{X} \rangle| + |\mathbf{b}^\top \mathbf{y}|}.$$

SketchyCGAL and Sedumi do not maintain the third variable \mathbf{Z} explicitly. We set $\mathbf{Z} = \mathbf{C} + \mathcal{A}^*\mathbf{y}$ as suggested in the guideline of the 7th DIMACS Implementation Challenge [SM1, SM14].

We use a testbed similar to the MaxCut experiments summarized in subsection 7.2. For ease of comparison, we replace $\text{tr } \mathbf{X} = n$ constraint with $\text{tr } \mathbf{X} \leq 1.05 \cdot n$. We explain the reason of this modification in the next subsection. We also set a stronger target accuracy by choosing $\varepsilon = 10^{-3}$ for each solver. Table SM7 reports the numerical outcomes.

SM5.3.2. The dual of the model problem. The dual function for the model problem (SM1.1) with the trace constraint is

$$\begin{aligned} \text{(SM5.2)} \quad \phi(\mathbf{y}) &:= \min_{\mathbf{X} \in \alpha \Delta_n} L(\mathbf{X}, \mathbf{y}) = \min_{\mathbf{X} \in \alpha \Delta_n} \langle \mathbf{C} + \mathcal{A}^*\mathbf{y}, \mathbf{X} \rangle - \langle \mathbf{y}, \mathbf{b} \rangle \\ &= \alpha \cdot \lambda_{\min}(\mathbf{C} + \mathcal{A}^*\mathbf{y}) - \langle \mathbf{y}, \mathbf{b} \rangle. \end{aligned}$$

Therefore, the dual problem is an unconstrained nonsmooth concave maximization:

$$\text{(SM5.3)} \quad \underset{\mathbf{y} \in \mathbb{R}^n}{\text{maximize}} \quad \phi(\mathbf{y}) := \alpha \cdot \lambda_{\min}(\mathbf{C} + \mathcal{A}^*\mathbf{y}) - \langle \mathbf{y}, \mathbf{b} \rangle.$$

Under strong duality, $\phi(\mathbf{y}_\star) = p_\star = \langle \mathbf{C}, \mathbf{X}_\star \rangle$.

Note that (SM5.3) is slightly different from the dual of the standard-form SDP. In the standard-form dual, the maximization in (SM5.2) takes place over the psd cone without any further constraints, and the term $\min_{\{\mathbf{X} \text{ is psd}\}} \langle \mathbf{C} + \mathcal{A}^*\mathbf{y}, \mathbf{X} \rangle$ becomes an indicator function that constraints $\mathbf{C} + \mathcal{A}^*\mathbf{y}$ to the psd cone. We can formulate the standard-form dual problem as

$$\text{(SM5.4)} \quad \underset{\mathbf{y} \in \mathbb{R}^n}{\text{maximize}} \quad -\langle \mathbf{y}, \mathbf{b} \rangle \quad \text{subject to} \quad \mathbf{C} + \mathcal{A}^*\mathbf{y} \text{ is psd},$$

In particular, the problem (SM5.4) is constrained, while (SM5.3) is unconstrained.

DIMACS measures are defined to evaluate convergence to a solution of (SM5.4), hence they do not fit well for evaluating (SM5.3) in general. In our experiments, we empirically observed that $\phi(\mathbf{y})$ converges to p_* for the dual sequence of SketchyCGAL, but \mathbf{y} does not converge to a solution of (SM5.4). However, the following lemma from our concurrent work with Ding [SM8] shows that the solution sets of the duals of the standard SDP and the trace-bounded SDP (subsection SM4.2) are the same under some technical conditions.

Lemma SM5.1 (Lemma 6.1 in [SM8]). *Suppose that the standard-form SDP has a solution \mathbf{X}_* , satisfies strong duality, and \mathbf{y}_* is the unique solution of the standard-form dual problem (SM5.4). Consider the trace bounded SDP with $\mathcal{X} = \{\mathbf{X} \in \mathbb{S}_n : \text{tr } \mathbf{X} \leq \alpha \text{ and } \mathbf{X} \text{ is psd}\}$ for some $\alpha > \text{tr } \mathbf{X}_*$. The dual problem is*

$$(SM5.5) \quad \underset{\mathbf{y} \in \mathbb{R}^n}{\text{maximize}} \quad \alpha \cdot \min\{0, \lambda_{\min}(\mathbf{C} + \mathcal{A}^* \mathbf{y})\} - \langle \mathbf{y}, \mathbf{b} \rangle.$$

Suppose $\|\mathbf{b}\| \neq 0$. Then, \mathbf{y}_* is the unique solution of (SM5.5).

We refer to [SM8] for the proof.

SM5.4. Failure of the Burer–Monteiro heuristic. This section presents empirical evidence that Burer–Monteiro (BM) factorization methods cannot support storage costs better than $\Omega(n\sqrt{d})$. Our approach is based on the paper of Waldspurger and Waters [SM25], which proves that the BM heuristic (8.1) can produce incorrect results unless $R = \Omega(n\sqrt{d})$.

Waldspurger provided us code that generates a random symmetric $\mathbf{C} \in \mathbb{S}_n(\mathbb{R})$. The MaxCut SDP (1.3) with objective \mathbf{C} has a unique solution, and the solution has rank 1. If the factorization rank R satisfies $R(R+3) \leq 2n$, then the BM formulation (8.1) has second-order critical points that are not optimal points of the original SDP (1.3). As a consequence, the Burer–Monteiro approach is reliable only if the storage budget is $\Theta(n^{3/2})$. In contrast, for the same problem instances, our analysis (Theorem Theorem 6.3) shows that SketchyCGAL succeeds with factorization rank $R = 2$ and storage budget $\Theta(n)$.

We will demonstrate numerically that Waldspurger and Waters [SM25] have identified a serious obstruction to using the Burer–Monteiro approach. Moreover, we will see that SketchyCGAL resolves the issue. See the code supplement for scripts to reproduce these experiments.

We use the Manopt software [SM3] to solve the Burer–Monteiro formulation of the MaxCut SDP. For $n = 100$, we drew 10 random matrices $\mathbf{C}_1, \dots, \mathbf{C}_{10}$ using Waldspurger’s code. For each instance, we sweep the factorization rank $R = 2, 3, 4, \dots, 13$. (For $R \geq 13$, we anticipate that each second-order critical point of the Burer–Monteiro problem is a solution to the original SDP, owing to the analysis in [SM4].) In each experiment, we ran Manopt with 100 random initializations, and we counted the number of times the algorithm failed. We declared failure if Manopt converged to a second-order stationary point whose objective value is 10^{-3} larger than the true optimal value. See Table SM1 for the statistics.

In contrast, SketchyCGAL can solve all of these instances, even when the sketch size $R = 2$. For these problems, we use the default parameter choices for SketchyCGAL, but we do not pre-scale the data or perform tuning. Figure SM7.1 compares the convergence trajectory of SketchyCGAL and Manopt for one problem instance. The difference is evident.

SM6. Details of phase retrieval experiments. This section presents further details about the phase retrieval experiments presented in [section 7](#).

SM6.1. Synthetic phase retrieval data. This section provides additional details on the construction of synthetic datasets for the abstract phase retrieval SDP.

For each $n \in \{10^2, 10^3, \dots, 10^6\}$, we generate 20 independent datasets as follows. First, draw $\chi_{\natural} \in \mathbb{C}^n$ from the complex standard normal distribution. We acquire $d = 12n$ phaseless measurements (7.2) using the coded diffraction pattern model [SM7].

To do so, we randomly draw 12 independent modulating waveforms ψ_j for $j = 1, 2, \dots, 12$. Each entry of ψ_j is drawn as the product of two independent random variables, one chosen uniformly from $\{1, i, -1, -i\}$, and the other from $\{\sqrt{2}/2, \sqrt{3}\}$ with probabilities 0.8 and 0.2 respectively. Then, we modulate χ_{\natural} with these waveforms and take its Fourier transform. Each \mathbf{a}_i corresponds to computing a single entry of this Fourier transform:

$$\mathbf{a}_{(j-1)n+\ell} = \mathbf{W}_n(\ell, :) \operatorname{diag}^*(\psi_j), \quad 1 \leq j \leq 12 \quad \text{and} \quad 1 \leq \ell \leq n,$$

where $\mathbf{W}_n(\ell, :)$ is the ℓ th row of the $n \times n$ discrete Fourier transform matrix. We use the fast Fourier transform to implement the measurement operator.

SM6.2. Fourier ptychography. We study a more realistic measurement setup, Fourier ptychography (FP), for the phase retrieval problem. In this setup, $\chi_{\natural} \in \mathbb{C}^n$ corresponds to an unknown high resolution image (vectorized) from a microscopic sample in the Fourier domain. One cannot directly acquire a high resolution image from this sample because of the physical limitations of optical systems. Any measurement is subject to a filter caused by the lens aperture. We can represent this filter by a sparse matrix $\Phi \in \mathbb{C}^{m \times n}$ with $m \leq n$, each row of which has only one non-zero coefficient. Because of this filter, we can acquire only low-resolution images, through m -dimensional Fourier transform.

FP enlightens the sample from L different angles using a LED grid. This lets us to obtain L different aperture matrices Φ_j , $j = 1, 2, \dots, L$. Then, we acquire phaseless measurements from the sample using the following transmission matrices:

$$\mathbf{a}_{(j-1)m+\ell} = \mathbf{W}_m^*(\ell, :) \Phi_j, \quad 1 \leq j \leq L \quad \text{and} \quad 1 \leq \ell \leq m.$$

Here, $\mathbf{W}_m^*(\ell, :)$ is the ℓ th row of the conjugate transpose of discrete Fourier transform matrix.

The aim in FP is to reconstruct complex valued χ_{\natural} from these phaseless measurements. Once we construct χ_{\natural} , we can generate a high resolution image by taking its inverse Fourier transform.

SM7. Details for QAP. This section gives further information about the QAP experiments summarized in [subsection 7.5](#). [Tables SM4](#) and [SM6](#) display the performance of SketchyCGAL, by presenting the upper bound after rounding ([subsection 7.5.3](#)), the objective value $\langle \mathbf{B} \otimes \mathbf{A}, \mathbf{X} \rangle$, feasibility gap, total number of iterations, memory usage, and computation time. Feasibility gap is defined as $\operatorname{dist}_{\mathbf{K}}(\mathcal{A}\mathbf{X})$ where \mathbf{K} is the Cartesian product between a singleton (for equality constraints) and the nonnegative orthant (for inequality constraints). It is evaluated with respect to the original (not rescaled) problem data and without normalization.

Note that the objective value of SketchyCGAL is not a lower bound for the optimum since the iterates are infeasible. Due to sublinear convergence, SketchyCGAL might be impractical when high accuracy is required, for example within a branch and bound procedure.

Tables SM3 and SM5 compare the relative gap (7.7) obtained by SketchyCGAL with the values for the CSDP method [SM6] with clique size $k = \{2, 3, 4\}$ and the PATH method [SM29] reported in [SM6, Tab. 6]. A graphical view of these results appears in Figure 7.4.

REFERENCES

- [1] 7th DIMACS implementation challenge: Semidefinite and related optimization problems, 2000, <http://archive.dimacs.rutgers.edu/Challenges/Seventh/> (accessed July 2020).
- [2] E. D. ANDERSEN, C. ROOS, AND T. TERLAKY, *On implementing a primal-dual interior-point method for conic quadratic optimization*, Mathematical Programming, (2003), pp. 249–277.
- [3] N. BOUMAL, B. MISHRA, P.-A. ABSIL, AND R. SEPULCHRE, *Manopt, a Matlab toolbox for optimization on manifolds*, J. Mach. Learn. Res., 15 (2014), pp. 1455–1459.
- [4] N. BOUMAL, V. VORONINSKI, AND A. BANDEIRA, *The non-convex Burer–Monteiro approach works on smooth semidefinite programs*, in Adv. Neural Information Processing Systems 29, 2016, pp. 2757–2765.
- [5] S. BOYD, N. PARIKH, E. CHU, B. PELEATO, AND J. ECKSTEIN, *Distributed optimization and statistical learning via the alternating direction method of multipliers*, Found. Trends Mach. Learn., 3 (2011), pp. 1–122.
- [6] J. F. S. BRAVO FERREIRA, Y. KHOO, AND A. SINGER, *Semidefinite programming approach for the quadratic assignment problem with a sparse graph*, Comput. Optim. Appl., 69 (2018), pp. 677–712.
- [7] E. J. CANDÈS, X. LI, AND M. SOLTANOLKOTABI, *Phase retrieval from coded diffraction patterns*, Appl. Comp. Harmonic Anal., 39 (2015), pp. 277 – 299.
- [8] L. DING, A. YURTSEVER, V. CEVHER, J. A. TROPP, AND M. UDELL, *An optimal-storage approach to semidefinite programming using approximate complementarity*, 2019, <https://arxiv.org/abs/1902.03373>.
- [9] A. GITTENS, *Topics in Randomized Numerical Linear Algebra*, PhD thesis, California Institute of Technology, Pasadena, CA, USA, 2013.
- [10] G. H. GOLUB AND C. F. VAN LOAN, *Matrix computations*, Johns Hopkins Univ. Press, Baltimore, MD, fourth ed., 2013.
- [11] N. HALKO, P. G. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Rev., 53 (2011), pp. 217–288.
- [12] M. JAGGI, *Revisiting Frank–Wolfe: Projection-free sparse convex optimization*, in Proc. 30th Int. Conf. Machine Learning, vol. 28(1), 2013, pp. 427–435.
- [13] H. LI, G. C. LINDERMAN, A. SZLAM, K. P. STANTON, Y. KLUGER, AND M. TYGERT, *Algorithm 971: an implementation of a randomized algorithm for principal component analysis*, ACM Trans. Math. Software, 43 (2017), pp. Art. 28, 14.
- [14] H. D. MITTELMANN, *An independent benchmarking of SDP and SOCP solvers*, Mathematical Programming, 95 (2003), pp. 407–430.
- [15] MOSEK APS, *MOSEK version 8.1.0.64 (oct. 8, 2018)*, <https://www.mosek.com/downloads/8.1.0.64/> (accessed July 2020).
- [16] M. L. OVERTON AND R. S. WOMERSLEY, *On the sum of the largest eigenvalues of a symmetric matrix*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 41–45, <https://doi.org/10.1137/0613006>.
- [17] N. SREBRO, *Learning with matrix factorizations*, PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 2004.
- [18] J. STURM, *SeDuMi 1.3*, http://sedumi.ie.lehigh.edu/?page_id=58 (accessed July 2020).
- [19] D. SUN, K.-C. TOH, Y. YUAN, AND X.-Y. ZHAO, *SDPNAL+: A Matlab software for semidefinite programming with bound constraints (version 1.0)*, Optimization Methods and Software, 35 (2020), pp. 87–115.
- [20] Y. SUN, Y. GUO, J. A. TROPP, AND M. UDELL, *Tensor random projection for low memory dimension reduction*, in NeurIPS Workshop on Relational Representation Learning, 2018.
- [21] K.-C. TOH, M. J. TODD, AND R. H. TÜTÜNCÜ, *SDPT3 version 4.0 - a MATLAB software for semidefinite-quadratic-linear programming*, <https://blog.nus.edu.sg/mattohkc/software/sdpt3/> (ac-

