

An Inexact Projected Gradient Method with Rounding and Lifting by Nonlinear Programming for Solving Rank-One Semidefinite Relaxation of Polynomial Optimization

Heng Yang · Ling Liang
Luca Carlone · Kim-Chuan Toh

October 27, 2021

Abstract We consider solving high-order semidefinite programming (SDP) relaxations of nonconvex polynomial optimization problems (POPs) that often admit degenerate rank-one optimal solutions. Instead of solving the SDP alone, we propose a new algorithmic framework that blends *local search* using the nonconvex POP into *global descent* using the convex SDP. In particular, we first design a globally convergent *inexact* projected gradient method (iPGM) for solving the SDP that serves as the backbone of our framework. We then accelerate iPGM by taking long, but *safeguarded*, rank-one steps generated by fast nonlinear programming algorithms. We prove that the new framework is still globally convergent for solving the SDP. To solve the iPGM subproblem of projecting a given point onto the feasible set of the SDP, we design a two-phase algorithm with phase one using a symmetric Gauss-Seidel based accelerated proximal gradient method (sGS-APG) to generate a good initial point, and phase two using a modified limited-memory BFGS (L-BFGS) method to obtain an accurate solution. We analyze the convergence for both phases and establish a novel global convergence result for the