- cessed July 2020).
- [22] J. A. TROPP, A. YURTSEVER, M. UDELL, AND V. CEVHER, *Fixed-rank approximation of a positive-semidefinite matrix from streaming data*, in Adv. Neural Information Processing Systems 30, 2017, pp. 1225–1234.
 - [23] J. A. TROPP, A. YURTSEVER, M. UDELL, AND V. CEVHER, *Practical sketching algorithms for low-rank matrix approximation*, SIAM J. Matrix Anal. Appl., 38 (2017), pp. 1454–1485.
 - [24] J. A. TROPP, A. YURTSEVER, M. UDELL, AND V. CEVHER, *Streaming low-rank matrix approximation with an application to scientific simulation*, SIAM J. Sci. Comput., 41 (2019), pp. A2430–A2463.
 - [25] I. WALDSPURGER AND A. WATERS, *Rank optimality for the Burer–Monteiro factorization*, 2018, <https://arxiv.org/abs/1812.03046>.
 - [26] Y. YE, M. J. TODD, AND S. MIZUNO, *An $(o(\sqrt{nl}))$ -iteration homogeneous and self-dual linear programming algorithm*, Mathematics of Operations Research, (1994), pp. 53–67.
 - [27] A. YURTSEVER, O. FERCOQ, AND V. CEVHER, *A conditional-gradient-based augmented Lagrangian framework*, in Proc. 36th Int. Conf. Machine Learning, vol. 97, 09–15 Jun 2019, pp. 7272–7281.
 - [28] A. YURTSEVER, O. FERCOQ, F. LOCATELLO, AND V. CEVHER, *A conditional gradient framework for composite convex minimization with applications to semidefinite programming*, in Proc. 35th Intl. Conf. Machine Learning, vol. 80, 2018, pp. 5727–5736.
 - [29] M. ZASLAVSKIY, F. BACH, AND J. VERT, *A path following algorithm for the graph matching problem*, IEEE Trans. Pattern Anal. Mach. Intel., 31 (2009), pp. 2227–2242.

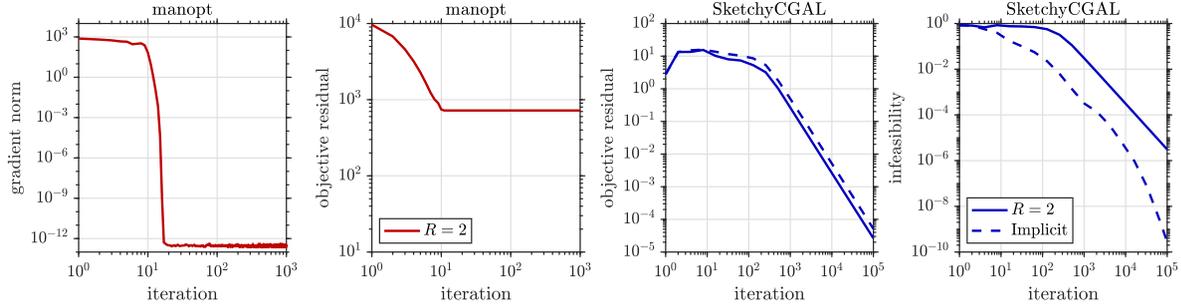


Figure SM7.1. MaxCut SDP: Failure of the Burer–Monteiro heuristic. We apply *Manopt* and *SketchyCGAL* for solving the MaxCut SDP with the dataset C_1 . The subplots show the gradient norm and suboptimality for *Manopt* [left] and the suboptimality and infeasibility for *SketchyCGAL* [right]. *Manopt* with $R = 2$ converges to a spurious solution, whereas *SketchyCGAL* successfully computes a rank-1 approximation of the global optimum. The dashed line describes the convergence of the *SketchyCGAL* implicit iterates. For details, see [subsection SM5.4](#).

Table SM1

We run *Manopt* for solving hard instances of the MaxCut SDP. We consider 10 datasets C_1, \dots, C_{10} . For each dataset, we run *Manopt* with 100 random initializations and report the number of failures. We declare failure when *Manopt* converges to second-order critical point that is not a global optimum. For details, see [subsection SM5.4](#).

Dataset / R	R = 2	3	4	5	6	7	8	9	10	11	12	13
C_1	82	69	63	53	35	32	24	12	11	1	4	0
C_2	77	56	56	36	19	17	12	2	0	0	0	0
C_3	89	65	54	47	44	46	23	11	5	0	3	0
C_4	84	69	50	40	27	23	18	17	1	0	9	0
C_5	85	68	52	51	43	30	31	20	14	3	4	0
C_6	81	68	53	41	23	22	10	10	2	0	1	0
C_7	83	76	60	39	19	19	19	3	0	0	1	0
C_8	81	73	44	34	41	25	8	12	5	4	10	0
C_9	84	64	46	35	25	17	1	10	0	2	4	0
C_{10}	83	71	54	50	31	25	24	16	13	0	8	0

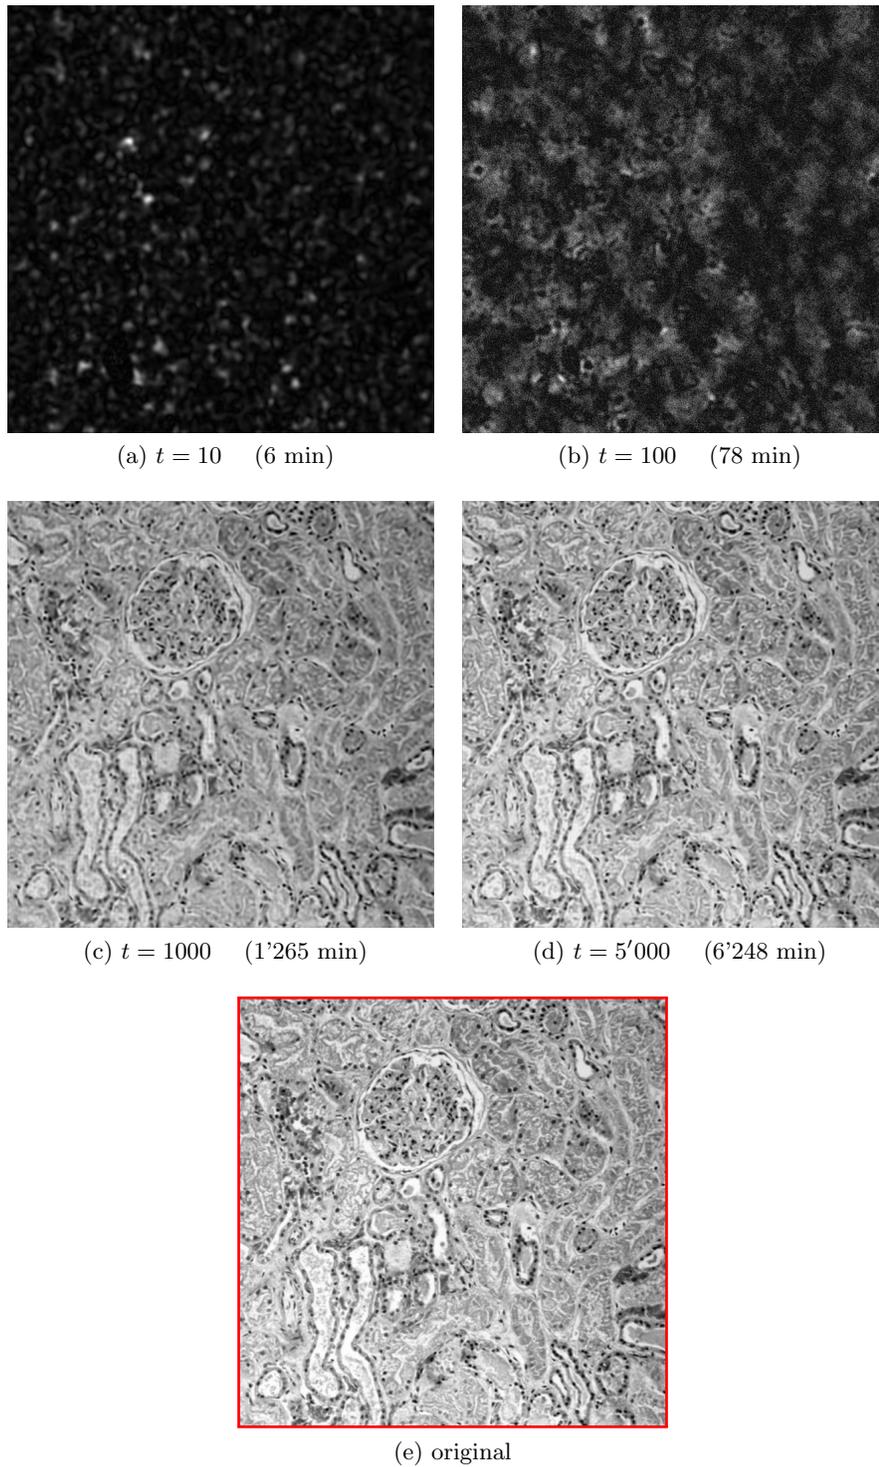


Figure SM7.2. Phase retrieval SDP: Imaging. Reconstruction of an $n = 640^2$ pixel image from Fourier ptychography data. We solve an $n \times n$ phase retrieval SDP via *SketchyCGAL* with rank parameter $R = 5$ and show the images obtained at iterations $t = 10, 10^2, 10^3, 5 \cdot 10^3$. The last subfigure is the original. See [subsection 7.4.3](#).

Table SM2: Numerical outcomes from the MaxCut SDP experiment with GSET Benchmark.

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
Data Name Size (n)	relative cut weight suboptimality infeasibility storage (MB) time (hh:mm:ss)				
G1 800	-1.0861e-02 1.3811e-02 9.5917e-02 8 00:00:01	1.1387e-03 -3.1293e-02 4.6070e-02 226 00:00:04	-1.0861e-02 1.4254e-01 6.0389e-09 180 00:00:04	-8.4961e-03 8.1450e-02 1.9172e-02 245 00:00:02	2.7152e-03 6.4228e-03 1.4574e-02 288 00:00:11
G2 800	-2.2936e-03 1.4340e-02 9.6277e-02 8 00:00:01	6.1750e-04 -3.3205e-02 4.8907e-02 226 00:00:04	-1.3232e-03 5.5569e-02 1.8859e-10 180 00:00:05	-2.6464e-03 8.2332e-02 1.8822e-02 245 00:00:02	2.0289e-03 6.1814e-03 1.4745e-02 288 00:00:11
G3 800	-9.5622e-03 1.5032e-02 9.6560e-02 8 00:00:01	-4.3863e-04 -3.6397e-02 5.3604e-02 226 00:00:05	-6.4041e-03 5.8755e-02 2.4555e-10 180 00:00:04	-7.8077e-03 8.1880e-02 1.9094e-02 245 00:00:02	-2.0177e-03 9.2430e-03 1.9138e-02 288 00:00:11
G4 800	-5.1115e-03 1.6379e-02 9.6255e-02 8 00:00:01	1.4101e-03 -4.0176e-02 5.9377e-02 226 00:00:05	-1.4101e-03 6.2225e-02 2.9489e-10 180 00:00:08	-3.8777e-03 8.3752e-02 1.9058e-02 180 00:00:02	1.2338e-03 1.5740e-02 2.6591e-02 288 00:00:22
G5 800	-7.3504e-03 1.6157e-02 9.6365e-02 8 00:00:01	-6.1253e-04 -3.9680e-02 5.8582e-02 226 00:00:05	-7.5254e-03 6.2757e-02 3.7160e-10 180 00:00:05	-7.3504e-03 8.3001e-02 1.9091e-02 180 00:00:04	-3.4127e-03 1.0683e-02 2.1958e-02 288 00:00:11
G6 800	-2.8718e-02 3.1505e-02 9.3371e-02 8 00:00:01	-1.7436e-02 3.8635e-01 3.0300e-01 226 00:00:03	-8.2051e-03 1.5088e-01 4.8995e-11 180 00:00:05	-3.0769e-02 1.0956e-01 4.9102e-02 178 00:00:05	3.5897e-03 8.4988e-02 2.7171e-02 352 00:00:12
G7 800	2.8802e-03 2.7527e-02 9.5193e-02 8 00:00:01	-2.3618e-02 4.2087e-01 2.9762e-01 226 00:00:03	-1.7857e-02 1.5116e-01 5.0094e-11 180 00:00:04	-1.2097e-02 1.1725e-01 5.0117e-02 178 00:00:03	-1.4977e-02 6.7074e-02 2.3929e-02 288 00:00:12
G8 800	-7.9320e-03 2.6436e-02 9.4596e-02 8 00:00:01	-1.8130e-02 4.1639e-01 2.9784e-01 226 00:00:02	5.6657e-04 1.5401e-01 4.9526e-11 180 00:00:04	-3.9093e-02 1.2247e-01 4.9573e-02 178 00:00:03	-7.3654e-03 8.2034e-02 2.4421e-02 288 00:00:13

Continued on the next page

Table SM2: MaxCut Benchmark with GSet (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G9 800	-2.0533e-02	-1.6648e-02	1.4983e-02	-5.0499e-02	1.6648e-02
	2.8791e-02	4.2360e-01	7.6875e-02	1.1126e-01	8.6664e-02
	9.6469e-02	3.0737e-01	3.4834e-12	4.8213e-02	2.7624e-02
	8	226	180	178	286
	00:00:01	00:00:03	00:00:05	00:00:03	00:00:11
G10 800	-3.3409e-02	-5.8324e-02	-3.3975e-03	-7.5878e-02	-1.1891e-02
	3.0233e-02	4.1742e-01	1.4609e-01	1.1787e-01	6.9872e-02
	9.6215e-02	2.9681e-01	5.0049e-11	4.8529e-02	2.3584e-02
	8	226	180	178	286
	00:00:01	00:00:03	00:00:06	00:00:03	00:00:11
G11 800	-1.5444e-02	-2.3166e-02	1.5444e-02	-7.3359e-02	1.5444e-02
	2.6265e-02	6.1699e-01	7.2209e-02	4.4688e-02	7.9842e-02
	9.5640e-02	0	5.3829e-11	3.7060e-02	2.5442e-02
	8	226	185	183	288
	00:00:01	00:00:02	00:00:04	00:00:03	00:00:10
G12 800	-3.9062e-03	-1.9531e-02	3.5156e-02	-7.8125e-03	1.9531e-02
	2.4141e-02	6.5004e-01	8.2364e-02	5.9988e-02	9.3180e-02
	9.5875e-02	0	5.0553e-11	4.8779e-02	2.8291e-02
	8	226	185	183	288
	00:00:01	00:00:01	00:00:04	00:00:03	00:00:10
G13 800	3.7736e-03	-1.8868e-02	2.6415e-02	-4.1509e-02	3.0189e-02
	2.3104e-02	6.3325e-01	8.5197e-02	5.7293e-02	8.2228e-02
	9.5982e-02	0	5.5987e-11	4.9942e-02	2.5854e-02
	8	226	185	178	288
	00:00:01	00:00:01	00:00:03	00:00:03	00:00:10
G14 800	-7.0779e-03	-1.3819e-02	-2.5278e-02	-6.2016e-02	-3.8760e-02
	4.7170e-02	-2.1354e-02	8.2634e-02	8.7609e-02	1.3110e-01
	7.6060e-02	5.6505e-02	4.3439e-11	1.8669e-02	2.6233e-02
	8	226	181	190	288
	00:00:02	00:00:05	00:00:03	00:00:02	00:00:09
G15 800	-1.4137e-02	-1.8849e-02	-2.2215e-02	-7.1693e-02	-4.7122e-02
	4.0370e-02	-2.5681e-02	8.7010e-02	8.8979e-02	1.3943e-01
	6.8233e-02	6.7477e-02	4.1721e-11	1.8486e-02	2.4687e-02
	8	226	186	180	288
	00:00:02	00:00:04	00:00:04	00:00:02	00:00:09
G16 800	-7.0994e-03	-9.4659e-03	-1.9608e-02	-5.0372e-02	-3.9892e-02
	3.8136e-02	-2.2376e-02	8.5358e-02	8.6831e-02	1.3028e-01
	7.0830e-02	5.9677e-02	4.8317e-11	1.8906e-02	2.8522e-02
	8	226	181	180	288
	00:00:02	00:00:04	00:00:04	00:00:02	00:00:09
G17 800	-1.2860e-02	-2.3689e-03	-1.8274e-02	-5.9560e-02	-3.6548e-02
	4.7458e-02	-2.2381e-02	7.8986e-02	8.7610e-02	1.3277e-01
	6.8409e-02	5.9613e-02	3.6482e-11	1.8555e-02	2.4114e-02
	8	226	186	180	286
	00:00:02	00:00:04	00:00:04	00:00:02	00:00:09

Continued on the next page

Table SM2: MaxCut Benchmark with GSet (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G18 800	-3.6954e-02	-1.0414e-01	-1.0078e-02	-1.1646e-01	-6.7189e-03
	4.2249e-02	3.7409e-01	1.1502e-01	1.0608e-01	1.0740e-01
	7.1076e-02	1.9476e-01	2.4288e-12	6.1771e-02	2.3304e-02
	8	226	186	178	289
	00:00:01	00:00:03	00:00:05	00:00:02	00:00:17
G19 800	-2.5031e-02	-9.3867e-02	1.2516e-03	-1.2390e-01	0
	6.0931e-02	4.4872e-01	1.4566e-01	1.1875e-01	1.4665e-01
	8.0855e-02	2.0949e-01	2.4037e-12	6.6121e-02	2.3218e-02
	8	226	186	178	355
	00:00:01	00:00:03	00:00:05	00:00:03	00:00:17
G20 800	-3.8232e-02	-1.0036e-01	-1.7921e-02	-1.2067e-01	-8.3632e-03
	5.5135e-02	4.0152e-01	1.3282e-01	1.1077e-01	1.2889e-01
	8.5073e-02	1.9516e-01	2.2402e-12	6.4826e-02	2.9589e-02
	8	226	186	178	289
	00:00:01	00:00:03	00:00:05	00:00:02	00:00:17
G21 800	-6.1728e-03	-5.6790e-02	4.4444e-02	-1.0617e-01	4.1975e-02
	5.7118e-02	4.1582e-01	1.3161e-01	1.1080e-01	1.3462e-01
	8.4940e-02	1.9936e-01	2.4262e-12	6.6784e-02	2.7448e-02
	8	226	186	178	289
	00:00:01	00:00:03	00:00:05	00:00:02	00:00:17
G22 2000	-1.2115e-02	1.0032e-03	-9.1056e-03	-1.0417e-02	-6.9450e-04
	1.5681e-02	-8.1412e-03	9.1982e-02	6.8851e-02	1.1905e-02
	9.6934e-02	4.1194e-02	2.6086e-10	2.8216e-02	1.2584e-02
	9	984	670	492	1300
	00:00:01	00:01:14	00:00:32	00:00:24	00:03:18
G23 2000	-8.0850e-03	-5.3900e-03	-1.3706e-02	-1.3398e-02	-6.0060e-03
	1.3379e-02	-7.8396e-03	9.2914e-02	6.8894e-02	3.0420e-02
	9.6457e-02	3.9683e-02	3.3246e-10	2.8258e-02	3.3485e-02
	9	984	704	492	1300
	00:00:01	00:01:15	00:00:33	00:00:24	00:03:43
G24 2000	-1.1172e-02	-1.2328e-03	-1.1557e-02	-1.2405e-02	-4.0835e-03
	1.3818e-02	-7.3304e-03	9.2312e-02	6.8781e-02	2.8684e-02
	9.6751e-02	3.7142e-02	4.1922e-10	2.8311e-02	2.9243e-02
	9	984	703	492	1300
	00:00:01	00:01:21	00:00:33	00:00:34	00:03:43
G25 2000	-3.4973e-03	2.4870e-03	-3.4196e-03	-4.6631e-03	1.6321e-03
	1.3900e-02	-7.7270e-03	1.0041e-01	6.9165e-02	1.2113e-02
	9.6988e-02	3.9683e-02	7.3474e-10	2.8242e-02	1.1920e-02
	9	984	670	492	1235
	00:00:01	00:01:21	00:00:31	00:00:34	00:04:26
G26 2000	-4.5073e-03	1.0880e-03	-4.2742e-03	-6.6832e-03	1.5542e-03
	1.6124e-02	-7.7421e-03	1.0441e-01	6.8630e-02	3.1191e-02
	9.6396e-02	3.9605e-02	9.9323e-10	2.8347e-02	2.9950e-02
	9	984	670	492	1300
	00:00:01	00:01:01	00:00:31	00:00:28	00:07:28

Continued on the next page

Table SM2: MaxCut Benchmark with GSet (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G27 2000	-1.4187e-02	-3.9792e-02	-6.9204e-04	-3.5294e-02	-1.0035e-02
	3.6489e-02	3.4331e-01	1.3416e-01	1.2784e-01	5.9977e-02
	9.7233e-02	2.2480e-01	5.3716e-12	4.7806e-02	1.7503e-02
	9	984	769	558	1300
	00:00:01	00:00:48	00:00:42	00:00:29	00:03:50
G28 2000	-9.4970e-03	-2.2511e-02	2.8139e-03	-2.5677e-02	-4.5726e-03
	3.7008e-02	3.7142e-01	1.2557e-01	1.2923e-01	1.2780e-01
	9.7469e-02	2.3601e-01	5.4840e-12	4.7381e-02	3.6643e-02
	9	984	735	561	1300
	00:00:01	00:00:56	00:00:48	00:00:54	00:03:17
G29 2000	-2.0168e-02	-3.1261e-02	-1.6807e-03	-3.1933e-02	-9.0756e-03
	3.4436e-02	3.6747e-01	1.2288e-01	1.2817e-01	1.1352e-01
	9.7664e-02	2.4240e-01	5.1238e-12	4.7327e-02	3.7082e-02
	9	984	704	558	1300
	00:00:01	00:00:56	00:00:48	00:00:22	00:03:17
G30 2000	-2.6605e-02	-3.3588e-02	9.9767e-04	-3.7912e-02	-3.9907e-03
	3.1683e-02	3.5612e-01	1.3726e-01	1.2752e-01	5.3960e-02
	9.7498e-02	2.3418e-01	5.7873e-12	4.7109e-02	1.5929e-02
	9	984	671	558	1300
	00:00:01	00:00:56	00:00:48	00:00:24	00:05:14
G31 2000	-1.2170e-02	-4.3811e-02	-1.7385e-03	-3.0250e-02	7.9972e-03
	3.2478e-02	4.1374e-01	1.2072e-01	1.2853e-01	1.3652e-01
	9.7291e-02	2.5613e-01	4.4749e-12	4.7889e-02	4.0526e-02
	9	984	704	558	1204
	00:00:01	00:00:47	00:00:48	00:00:24	00:04:19
G32 2000	-4.7022e-03	-2.5078e-02	2.5078e-02	-9.2476e-02	3.2915e-02
	2.3071e-02	6.3823e-01	7.5100e-02	4.6005e-02	8.7684e-02
	9.7150e-02	0	3.3059e-11	4.2067e-02	2.7141e-02
	9	984	558	624	1235
	00:00:01	00:00:30	00:00:21	00:00:36	00:02:56
G33 2000	2.5974e-02	-2.7597e-02	4.5455e-02	-3.0844e-02	5.5195e-02
	2.4605e-02	6.4997e-01	7.3805e-02	4.5305e-02	8.8476e-02
	9.6924e-02	0	3.2314e-11	4.2924e-02	2.6904e-02
	9	984	595	561	1300
	00:00:01	00:02:12	00:00:17	00:01:00	00:02:57
G34 2000	-1.6000e-03	-1.7600e-02	3.3600e-02	-2.4000e-02	4.6400e-02
	3.1015e-02	6.6216e-01	7.5042e-02	5.0275e-02	9.5898e-02
	9.7385e-02	0	2.8472e-11	4.1950e-02	2.9024e-02
	9	984	595	624	1300
	00:00:01	00:00:32	00:00:18	00:00:27	00:02:53
G35 2000	-5.9347e-03	-2.4278e-03	-2.6976e-02	-7.4049e-02	-5.6110e-02
	3.6621e-02	-2.8394e-02	1.0806e-01	9.4208e-02	1.5892e-01
	5.3605e-02	7.6066e-02	5.0420e-11	1.7448e-02	2.3757e-02
	9	984	682	592	1300
	00:00:13	00:01:26	00:00:29	00:00:19	00:02:44

Continued on the next page

Table SM2: MaxCut Benchmark with GSet (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G36 2000	-1.1345e-02	-3.2415e-03	-1.9044e-02	-6.9692e-02	-6.3614e-02
	3.6011e-02	-1.4296e-02	8.0139e-02	8.9393e-02	1.6699e-01
	4.9253e-02	3.8655e-02	1.0676e-11	8.3391e-03	1.4397e-02
	9	984	591	592	1300
	00:00:11	00:01:43	00:00:29	00:00:24	00:03:40
G37 2000	-1.6125e-02	-2.4187e-03	-3.3593e-02	-7.8877e-02	-6.7052e-02
	4.7562e-02	-1.2850e-02	1.1191e-01	9.4806e-02	1.6740e-01
	5.7891e-02	3.5075e-02	4.7745e-11	1.7306e-02	1.9096e-02
	9	984	703	592	1269
	00:00:16	00:01:29	00:00:27	00:00:19	00:03:40
G38 2000	-5.4039e-03	-1.7563e-03	-2.8641e-02	-7.3629e-02	-5.6336e-02
	4.3352e-02	-1.2564e-02	1.1024e-01	9.4460e-02	1.6723e-01
	5.3739e-02	3.4143e-02	4.8473e-11	1.7351e-02	2.2529e-02
	9	984	639	594	1300
	00:00:14	00:01:42	00:00:27	00:00:17	00:03:39
G39 2000	-6.1445e-02	-1.2899e-01	-1.6417e-02	-1.4071e-01	-3.2364e-02
	6.6946e-02	4.4741e-01	1.2226e-01	1.2104e-01	2.9167e-01
	9.6734e-02	2.3043e-01	6.2560e-13	6.8073e-02	3.4601e-02
	9	984	661	627	1498
	00:00:01	00:00:55	00:00:50	00:00:31	00:06:45
G40 2000	-5.1092e-02	-1.9368e-01	-2.8333e-02	-1.8346e-01	-3.1119e-02
	6.0931e-02	4.8275e-01	1.4461e-01	1.0003e-01	1.9028e-01
	8.1869e-02	2.4701e-01	6.1378e-13	6.7156e-02	2.8995e-02
	9	984	692	624	1564
	00:00:01	00:00:55	00:00:40	00:00:33	00:07:50
G41 2000	-3.2640e-02	-1.3907e-01	-2.3179e-02	-1.6509e-01	-2.2233e-02
	5.3490e-02	4.6215e-01	1.2991e-01	1.1735e-01	1.9720e-01
	9.7598e-02	2.3489e-01	5.6737e-13	6.9128e-02	3.7320e-02
	9	984	742	624	1367
	00:00:01	00:00:43	00:00:40	00:00:34	00:06:28
G42 2000	-3.9796e-02	-1.2772e-01	8.7922e-03	-1.4715e-01	-7.8667e-03
	5.3904e-02	4.4115e-01	1.2281e-01	1.1432e-01	2.7490e-01
	9.7783e-02	2.3361e-01	6.2319e-13	6.9205e-02	3.3145e-02
	9	984	607	624	1498
	00:00:01	00:01:06	00:00:41	00:00:34	00:06:44
G43 1000	-1.2117e-02	-1.6871e-03	-1.7025e-02	-2.1012e-02	-4.7546e-03
	1.5552e-02	-7.0064e-03	8.6541e-02	6.5582e-02	1.7807e-02
	9.6716e-02	3.3450e-02	6.1918e-11	2.7835e-02	2.1330e-02
	8	299	245	246	417
	00:00:01	00:00:12	00:00:07	00:00:06	00:00:20
G44 1000	1.7025e-03	-1.2382e-03	-8.0483e-03	-1.5942e-02	-2.0121e-03
	1.3966e-02	-6.4682e-03	8.5723e-02	6.4779e-02	1.7696e-02
	9.6710e-02	3.1378e-02	6.1629e-11	2.8081e-02	1.9844e-02
	8	299	311	246	417
	00:00:01	00:00:13	00:00:07	00:00:05	00:00:20

Continued on the next page

Table SM2: MaxCut Benchmark with GSet (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G45 1000	-8.0695e-03	1.0863e-03	-8.6903e-03	-1.1173e-02	7.7592e-04
	1.5134e-02	-6.2559e-03	8.7887e-02	6.4124e-02	2.0988e-02
	9.6247e-02	3.0095e-02	6.1268e-11	2.8040e-02	2.6194e-02
	8	299	311	180	417
	00:00:01	00:00:13	00:00:06	00:00:04	00:00:29
G46 1000	-9.3516e-04	1.4027e-03	-2.6496e-03	-4.2082e-03	2.1820e-03
	1.5542e-02	-6.0780e-03	8.5848e-02	6.4099e-02	1.6751e-02
	9.6680e-02	2.9203e-02	6.0013e-11	2.8459e-02	2.0716e-02
	8	299	311	188	417
	00:00:01	00:00:11	00:00:06	00:00:04	00:00:29
G47 1000	-3.4098e-03	2.0149e-03	-8.6795e-03	-6.6646e-03	-9.2994e-04
	1.5588e-02	-6.5081e-03	8.7395e-02	6.5387e-02	1.7486e-02
	9.6696e-02	3.1407e-02	6.0498e-11	2.8090e-02	2.2294e-02
	8	299	311	261	417
	00:00:01	00:00:08	00:00:08	00:00:11	00:00:27
G48 3000	-1.3667e-02	0	0	0	0
	1.0092e-02	4.7282e-02	1.2901e-01	4.1637e-02	1.5852e-02
	9.4775e-02	2.2182e-15	3.5505e-11	6.8384e-02	1.5935e-02
	10	2087	1024	879	2339
	00:00:01	00:03:26	00:00:44	00:01:27	00:12:26
G49 3000	-1.4000e-02	0	0	0	0
	9.7028e-03	4.7281e-02	1.2905e-01	4.1763e-02	4.7003e-02
	9.7780e-02	1.8776e-15	3.5624e-11	6.8213e-02	4.7125e-02
	10	2087	958	879	2273
	00:00:01	00:03:14	00:00:43	00:01:24	00:09:31
G50 3000	-1.3946e-02	0	0	0	0
	8.0999e-03	4.2954e-02	1.2508e-01	3.9432e-02	3.2078e-02
	9.7044e-02	1.5870e-15	3.5298e-11	6.8811e-02	3.6382e-02
	10	2087	958	879	2414
	00:00:01	00:03:51	00:00:46	00:01:25	00:09:45
G51 1000	-1.6319e-02	-9.3633e-03	-2.6752e-02	-7.7582e-02	-5.1899e-02
	4.4961e-02	-2.4666e-02	8.9620e-02	8.9975e-02	1.3953e-01
	6.1555e-02	6.4965e-02	4.7564e-11	1.7660e-02	2.5262e-02
	8	299	311	196	417
	00:00:04	00:00:08	00:00:07	00:00:03	00:00:22
G52 1000	-8.8901e-03	-3.7716e-03	-2.1013e-02	-5.4149e-02	-3.8793e-02
	3.5672e-02	-3.0821e-02	8.8583e-02	8.7957e-02	1.4374e-01
	6.1392e-02	7.7417e-02	3.9779e-11	1.8136e-02	2.4383e-02
	8	299	245	196	417
	00:00:03	00:00:08	00:00:10	00:00:03	00:00:15
G53 1000	-4.8767e-03	-2.9802e-03	-5.1477e-03	-6.0688e-02	-3.7930e-02
	4.9023e-02	-2.9221e-02	8.7032e-02	8.8963e-02	1.3194e-01
	7.5129e-02	7.3965e-02	4.9598e-11	1.7903e-02	2.6304e-02
	8	299	245	196	417
	00:00:04	00:00:11	00:00:06	00:00:03	00:00:16

Continued on the next page

Table SM2: MaxCut Benchmark with GSet (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G54 1000	-1.0243e-02	0	-1.7251e-02	-6.9542e-02	-3.6927e-02
	4.7776e-02	-2.3638e-02	8.2065e-02	8.7252e-02	1.2880e-01
	7.5128e-02	6.1915e-02	4.3089e-11	1.8259e-02	2.5147e-02
	8	299	245	196	417
	00:00:06	00:00:10	00:00:06	00:00:03	00:00:16
G55 5000	-1.5392e-02	-7.2911e-03	-9.2152e-03	-1.9949e-02	-1.7215e-03
	3.2784e-02	3.8322e-02	9.8695e-02	7.7692e-02	5.5184e-02
	9.8225e-02	5.9028e-02	2.7354e-11	3.7645e-02	2.9373e-02
	12	6541	2544	2444	6080
	00:00:01	00:18:23	00:04:29	00:05:04	00:59:13
G56 5000	-1.3081e-02	-7.5146e-02	-8.3496e-04	-4.4531e-02	1.1133e-03
	4.2961e-02	6.0115e-01	1.5219e-01	1.4778e-01	1.3967e-01
	9.8497e-02	2.4789e-01	6.0728e-12	3.5298e-02	3.1123e-02
	12	6476	2736	2432	6080
	00:00:02	00:09:19	00:06:02	00:05:35	00:59:12
G57 5000	-1.5873e-02	-3.6190e-02	2.0317e-02	-2.6984e-01	2.5397e-02
	3.4139e-02	6.5144e-01	9.2966e-02	4.9293e-02	1.0633e-01
	9.8470e-02	0	3.3474e-11	4.2377e-02	3.2602e-02
	12	6476	2652	2349	6080
	00:00:01	00:09:20	00:04:52	00:11:27	00:44:54
G58 5000	-9.9085e-03	-6.0851e-03	-3.7641e-02	-7.7868e-02	-9.0361e-02
	4.0681e-02	-1.5515e-02	1.0405e-01	7.7087e-02	2.1977e-01
	4.6279e-02	4.2906e-02	1.3232e-11	7.2850e-03	2.5553e-02
	12	6541	2749	2354	6080
	00:02:31	00:34:53	00:09:44	00:08:31	00:28:48
G59 5000	-4.4515e-02	-2.1826e-01	-1.1458e-02	-1.9591e-01	-1.5026e-02
	4.7775e-02	5.1949e-01	1.3418e-01	1.0923e-01	1.9235e-01
	9.8420e-02	3.0113e-01	3.7377e-14	6.7124e-02	3.1820e-02
	79	6541	2827	2349	6862
	00:00:02	00:14:27	00:14:07	00:06:30	02:04:37
G60 7000	-1.5879e-02	-5.9546e-03	-6.1016e-03	-2.0290e-02	1.3968e-03
	3.4044e-02	4.0277e-02	1.0393e-01	7.8355e-02	6.5477e-02
	9.8489e-02	5.9479e-02	3.0375e-11	3.7698e-02	3.6393e-02
	17	12682	5017	4756	11705
	00:00:02	00:38:37	00:11:47	00:15:47	02:41:03
G61 7000	-1.3506e-02	-8.9138e-02	-1.2541e-02	-3.8009e-02	6.5599e-03
	3.7879e-02	5.7665e-01	1.9524e-01	1.4471e-01	1.5498e-01
	9.8261e-02	2.4786e-01	2.5011e-11	3.5055e-02	3.7988e-02
	17	12682	5329	4734	11705
	00:00:03	00:19:08	00:12:54	00:17:20	03:49:14
G62 7000	-2.0464e-02	-6.0482e-02	1.0459e-02	-2.9013e-01	2.5466e-02
	3.5993e-02	6.5473e-01	9.1622e-02	3.0213e-02	9.8260e-02
	9.8628e-02	0	3.8092e-11	4.5815e-02	3.0209e-02
	17	12682	5155	4571	11705
	00:00:02	00:15:42	00:26:48	00:21:52	02:36:22

Continued on the next page

Table SM2: MaxCut Benchmark with GSet (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G63 7000	-6.6179e-03	-4.5402e-03	-3.6553e-02	-7.7299e-02	-9.0535e-02
	3.3992e-02	-1.5066e-02	1.0744e-01	7.8815e-02	2.2130e-01
	3.5469e-02	4.2551e-02	1.3603e-11	7.5063e-03	2.5199e-02
	17	12682	5236	4678	11705
	00:04:40	00:59:53	00:14:56	00:14:15	01:42:50
G64 7000	-6.0457e-02	-2.4390e-01	-1.8344e-02	-2.2181e-01	-1.8861e-02
	5.5067e-02	5.2060e-01	1.1120e-01	1.0273e-01	1.5301e-01
	9.8677e-02	3.1977e-01	1.4076e-14	6.8433e-02	3.3139e-02
	17	12682	5458	4714	13237
	00:00:04	00:25:34	00:25:26	00:24:07	06:54:36
G65 8000	-4.4426e-03	-4.1599e-02	2.4637e-02	-2.7746e-01	3.6349e-02
	3.1221e-02	6.5182e-01	9.4792e-02	2.9541e-02	1.0447e-01
	9.8691e-02	0	3.8620e-11	4.6184e-02	3.2221e-02
	17	15496	6863	6164	15221
	00:00:02	00:27:01	00:11:39	00:40:02	03:27:59
G66 9000	-1.8544e-02	-4.0588e-02	1.5395e-02	-3.1176e-01	-
	3.3867e-02	6.4642e-01	8.6972e-02	2.8440e-02	-
	9.8820e-02	0	3.7667e-11	4.6642e-02	-
	17	21970	8340	7755	-
	00:00:04	00:42:07	00:16:41	00:55:07	-
G67 10000	-1.8971e-02	-	3.2797e-02	-3.2990e-01	-
	3.0792e-02	-	4.8610e-02	3.3673e-02	-
	9.8756e-02	-	8.0771e-12	4.4400e-02	-
	17	-	10269	9540	-
	00:00:03	-	00:25:29	01:13:18	-

Table SM3

We solve SDP relaxations of QAP instances from QAPLIB using SketchyCGAL. We compute the relative gap and compare it with the values for the CSDP method [SM6] with clique size $k = \{2, 3, 4\}$ and the PATH method [SM29] reported in [SM6, Tab. 4]. Smaller is better. See subsection 7.5.

Dataset	Optimum	SketchyCGAL	CSDP2	CSDP3	CSDP4	PATH
chr12a	9552	0	34.5	6	0	42.7
chr12b	9742	0	38.9	25.4	11.9	38.1
chr12c	11156	0	5.8	2.3	2.3	18.6
chr15a	9896	0.4	2.1	2.1	2.1	52
chr15b	7990	0	26.3	34.5	29.2	158.6
chr15c	9504	0	0	0	0	63.3
chr18a	11098	1.5	69.8	0.2	0.2	76.3
chr18b	1534	8.7	8.9	22.9	29.5	99.3
chr20a	2192	6.6	122.5	76.1	43.8	95.4
chr20b	2298	0	62.9	9.3	9.3	82.2
chr20c	14142	24.7	173	100.1	111.5	88.9
chr22a	6156	0	17.2	7.6	3	38.3
chr22b	6194	7.1	7.3	2.3	1	40.4
chr25a	3796	21.8	107	49.2	25.2	69.9
esc16a	68	2.9	8.8	11.8	11.8	11.8
esc16b	292	0	0	0.7	0	2.7
esc16c	160	5	5	7.5	8.7	6.3
esc16d	16	12.5	12.5	50	25	75
esc16e	28	0	14.3	7.1	14.3	21.4
esc16g	26	7.7	7.7	0	15.4	15.4
esc16h	996	0	1.6	0	1.6	16.9
esc16i	14	0	0	0	0	57.1
esc16j	8	0	0	0	0	75
esc32a	130	52.3	115.4	124.6	113.8	93.8
esc32b	168	57.1	109.5	114.3	111.9	88.1
esc32c	642	2.2	12.8	15.9	13.7	7.8
esc32d	200	14	38	36	39	21
esc32e	2	0	0	0	0	600
esc32g	6	0	0	0	0	366.7
esc32h	438	5	24.7	26.9	22.8	18.3
esc64a	116	6.9	60.3	53.4	60.3	106.9
esc128	64	28.1	250	206.3	175	221.9
ste36a	9526	15.6	70.2	74.7	74.2	76.3
ste36b	15852	15.7	188.8	204.3	211.9	158.6
ste36c	8239110	9	66	62.8	63.7	83.2

Table SM4

We run *SketchyCGAL* for (the first of) 10^6 iterations or 72 hours of runtime, for solving SDP relaxations of QAP instances from QAPLIB. We report the upper bound after rounding ([subsection 7.5.3](#)), the objective value, the feasibility gap, the number of iterations, the memory usage (in MB), and the cpu time ('hh:mm:ss'). See [subsection 7.5](#).

Dataset	Optimum	Upper bnd.	Objective	Feas. gap	Iteration	Memory	Time
chr12a	9552	9552	9 576,13	5.56e−5	1000000	201	05 : 21 : 44
chr12b	9742	9742	9 759,34	2.53e−5	1000000	201	05 : 22 : 35
chr12c	11156	11156	11 021,45	2.04e−4	1000000	201	05 : 24 : 09
chr15a	9896	9936	9 519,53	1.92e−4	1000000	214	09 : 20 : 38
chr15b	7990	7990	7 469,92	1.73e−4	1000000	214	09 : 58 : 04
chr15c	9504	9504	9 517,71	3.55e−5	1000000	214	08 : 40 : 31
chr18a	11098	11262	10 700,65	2.38e−4	1000000	214	14 : 04 : 57
chr18b	1534	1668	1 534,29	7.64e−5	1000000	214	13 : 52 : 37
chr20a	2192	2336	2 183,68	1.13e−4	1000000	201	18 : 59 : 05
chr20b	2298	2298	2 288,94	1.08e−4	1000000	201	18 : 47 : 35
chr20c	14142	17630	13 316,48	1.91e−4	1000000	201	20 : 20 : 08
chr22a	6156	6156	6 152,80	9.99e−5	1000000	193	27 : 47 : 48
chr22b	6194	6634	6 198,39	1.05e−4	1000000	193	27 : 35 : 40
chr25a	3796	4624	3 692,50	2.08e−4	1000000	197	45 : 30 : 14
esc16a	68	70	60,36	5.96e−5	1000000	201	15 : 36 : 34
esc16b	292	292	287,64	9.65e−4	1000000	198	24 : 47 : 46
esc16c	160	168	145,01	4.71e−5	1000000	197	17 : 28 : 31
esc16d	16	18	13,00	6.16e−5	1000000	201	11 : 37 : 42
esc16e	28	28	25,41	6.11e−5	1000000	201	11 : 32 : 58
esc16g	26	28	22,46	7.03e−5	1000000	201	12 : 00 : 45
esc16h	996	996	975,41	2.10e−4	1000000	198	26 : 53 : 14
esc16i	14	14	11,37	7.43e−5	1000000	201	10 : 42 : 41
esc16j	8	8	7,11	9.00e−5	1000000	201	10 : 05 : 17
esc32a	130	198	99,61	3.62e−4	569710	262	72 : 00 : 00
esc32b	168	264	118,36	2.86e−4	449002	197	72 : 00 : 00
esc32c	642	656	610,84	7.01e−4	402329	197	72 : 00 : 00
esc32d	200	228	187,33	5.92e−4	494696	273	72 : 00 : 00
esc32e	2	2	1,90	4.16e−4	1000000	197	60 : 25 : 27
esc32g	6	6	5,83	4.37e−4	1000000	197	66 : 37 : 07
esc32h	438	460	417,79	3.48e−4	363879	197	72 : 00 : 00
esc64a	116	124	97,85	7.28e−3	126620	288	72 : 00 : 00
esc128	64	82	53,95	4.06e−2	23882	587	72 : 00 : 00
ste36a	9526	11012	8 995,09	1.03e−3	256265	263	72 : 00 : 00
ste36b	15852	18336	15 344,24	1.03e−3	257781	263	72 : 00 : 00
ste36c	8239110	8983468	7 984 574,51	1.05e−3	257657	263	72 : 00 : 00

Table SM5

We solve SDP relaxations of QAP instances from TSPLIB using SketchyCGAL. We compute the relative gap and compare it with the values for the CSDP method [SM6] with clique size $k = \{2, 3, 4\}$ and the PATH method [SM29] reported in [SM6, Tab. 6]. Smaller is better.

Dataset	Optimum	SketchyCGAL	CSDP2	CSDP3	CSDP4	PATH
att48	10628	73.8	213	236.5	233.6	329.8
bayg29	1610	38	114.3	115.8	114.3	210.1
bays29	2020	44	107.6	118.3	115.4	164.8
berlin52	7542	52	175	127.2	127.2	280.6
bier127	118282	73.2	216.4	193.8	193.8	234.2
brazil58	25395	86.5	248	200.8	200.8	337
burma14	3323	10.8	24.6	28.4	32.3	95.5
ch130	6110	102.8	352.4	380.6	380.6	621.3
ch150	6528	121.1	346.9	318.2	318.2	689.3
dantzig42	699	73.5	193.1	174	174	82
eil101	629	70.6	227.3	235.3	235.3	437.7
eil51	426	57.3	203.6	205.4	205.5	244.4
eil76	538	65.4	282.9	183	183	328.2
fri26	937	32.7	91.6	39.4	39.4	41.6
gr120	6942	123.6	445.2	261.6	261.6	617.6
gr137	69853	147.5	264.6	220.3	220.3	38.9
gr17	2085	15.3	46.8	32.4	44.9	86.9
gr21	2707	18.8	94.5	69.7	66.3	185.7
gr24	1272	30	89.2	86.2	73.9	129.4
gr48	5046	59.7	210.2	187.4	187.4	270.4
gr96	55209	119.3	228.9	201.7	201.7	46
hk48	11461	43.9	222.4	207.7	207.7	281.6
kroA100	21282	125	469.6	469	469	720.2
kroA150	26524	168.4	411	467.4	467.4	945.8
kroB100	22141	164.3	411.9	313.6	313.6	624.2
kroB150	26130	166.8	417.3	353.7	353.7	844.7
kroC100	20749	163.4	507.4	445.1	445.1	763
kroD100	21294	118.5	504.2	349.8	349.8	654.4
kroE100	22068	137.1	489.5	346.3	346.3	684.2
lin105	14379	167.2	303.1	234.8	234.8	248.4
pr107	44303	195.3	181.5	207.9	207.9	41.6
pr124	59030	192.4	293.8	180.2	180.2	67.6
pr136	96772	148.6	325.5	164.7	164.7	196.6
pr144	58537	270.4	255	283.7	283.7	59.8
pr76	108159	81	192.2	194	194	39.4
rat99	1211	95.3	236.4	161.5	161.5	444.1
rd100	7910	89	438.4	375.3	375.3	506.5
st70	675	70.8	300.9	320	317.9	387.9
swiss42	1273	35.5	163.2	190.4	190.8	194
ulysses16	6859	11.5	23.6	20.2	23.2	82.7
ulysses22	7013	26.4	64.5	57	59.7	126.3

Table SM6

We run *SketchyCGAL* for (the first of) 10^6 iterations or 72 hours of runtime, for solving SDP relaxations of QAP instances from *TSPLIB*. We report the upper bound after rounding ([subsection 7.5.3](#)), the objective value, the feasibility gap, the number of iterations, the memory usage (in MB), and the cpu time ('hh:mm:ss'). See [subsection 7.5](#).

Dataset	Optimum	Upper bnd.	Objective	Feas. gap	Iteration	Memory	Time
att48	10628	18474	9 076,89	8.18e-4	274725	207	72 : 00 : 00
bayg29	1610	2222	1 498,61	5.43e-4	1000000	197	69 : 06 : 57
bays29	2020	2908	1 855,26	6.49e-4	1000000	193	69 : 03 : 56
berlin52	7542	11463	6 662,82	2.10e-3	228895	267	72 : 00 : 00
bier127	118282	204819	112 696,76	6.14e-2	19905	1006	72 : 00 : 00
brazil58	25395	47362	18 379,06	1.87e-3	173449	262	72 : 00 : 00
burma14	3323	3682	3 153,62	5.15e-4	1000000	219	08 : 22 : 03
ch130	6110	12389	7 874,20	6.79e-2	15373	922	72 : 00 : 00
ch150	6528	14432	13 316,42	2.36e-1	10047	1592	72 : 00 : 00
dantzig42	699	1213	589,86	5.57e-4	350934	262	72 : 00 : 00
eil101	629	1073	618,22	1.60e-2	34377	529	72 : 00 : 00
eil51	426	670	407,33	4.52e-3	213654	202	72 : 00 : 00
eil76	538	890	534,11	1.05e-2	77730	287	72 : 00 : 00
fri26	937	1243	858,27	6.59e-4	1000000	197	56 : 04 : 52
gr120	6942	15525	7 409,45	2.75e-2	23936	742	72 : 00 : 00
gr137	69853	172902	101 276,79	6.71e-2	16900	1206	72 : 00 : 00
gr17	2085	2403	1 800,49	3.84e-4	1000000	214	12 : 35 : 03
gr21	2707	3217	2 570,13	5.52e-4	1000000	193	22 : 26 : 13
gr24	1272	1653	1 136,69	3.47e-4	1000000	197	33 : 36 : 35
gr48	5046	8058	4 460,17	1.09e-3	289217	197	72 : 00 : 00
gr96	55209	121079	51 885,99	1.16e-2	42664	451	72 : 00 : 00
hk48	11461	16491	10 520,26	1.02e-3	257800	197	72 : 00 : 00
kroA100	21282	47891	20 403,88	1.24e-2	38581	529	72 : 00 : 00
kroA150	26524	71197	54 112,35	1.90e-1	10528	1592	72 : 00 : 00
kroB100	22141	58529	20 638,06	1.08e-2	40757	557	72 : 00 : 00
kroB150	26130	69717	46 897,55	1.57e-1	11642	1592	72 : 00 : 00
kroC100	20749	54643	20 017,67	1.15e-2	40813	529	72 : 00 : 00
kroD100	21294	46518	19 866,95	1.48e-2	40801	491	72 : 00 : 00
kroE100	22068	52316	20 819,92	1.77e-2	32288	464	72 : 00 : 00
lin105	14379	38417	13 473,25	1.67e-2	29375	506	72 : 00 : 00
pr107	44303	130840	40 106,77	2.44e-2	24450	512	72 : 00 : 00
pr124	59030	172577	87 856,31	5.25e-2	15894	885	72 : 00 : 00
pr136	96772	240538	151 852,46	1.27e-1	12598	1183	72 : 00 : 00
pr144	58537	216827	172 290,76	2.21e-1	10624	1383	72 : 00 : 00
pr76	108159	195718	87 972,52	6.71e-3	68843	287	72 : 00 : 00
rat99	1211	2365	1 217,21	1.40e-2	33688	529	72 : 00 : 00
rd100	7910	14948	7 460,67	1.03e-2	44344	491	72 : 00 : 00
st70	675	1153	597,10	4.59e-3	114035	267	72 : 00 : 00
swiss42	1273	1725	1 115,10	7.29e-4	416682	262	72 : 00 : 00
ulysses16	6859	7650	5 968,50	8.01e-4	1000000	206	10 : 47 : 06
ulysses22	7013	8865	5 999,69	1.03e-3	1000000	198	26 : 07 : 38

Table SM7: Primal–dual errors from the MaxCut SDP experiment with GSET Benchmark.

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
Data Name Size (n)	err1				
	err2				
	err3				
	err4				
	err5				
	err6				
	time (dd:hh:mm:ss)				
G1 800	9.2240e-04	1.6451e-04	5.6194e-13	9.3543e-04	1.0055e-04
	0	0	0	1.6028e-14	0
	0	7.4893e-05	6.3556e-10	2.3076e-04	0
	8.8644e-06	0	0	0	2.2546e-06
	4.3895e-05	-9.0090e-05	1.3200e-04	-5.8070e-04	-6.5269e-08
	1.1514e-04	3.0968e-05	1.3200e-04	0	5.1987e-05
	00:00:00:20	00:00:00:08	00:00:00:06	00:00:00:56	00:00:00:35
G2 800	9.6153e-04	1.6864e-04	6.3048e-13	2.0071e-05	9.7060e-05
	0	0	0	1.4523e-14	0
	0	7.6757e-05	6.5615e-10	3.4207e-04	0
	6.8492e-06	0	0	0	2.1165e-06
	-4.9289e-05	-9.2305e-05	1.1758e-04	-5.2680e-04	-5.9330e-08
	1.5580e-04	3.1789e-05	1.1758e-04	0	5.0184e-05
	00:00:00:19	00:00:00:08	00:00:00:06	00:00:00:59	00:00:00:34
G3 800	9.6575e-04	2.4802e-04	5.0941e-13	2.8275e-04	8.5287e-05
	0	0	0	1.7220e-14	0
	0	1.1287e-04	6.3521e-10	4.6705e-04	0
	1.4693e-05	0	0	0	2.0824e-06
	2.2749e-04	-1.3581e-04	1.6284e-04	1.3902e-03	-5.8110e-08
	1.2655e-04	4.6700e-05	1.6284e-04	-3.3075e-15	4.4091e-05
	00:00:00:15	00:00:00:08	00:00:00:07	00:00:01:14	00:00:00:22
G4 800	8.9262e-04	2.5869e-04	5.5903e-13	1.8571e-04	9.9246e-05
	0	0	0	2.9142e-14	0
	0	1.1772e-04	6.6670e-10	2.0840e-04	0
	2.3408e-06	0	0	0	2.7813e-06
	2.0978e-04	-1.4133e-04	1.3819e-04	2.6590e-04	-1.3404e-07
	2.3739e-04	4.8998e-05	1.3819e-04	0	5.1241e-05
	00:00:00:18	00:00:00:08	00:00:00:07	00:00:01:03	00:00:00:27
G5 800	9.6447e-04	1.6316e-04	5.4488e-13	6.7925e-04	1.2531e-04
	0	0	0	1.3025e-14	0
	0	7.4251e-05	6.2467e-10	8.3802e-05	0
	1.9874e-06	0	0	0	3.1253e-06
	1.0526e-04	-8.9228e-05	1.1856e-04	1.6789e-04	-8.5141e-08
	2.1343e-04	3.0827e-05	1.1856e-04	0	6.4780e-05
	00:00:00:15	00:00:00:08	00:00:00:06	00:00:01:16	00:00:00:36

Continued on the next page

Table SM7: Primal–dual errors from the MaxCut SDP experiment with GSET Benchmark (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G6 800	9.3303e-04	1.0866e-03	5.5518e-13	7.5832e-05	1.2792e-04
	0	0	0	1.5318e-14	0
	0	1.4492e-03	2.8708e-09	7.9668e-04	0
	3.1301e-05	0	0	0	2.1511e-05
	2.1382e-04	-3.3339e-05	7.7761e-04	-1.3520e-03	-4.4192e-07
	1.1763e-04	1.4312e-03	7.7761e-04	3.9961e-15	6.5767e-05
	00:00:00:24	00:00:00:09	00:00:00:09	00:00:00:55	00:00:00:31
G7 800	8.5782e-04	1.2372e-03	5.5790e-13	6.0123e-05	9.4592e-05
	0	0	0	2.6558e-14	0
	0	1.5941e-03	2.5753e-09	7.8940e-04	0
	3.1960e-05	0	0	0	1.3177e-05
	-3.7838e-05	3.7809e-05	9.3721e-04	-1.4522e-03	-2.7617e-07
	6.9057e-05	1.7315e-03	9.3721e-04	4.2082e-15	4.8682e-05
	00:00:00:26	00:00:00:07	00:00:00:06	00:00:00:56	00:00:00:24
G8 800	9.3260e-04	1.0938e-03	5.5102e-13	1.7077e-04	1.2127e-04
	0	0	0	1.4665e-14	0
	0	1.4105e-03	2.4743e-09	7.9436e-04	0
	4.1592e-05	0	0	0	2.0611e-05
	-3.1418e-05	3.7726e-05	8.5330e-04	-1.4241e-03	-4.5087e-07
	-8.7893e-05	1.5324e-03	8.5331e-04	3.0468e-15	6.2313e-05
	00:00:00:23	00:00:00:08	00:00:00:10	00:00:01:01	00:00:00:28
G9 800	9.5542e-04	1.2435e-03	2.2281e-14	1.0045e-04	1.2441e-04
	0	0	0	1.9571e-14	0
	0	3.1113e-04	2.5145e-10	8.5002e-04	0
	5.0016e-05	0	0	0	2.0214e-05
	-1.2117e-04	1.5484e-05	5.3311e-05	-1.6243e-03	-4.2149e-07
	3.3033e-06	1.6621e-03	5.3311e-05	-1.0117e-14	6.3972e-05
	00:00:00:26	00:00:00:07	00:00:00:07	00:00:00:49	00:00:00:27
G10 800	9.6338e-04	1.1629e-03	8.5789e-13	1.0744e-04	3.7687e-04
	0	0	0	8.3036e-14	0
	0	1.5074e-03	2.7911e-09	4.4804e-04	0
	3.6980e-05	0	0	0	5.6003e-05
	7.1590e-05	3.8077e-05	6.5403e-04	3.4374e-04	-1.2148e-06
	-1.0780e-05	1.6306e-03	6.5403e-04	-2.8520e-14	1.9384e-04
	00:00:00:23	00:00:00:07	00:00:00:06	00:00:01:06	00:00:00:21
G11 800	9.4272e-04	2.5133e-15	1.0470e-13	2.0571e-05	2.2880e-04
	0	0	0	1.1073e-14	0
	0	1.9802e-03	5.0100e-10	2.2129e-04	0
	9.4851e-06	0	0	0	3.1321e-05
	6.8661e-04	-6.8224e-05	3.2120e-04	2.3467e-05	-7.2874e-07
	5.9825e-04	4.6042e-03	3.2120e-04	9.4460e-15	1.1762e-04
	00:00:00:17	00:00:00:06	00:00:00:04	00:00:01:03	00:00:00:18

Continued on the next page

Table SM7: Primal–dual errors from the MaxCut SDP experiment with GSET Benchmark (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G12 800	9.6198e-04	1.8094e-15	3.7075e-14	1.7431e-05	2.3199e-04
	0	0	0	1.7518e-14	0
	0	2.3384e-03	5.0100e-10	1.5710e-04	0
	4.1334e-05	0	0	0	3.2627e-05
	2.5989e-04	7.0095e-06	4.1058e-04	-6.5189e-05	-7.3963e-07
	3.2035e-04	6.2413e-03	4.1058e-04	1.7019e-14	1.1926e-04
	00:00:00:16	00:00:00:09	00:00:00:09	00:00:01:16	00:00:00:18
G13 800	9.0704e-04	1.6417e-15	1.9662e-14	2.4604e-04	2.3859e-04
	0	0	0	2.3089e-14	0
	0	3.3393e-03	7.4695e-10	2.2677e-04	0
	4.6549e-05	0	0	0	3.3410e-05
	1.9049e-04	-1.2000e-04	1.6121e-04	-5.2514e-04	-7.2136e-07
	3.1058e-04	7.9402e-03	1.6121e-04	1.4203e-14	1.2269e-04
	00:00:00:11	00:00:00:09	00:00:00:08	00:00:00:55	00:00:00:20
G14 800	8.7483e-04	6.8671e-04	0	2.7915e-05	1.1104e-04
	0	0	0	2.9734e-14	0
	0	4.2273e-04	1.6712e-10	3.0881e-04	0
	2.1202e-05	0	0	0	2.9799e-05
	3.8840e-04	-3.4869e-04	6.5529e-04	-7.4014e-04	-8.6246e-07
	4.1983e-04	2.4403e-04	6.5529e-04	0	5.6610e-05
	00:00:00:16	00:00:00:08	00:00:00:06	00:00:01:27	00:00:00:23
G15 800	9.4137e-04	2.0233e-04	1.1360e-15	8.3307e-04	2.8862e-04
	0	0	0	5.5046e-14	0
	0	1.2268e-04	1.2565e-11	9.7876e-05	0
	3.8869e-05	0	0	0	8.9568e-05
	3.0556e-04	-1.0261e-04	4.5321e-04	-4.7425e-04	-2.9107e-06
	2.8753e-04	7.2245e-05	4.5321e-04	1.0909e-14	1.4647e-04
	00:00:00:19	00:00:00:12	00:00:00:06	00:00:02:16	00:00:00:23
G16 800	9.4043e-04	6.7439e-04	2.7704e-15	1.7499e-04	1.8725e-04
	0	0	0	9.7405e-14	0
	0	4.1500e-04	1.6493e-10	2.1924e-04	0
	2.2862e-05	0	0	0	4.7086e-05
	3.7868e-04	-3.4259e-04	6.8970e-04	-7.9071e-04	-1.5078e-06
	4.2238e-04	2.4011e-04	6.8970e-04	-3.4307e-15	9.5411e-05
	00:00:00:16	00:00:00:08	00:00:00:05	00:00:01:12	00:00:00:24
G17 800	9.3256e-04	7.6583e-04	1.1136e-14	8.0914e-04	1.3170e-04
	0	0	0	6.2389e-14	0
	0	4.6857e-04	1.7499e-10	2.2869e-04	0
	2.9837e-05	0	0	0	3.7355e-05
	3.4252e-04	-3.8909e-04	6.2508e-04	-9.5150e-04	-1.3348e-06
	3.9942e-04	2.7278e-04	6.2508e-04	0	6.6833e-05
	00:00:00:15	00:00:00:08	00:00:00:07	00:00:01:48	00:00:00:25

Continued on the next page

Table SM7: Primal–dual errors from the MaxCut SDP experiment with GSET Benchmark (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G18 800	9.4244e-04	3.8563e-03	1.6886e-15	2.8812e-05	8.5070e-05
	0	0	0	1.8694e-13	0
	0	1.1277e-03	6.0189e-11	3.1932e-04	0
	9.0768e-05	0	0	0	1.6922e-05
	-3.7838e-04	-1.1201e-04	3.9986e-04	4.5379e-04	-3.5880e-07
	-2.4600e-04	6.1459e-03	3.9986e-04	6.9407e-14	4.3661e-05
	00:00:00:22	00:00:00:08	00:00:00:06	00:00:01:11	00:00:00:25
G19 800	9.4555e-04	1.5963e-03	1.4078e-15	1.5129e-04	1.3630e-04
	0	0	0	4.0056e-14	0
	0	4.3819e-04	5.7485e-11	6.4481e-04	0
	1.1372e-04	0	0	0	3.6206e-05
	-3.2867e-04	8.5952e-05	7.3135e-04	-1.4201e-03	-7.6293e-07
	-3.3669e-04	2.8085e-03	7.3135e-04	8.2191e-15	6.9764e-05
	00:00:00:25	00:00:00:08	00:00:00:06	00:00:01:19	00:00:00:27
G20 800	9.3800e-04	1.0881e-03	2.5029e-15	4.4890e-04	2.6035e-04
	0	0	0	3.3210e-13	0
	0	2.0899e-03	6.0636e-11	2.3547e-04	0
	9.7939e-05	0	0	0	4.8009e-05
	-1.7534e-04	2.2938e-05	7.0134e-04	-6.7662e-05	-1.0140e-06
	-1.3401e-04	1.9626e-03	7.0134e-04	1.0248e-13	1.3370e-04
	00:00:00:20	00:00:00:08	00:00:00:06	00:00:01:22	00:00:00:27
G21 800	9.2523e-04	4.6484e-03	2.4371e-15	6.9572e-05	1.4658e-04
	0	0	0	2.8251e-14	0
	0	1.3595e-03	6.1223e-11	7.2098e-04	0
	1.0570e-04	0	0	0	3.0686e-05
	-2.7449e-04	1.4565e-04	5.4134e-04	-1.7175e-03	-6.6934e-07
	-3.0912e-04	7.8067e-03	5.4134e-04	3.8137e-14	7.5175e-05
	00:00:00:19	00:00:00:08	00:00:00:07	00:00:01:00	00:00:00:24
G22 2000	9.4106e-04	4.2850e-04	2.8563e-13	7.2831e-05	6.1785e-05
	0	0	0	5.0064e-14	0
	0	5.9683e-04	3.9451e-09	1.6903e-04	0
	6.3403e-06	0	0	0	1.8727e-06
	4.2900e-05	-3.0984e-04	7.3429e-04	-6.1514e-04	-2.4181e-08
	1.5465e-04	1.9025e-04	7.3429e-04	2.6991e-15	3.1558e-05
	00:00:00:37	00:00:01:55	00:00:00:54	00:00:13:00	00:00:08:59
G23 2000	9.4571e-04	4.5428e-04	2.8472e-13	2.1546e-05	4.5267e-04
	0	0	0	4.9131e-14	0
	0	6.3328e-04	4.0224e-09	1.7164e-04	0
	5.7582e-06	0	0	0	1.3276e-05
	8.9313e-05	-3.2834e-04	6.9341e-04	-6.0588e-04	-1.6808e-07
	1.0244e-04	2.0180e-04	6.9342e-04	1.2644e-14	2.3122e-04
	00:00:00:39	00:00:03:05	00:00:01:24	00:00:12:47	00:00:06:57

Continued on the next page

Table SM7: Primal–dual errors from the MaxCut SDP experiment with GSET Benchmark (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G24 2000	9.6919e-04	4.4984e-04	2.9846e-13	1.6084e-05	1.4380e-04
	0	0	0	4.4865e-14	0
	0	6.2656e-04	3.9715e-09	1.9607e-04	0
	4.4040e-06	0	0	0	4.4677e-06
	1.5659e-04	-3.2516e-04	7.0366e-04	-6.9019e-04	-6.4201e-08
	1.0454e-04	1.9981e-04	7.0366e-04	-9.7938e-15	7.3443e-05
	00:00:00:32	00:00:02:05	00:00:00:54	00:00:22:12	00:00:07:04
G25 2000	9.7551e-04	4.2290e-04	2.9191e-13	1.0653e-04	5.6048e-04
	0	0	0	2.3325e-14	0
	0	5.8914e-04	3.9520e-09	8.8316e-04	0
	5.0085e-06	0	0	0	1.8061e-05
	9.7544e-05	-3.0561e-04	7.3658e-04	-2.6564e-03	-2.2954e-07
	9.9098e-05	1.8791e-04	7.3658e-04	0	2.8627e-04
	00:00:00:34	00:00:08:45	00:00:00:53	00:00:29:46	00:00:07:34
G26 2000	9.4268e-04	3.8322e-04	3.0824e-13	2.6009e-05	1.4346e-04
	0	0	0	5.1837e-14	0
	0	5.3371e-04	4.0921e-09	1.8849e-04	0
	5.5424e-06	0	0	0	4.7426e-06
	1.2342e-04	-2.7715e-04	7.7996e-04	-6.7260e-04	-6.8321e-08
	1.3050e-04	1.7010e-04	7.7997e-04	7.0813e-15	7.3263e-05
	00:00:00:43	00:00:02:07	00:00:01:05	00:00:12:41	00:00:07:32
G27 2000	9.6526e-04	2.0634e-03	8.3648e-15	3.4442e-04	2.7095e-04
	0	0	0	4.8075e-14	0
	0	4.7227e-04	9.4357e-09	5.7485e-04	0
	1.5259e-05	0	0	0	2.8739e-05
	3.5349e-05	1.0351e-05	2.5244e-04	2.0859e-05	-3.8409e-07
	4.7595e-05	3.0249e-03	2.5245e-04	1.7521e-13	1.3810e-04
	00:00:00:48	00:00:01:56	00:00:01:12	00:00:16:57	00:00:05:47
G28 2000	9.7225e-04	1.9942e-03	1.1612e-14	2.9845e-04	2.3205e-04
	0	0	0	4.8939e-14	0
	0	4.6056e-04	9.6511e-09	6.0724e-04	0
	1.8238e-05	0	0	0	2.5077e-05
	1.3091e-04	2.5565e-05	2.7169e-04	2.5404e-05	-3.3463e-07
	1.0975e-04	2.9542e-03	2.7169e-04	-2.1276e-14	1.1827e-04
	00:00:01:03	00:00:02:53	00:00:01:13	00:00:18:03	00:00:06:29
G29 2000	9.4325e-04	2.1648e-03	1.0602e-14	5.1373e-05	2.3064e-04
	0	0	0	7.7832e-14	0
	0	4.8413e-04	9.5579e-09	1.3324e-04	0
	1.1656e-05	0	0	0	2.2448e-05
	7.8858e-05	-2.1517e-05	2.6210e-04	5.2998e-05	-2.8169e-07
	1.0417e-04	3.1277e-03	2.6210e-04	6.1585e-14	1.1760e-04
	00:00:01:00	00:00:02:53	00:00:01:12	00:00:24:10	00:00:04:31

Continued on the next page

Table SM7: Primal–dual errors from the MaxCut SDP experiment with GSET Benchmark (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G30 2000	9.5286e-04	2.4869e-03	8.3719e-15	2.4693e-04	2.4866e-04
	0	0	0	1.0288e-13	0
	0	5.6779e-04	9.4828e-09	2.4961e-04	0
	1.7223e-05	0	0	0	2.6003e-05
	-4.1280e-05	-3.0774e-05	2.8462e-04	6.6465e-05	-3.2359e-07
	1.3337e-05	3.5902e-03	2.8463e-04	2.7779e-13	1.2677e-04
	00:00:00:49	00:00:01:53	00:00:01:48	00:00:24:38	00:00:05:44
G31 2000	9.6727e-04	1.7409e-03	1.1184e-14	2.8342e-05	2.1105e-04
	0	0	0	9.7757e-14	0
	0	4.0157e-04	3.3898e-09	2.4348e-04	0
	1.8430e-05	0	0	0	2.2272e-05
	-3.6136e-05	1.7021e-05	2.4462e-04	-1.1642e-04	-2.8744e-07
	2.5746e-06	2.5715e-03	2.4462e-04	-2.6184e-13	1.0759e-04
	00:00:00:57	00:00:02:36	00:00:01:49	00:00:33:26	00:00:05:48
G32 2000	9.5091e-04	2.7476e-15	5.4729e-14	3.9317e-04	3.6993e-04
	0	0	0	1.1256e-14	0
	0	2.7594e-03	4.5010e-09	2.8428e-04	0
	6.8408e-06	0	0	0	3.2460e-05
	3.4139e-04	-2.7464e-05	6.0739e-04	-4.6221e-04	-4.6052e-07
	3.8942e-04	6.9408e-03	6.0739e-04	-3.4964e-15	1.8858e-04
	00:00:00:38	00:00:01:41	00:00:00:37	00:00:19:32	00:00:10:27
G33 2000	9.4110e-04	2.5624e-15	3.1971e-14	2.2252e-04	3.8940e-04
	0	0	0	1.0459e-14	0
	0	2.6593e-03	4.5010e-09	3.8443e-04	0
	1.1287e-05	0	0	0	3.4800e-05
	3.2243e-04	3.3035e-05	6.0105e-04	-7.4134e-04	-4.9724e-07
	3.6409e-04	6.8252e-03	6.0106e-04	-3.9643e-15	1.9849e-04
	00:00:00:29	00:00:01:40	00:00:00:31	00:00:16:18	00:00:10:27
G34 2000	9.6895e-04	2.2269e-15	1.0497e-15	8.0679e-05	1.2211e-04
	0	0	0	1.5713e-14	0
	0	2.8937e-03	4.5010e-09	2.6039e-04	0
	1.1397e-05	0	0	0	1.0901e-05
	2.9852e-04	6.0226e-05	8.6747e-04	-5.4601e-04	-1.5148e-07
	3.5070e-04	7.6479e-03	8.6747e-04	-7.3622e-15	6.2250e-05
	00:00:00:29	00:00:01:52	00:00:00:35	00:00:15:19	00:00:10:21
G35 2000	8.9543e-04	4.3815e-04	2.0904e-15	6.0172e-05	1.5840e-04
	0	0	0	6.7384e-14	0
	0	2.5141e-04	1.1770e-11	1.8310e-04	0
	1.4946e-05	0	0	0	3.7964e-05
	3.8998e-04	-2.1941e-04	1.3746e-04	4.7801e-04	-7.6303e-07
	4.6150e-04	1.5373e-04	1.3746e-04	1.6827e-15	8.0201e-05
	00:00:00:50	00:00:02:16	00:00:00:53	00:00:18:38	00:00:13:48

Continued on the next page

Table SM7: Primal-dual errors from the MaxCut SDP experiment with GSET Benchmark (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G36 2000	6.8005e-04	7.4971e-04	1.6344e-15	4.2775e-05	1.9278e-04
	0	0	0	1.7356e-13	0
	0	4.3186e-04	1.0779e-11	2.3336e-04	0
	1.1880e-05	0	0	0	7.3947e-05
	4.9472e-04	-3.7549e-04	5.9490e-04	4.9405e-04	-3.0945e-06
	5.0098e-04	2.6313e-04	5.9490e-04	-1.8953e-14	9.5444e-05
	00:00:01:19	00:00:02:23	00:00:00:52	00:00:19:36	00:00:12:52
G37 2000	7.2413e-04	5.7490e-04	2.1943e-15	3.0272e-04	2.2910e-04
	0	0	0	1.1485e-13	0
	0	3.2790e-04	1.1286e-11	2.0440e-04	0
	1.1905e-05	0	0	0	7.5591e-05
	5.5250e-04	-2.8794e-04	2.3988e-04	7.1309e-04	-2.6385e-06
	5.6725e-04	2.0161e-04	2.3988e-04	0	1.1446e-04
	00:00:01:05	00:00:02:37	00:00:00:53	00:00:25:50	00:00:09:27
G38 2000	6.9024e-04	5.9235e-04	1.3810e-15	5.2776e-04	1.1645e-04
	0	0	0	1.4494e-13	0
	0	3.3869e-04	1.0665e-11	1.5038e-04	0
	1.2879e-05	0	0	0	3.3208e-05
	4.9102e-04	-2.9665e-04	4.6436e-04	3.9093e-04	-6.5520e-07
	4.7173e-04	2.0780e-04	4.6436e-04	-8.2949e-15	5.8869e-05
	00:00:01:08	00:00:02:43	00:00:00:49	00:00:25:53	00:00:09:05
G39 2000	9.6781e-04	2.1901e-03	3.1846e-15	5.5948e-05	1.7300e-04
	0	0	0	9.9559e-14	0
	0	4.9220e-04	1.2180e-10	3.8443e-04	0
	4.7806e-05	0	0	0	4.1089e-05
	-1.3034e-04	-1.1277e-05	7.8432e-04	5.2063e-04	-5.2393e-07
	-5.0727e-05	3.6071e-03	7.8432e-04	-9.6073e-15	8.7894e-05
	00:00:01:46	00:00:02:30	00:00:01:03	00:00:24:44	00:00:10:02
G40 2000	9.6841e-04	3.2031e-03	1.6078e-14	5.5471e-04	1.0523e-04
	0	0	0	1.0969e-13	0
	0	7.2687e-04	1.3897e-11	7.4617e-04	0
	5.6218e-05	0	0	0	1.9443e-05
	-1.1338e-04	5.5916e-05	1.4461e-04	1.5733e-03	-2.4104e-07
	-1.0103e-04	5.3231e-03	1.4461e-04	9.3855e-14	5.3539e-05
	00:00:01:40	00:00:01:49	00:00:01:05	00:01:28:45	00:00:10:28
G41 2000	9.4959e-04	2.9907e-03	1.8685e-14	1.4761e-04	3.6781e-04
	0	0	0	7.8040e-14	0
	0	6.7766e-04	3.7649e-11	6.7791e-04	0
	7.0674e-05	0	0	0	5.4200e-05
	-1.5249e-04	8.6809e-06	8.1046e-05	1.2410e-03	-6.9022e-07
	-9.7060e-05	4.9304e-03	8.1046e-05	-1.8967e-15	1.8730e-04
	00:00:01:21	00:00:02:35	00:00:01:38	00:01:30:45	00:00:08:27

Continued on the next page

Table SM7: Primal–dual errors from the MaxCut SDP experiment with GSET Benchmark (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G42 2000	9.7592e-04	2.0139e-03	2.6746e-15	8.3607e-05	1.0746e-04
	0	0	0	1.4588e-13	0
	0	4.6025e-04	3.9803e-11	6.1700e-04	0
	6.1789e-05	0	0	0	2.5823e-05
	-1.2744e-04	-4.2775e-05	7.4182e-04	1.2828e-03	-3.2331e-07
	-1.7927e-04	3.2573e-03	7.4182e-04	1.4930e-13	5.4599e-05
	00:00:01:33	00:00:02:18	00:00:00:52	00:00:18:50	00:00:09:19
G43 1000	9.3475e-04	2.1860e-04	5.9232e-13	4.1395e-04	1.2988e-04
	0	0	0	4.8906e-14	0
	0	3.0691e-04	1.5921e-09	2.1816e-04	0
	1.0919e-05	0	0	0	5.1650e-06
	2.6443e-04	-1.6022e-04	2.1419e-04	3.3428e-04	-9.9318e-08
	1.6891e-04	9.7363e-05	2.1419e-04	-8.3658e-15	6.6887e-05
	00:00:00:14	00:00:00:24	00:00:00:10	00:00:02:00	00:00:00:53
G44 1000	9.5960e-04	1.8448e-04	5.2090e-13	6.6220e-04	9.1836e-05
	0	0	0	5.4161e-14	0
	0	2.5894e-04	1.4687e-09	1.8707e-04	0
	1.1915e-05	0	0	0	3.8575e-06
	1.4291e-04	-1.3529e-04	1.9904e-04	3.7178e-05	-7.2956e-08
	9.0926e-05	8.2099e-05	1.9904e-04	-8.1424e-15	4.7294e-05
	00:00:00:14	00:00:00:16	00:00:00:09	00:00:01:36	00:00:00:56
G45 1000	9.3935e-04	1.7449e-04	5.7531e-13	8.8850e-04	1.2776e-04
	0	0	0	5.3692e-14	0
	0	2.4507e-04	1.5085e-09	2.0479e-04	0
	8.5894e-06	0	0	0	4.9143e-06
	9.6997e-05	-1.2802e-04	2.8376e-04	-1.1397e-04	-9.9553e-08
	1.4242e-04	7.7608e-05	2.8376e-04	9.6044e-15	6.5796e-05
	00:00:00:15	00:00:00:15	00:00:00:10	00:00:01:45	00:00:00:54
G46 1000	9.5919e-04	1.7429e-04	5.7220e-13	1.2383e-04	9.9527e-05
	0	0	0	3.2220e-14	0
	0	2.4488e-04	1.5276e-09	3.0758e-04	0
	1.0933e-05	0	0	0	3.8482e-06
	1.1956e-04	-1.2778e-04	1.5929e-04	7.7609e-04	-7.2593e-08
	8.8945e-05	7.7594e-05	1.5929e-04	-7.0508e-15	5.1261e-05
	00:00:00:15	00:00:00:21	00:00:00:10	00:00:02:19	00:00:00:55
G47 1000	9.5642e-04	2.2700e-04	5.3453e-13	9.5704e-04	1.1095e-04
	0	0	0	5.1509e-14	0
	0	3.1874e-04	1.4995e-09	1.6724e-04	0
	8.0600e-06	0	0	0	4.1527e-06
	2.1827e-04	-1.6626e-04	1.9451e-04	-2.2679e-04	-7.9365e-08
	1.3089e-04	1.0118e-04	1.9451e-04	-9.8917e-15	5.7148e-05
	00:00:00:14	00:00:00:15	00:00:00:15	00:00:02:42	00:00:00:38

Continued on the next page

Table SM7: Primal–dual errors from the MaxCut SDP experiment with GSET Benchmark (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G48 3000	9.6092e-04	1.1491e-14	2.7859e-14	1.7221e-05	4.5896e-08
	0	0	0	1.1588e-14	0
	0	4.7998e-04	1.1881e-07	8.7379e-04	0
	7.2356e-05	0	0	0	2.1067e-09
	-8.1065e-04	-6.6175e-05	1.2623e-04	-1.5367e-03	-9.9465e-09
	-6.5493e-04	1.3226e-04	1.2627e-04	-2.1276e-14	1.3419e-08
	00:00:00:13	00:00:06:07	00:00:01:23	00:01:47:16	00:00:17:05
G49 3000	8.7711e-04	8.6008e-15	4.2404e-13	1.6246e-04	1.1879e-04
	0	0	0	1.2076e-14	0
	0	1.0779e-04	5.4479e-07	2.1347e-04	0
	3.1797e-05	0	0	0	3.7620e-06
	-4.0582e-04	-1.4857e-05	9.4205e-04	-1.8609e-03	-2.0698e-08
	-1.6411e-04	2.9698e-05	9.4220e-04	1.7790e-15	6.0453e-05
	00:00:00:14	00:00:07:37	00:00:01:26	00:00:34:06	00:00:19:11
G50 3000	9.0458e-04	1.1434e-14	4.0840e-13	1.3191e-04	7.0325e-04
	0	0	0	1.7726e-14	0
	0	1.8344e-03	5.3630e-07	6.7731e-04	0
	1.4867e-05	0	0	0	2.2364e-05
	2.8789e-05	-2.5364e-04	6.1843e-04	1.0880e-03	-1.2859e-07
	6.1924e-05	5.0555e-04	6.1859e-04	4.5250e-14	3.5788e-04
	00:00:00:15	00:00:05:51	00:00:01:13	00:01:03:59	00:00:19:15
G51 1000	9.4758e-04	1.8096e-04	0	9.4846e-05	3.2323e-04
	0	0	0	8.4322e-14	0
	0	1.0685e-04	1.2125e-11	6.6263e-05	0
	2.2294e-05	0	0	0	8.7008e-05
	5.4910e-04	-9.1780e-05	4.7349e-04	2.3462e-05	-2.8725e-06
	4.9088e-04	6.3721e-05	4.7349e-04	-5.7231e-15	1.6383e-04
	00:00:00:25	00:00:00:17	00:00:00:13	00:00:02:32	00:00:00:52
G52 1000	9.3422e-04	2.3942e-04	0	3.6346e-05	2.7082e-04
	0	0	0	8.7588e-15	0
	0	1.4337e-04	1.2359e-11	3.3592e-04	0
	2.0451e-05	0	0	0	7.9081e-05
	4.9357e-04	-1.2145e-04	3.8698e-04	-1.9403e-05	-2.8396e-06
	4.8520e-04	8.4252e-05	3.8698e-04	0	1.3683e-04
	00:00:00:22	00:00:00:20	00:00:00:10	00:00:03:19	00:00:00:51
G53 1000	8.9247e-04	2.9963e-04	0	2.7250e-05	7.1987e-05
	0	0	0	2.6300e-14	0
	0	1.7934e-04	1.2351e-11	1.1149e-04	0
	2.0807e-05	0	0	0	1.7338e-05
	4.4112e-04	-1.5194e-04	3.8368e-04	-2.4948e-04	-5.1172e-07
	4.6495e-04	1.0550e-04	3.8368e-04	2.8611e-15	3.6615e-05
	00:00:00:20	00:00:00:20	00:00:00:10	00:00:02:14	00:00:00:45

Continued on the next page

Table SM7: Primal–dual errors from the MaxCut SDP experiment with GSET Benchmark (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G54 1000	8.2442e-04	9.6590e-04	0	6.3149e-04	2.3269e-04
	0	0	0	7.8280e-14	0
	0	5.7862e-04	1.2365e-11	1.2067e-04	0
	1.9272e-05	0	0	0	5.3841e-05
	4.7339e-04	-4.9058e-04	4.3025e-04	-1.4418e-04	-2.0052e-06
	4.2979e-04	3.3945e-04	4.3025e-04	-5.4878e-15	1.1801e-04
	00:00:00:28	00:00:00:17	00:00:00:15	00:00:02:28	00:00:01:17
G55 5000	9.8552e-04	2.4798e-04	2.8764e-15	7.7904e-04	1.9044e-04
	0	0	0	3.4424e-14	0
	0	5.6271e-04	2.5720e-09	1.1896e-04	0
	5.0552e-06	0	0	0	7.9203e-06
	4.1808e-04	-1.4241e-04	7.4066e-05	-4.3644e-04	-5.9389e-08
	4.2793e-04	2.2303e-04	7.4067e-05	2.0123e-14	9.6502e-05
	00:00:05:54	00:00:42:38	00:00:10:21	00:23:04:37	00:03:04:42
G56 5000	9.7842e-04	9.7935e-04	2.5583e-15	8.0068e-05	1.5314e-04
	0	0	0	6.4741e-14	0
	0	2.6628e-03	1.3289e-08	1.2807e-04	0
	8.9039e-06	0	0	0	1.2366e-05
	1.1780e-04	5.5033e-06	5.5448e-04	-4.1735e-04	-9.7127e-08
	1.0618e-04	2.0459e-03	5.5448e-04	7.5565e-15	7.7549e-05
	00:00:03:35	00:00:31:46	00:00:08:57	00:23:11:02	00:03:00:38
G57 5000	9.7494e-04	3.1018e-15	3.7205e-15	3.8694e-04	2.5653e-04
	0	0	0	9.8838e-15	0
	0	3.5240e-03	4.0510e-09	2.9910e-04	0
	4.7070e-06	0	0	0	1.4371e-05
	2.1366e-04	2.1686e-05	5.9892e-04	-4.2607e-04	-1.2404e-07
	2.7322e-04	8.6078e-03	5.9892e-04	-3.4877e-14	1.2994e-04
	00:00:01:40	00:00:23:26	00:00:05:45	00:19:20:50	00:02:30:43
G58 5000	4.0780e-04	2.6793e-04	2.0718e-15	5.3746e-04	1.8343e-04
	0	0	0	6.6770e-14	0
	0	1.4522e-04	9.0746e-12	3.6710e-04	0
	4.9369e-06	0	0	0	6.9383e-05
	4.9332e-04	-1.3301e-04	9.4745e-04	2.4708e-04	-1.9770e-06
	4.9184e-04	9.3117e-05	9.4745e-04	1.7301e-14	9.1031e-05
	00:00:10:32	00:00:47:06	00:00:10:57	00:23:08:17	00:03:35:20
G59 5000	9.8314e-04	2.2738e-03	4.8543e-15	2.3647e-04	2.6866e-04
	0	0	0	1.4684e-13	0
	0	4.0049e-04	1.8306e-11	9.6139e-05	0
	2.6813e-05	0	0	0	2.8799e-05
	-7.4885e-05	-1.5350e-05	4.0915e-04	-1.4136e-04	-2.2324e-07
	3.1707e-05	3.6954e-03	4.0915e-04	9.3195e-15	1.3600e-04
	00:00:09:36	00:00:37:11	00:00:11:56	00:16:04:30	00:03:46:58

Continued on the next page

Table SM7: Primal–dual errors from the MaxCut SDP experiment with GSET Benchmark (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G60 7000	9.4222e-04	2.9488e-04	2.7808e-15	9.3599e-05	4.3747e-04
	0	0	0	4.3064e-14	0
	0	6.7677e-04	4.2801e-09	1.3113e-04	0
	1.9291e-06	0	0	0	1.5191e-05
	7.1170e-05	-1.6815e-04	1.6532e-04	-2.2004e-04	-8.6697e-08
	4.0690e-05	2.6756e-04	1.6533e-04	-4.5629e-14	2.2125e-04
	00:00:22:47	00:01:20:38	00:00:24:05	00:16:47:49	00:08:17:19
G61 7000	7.4498e-04	1.1123e-03	3.0992e-15	2.7651e-04	1.4060e-04
	0	0	0	4.9482e-14	0
	0	3.0721e-03	6.3286e-09	2.7933e-04	0
	5.7313e-06	0	0	0	9.0226e-06
	5.2060e-05	-3.0045e-05	1.8713e-04	-1.1403e-03	-6.6698e-08
	5.4104e-05	2.2545e-03	1.8713e-04	-9.7085e-14	7.1066e-05
	00:00:07:42	00:01:29:46	00:00:23:22	00:16:39:51	00:08:25:12
G62 7000	9.8167e-04	3.2885e-15	1.0568e-15	2.6213e-04	2.8277e-04
	0	0	0	1.4285e-14	0
	0	3.2915e-03	6.7510e-09	1.5326e-04	0
	3.3956e-06	0	0	0	1.3510e-05
	2.2564e-04	3.1229e-05	5.9666e-04	-2.4200e-04	-9.7909e-08
	2.1179e-04	8.1609e-03	5.9666e-04	4.1265e-14	1.4296e-04
	00:00:03:09	00:01:18:18	00:00:14:27	00:10:27:25	00:07:20:45
G63 7000	9.2555e-04	2.9923e-04	2.8620e-15	1.0476e-04	1.6669e-04
	0	0	0	9.4341e-14	0
	0	1.6003e-04	4.9605e-12	3.8707e-04	0
	5.6376e-06	0	0	0	5.4369e-05
	4.8804e-04	-1.4816e-04	4.8130e-04	1.3827e-05	-1.4318e-06
	7.6415e-04	1.0374e-04	4.8130e-04	-8.7045e-15	8.2908e-05
	00:00:17:25	00:02:10:33	00:00:31:58	00:14:22:49	00:10:26:11
G64 7000	9.8308e-04	3.8323e-03	3.6898e-14	8.8455e-04	2.1453e-04
	0	0	0	6.5761e-13	0
	0	6.5016e-04	1.7381e-11	1.3625e-04	0
	2.6912e-05	0	0	0	1.5266e-05
	3.6805e-05	-9.8637e-05	1.4654e-04	-4.6901e-05	-9.7600e-08
	-5.6791e-05	6.0347e-03	1.4654e-04	-3.2217e-13	1.0844e-04
	00:00:17:25	00:01:28:39	00:00:33:48	00:18:37:55	00:08:04:41
G65 8000	9.8854e-04	3.4940e-15	3.8001e-15	5.3185e-04	–
	0	0	0	7.4499e-15	–
	0	3.4543e-03	6.7510e-09	3.1610e-04	–
	3.8799e-06	0	0	0	–
	1.9849e-04	2.8883e-05	8.5582e-04	-5.2525e-04	–
	2.1539e-04	8.3666e-03	8.5582e-04	1.3871e-14	–
	00:00:02:49	00:01:34:04	00:00:49:44	00:22:46:02	–

Continued on the next page

Table SM7: Primal-dual errors from the MaxCut SDP experiment with GSET Benchmark (cont.).

	SketchyCGAL	MoSeK	SDPT3	SDPNAL+	Sedumi
G66 9000	9.8472e-04	3.5663e-15	6.7137e-15	4.6474e-04	-
	0	0	0	1.1480e-14	-
	0	3.4548e-03	6.7510e-09	2.2088e-04	-
	3.8099e-06	0	0	0	-
	1.9103e-04	-2.5418e-05	8.9009e-04	-5.4521e-04	-
	1.9216e-04	8.2468e-03	8.9009e-04	-1.4191e-14	-
	00:00:02:55	00:02:18:24	00:01:15:51	01:07:25:38	-
G67 10000	9.8870e-04	-	1.0890e-15	9.7122e-04	-
	0	-	0	1.5106e-14	-
	0	-	2.0260e-09	4.3451e-05	-
	3.7053e-06	-	0	0	-
	2.5490e-04	-	4.5646e-04	-2.3706e-04	-
	1.8960e-04	-	4.5646e-04	-2.6790e-15	-
	00:00:03:46	-	00:00:45:40	01:20:08:29	-

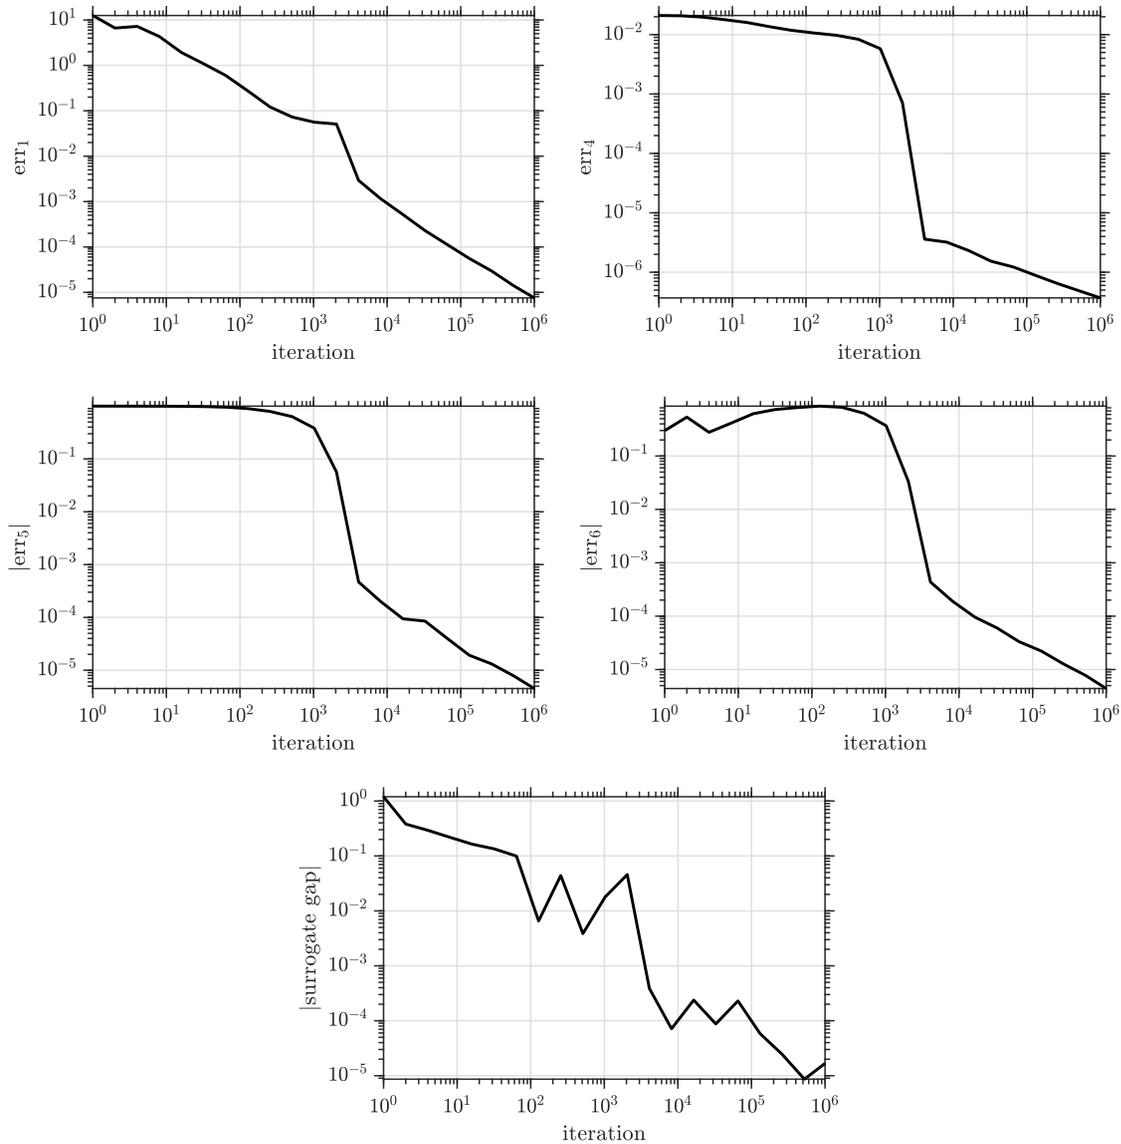


Figure SM7.3. MaxCut SDP: Primal-dual convergence. We solve the MaxCut SDP for the G67 dataset ($n = 10000$) with SketchyCGAL. The subplots show the magnitudes of the DIMACS errors and the surrogate gap of the implicit primal iterate and the dual iterate. We omit err_2 and err_3 since they are zero by construction. See [subsection SM5.3](#).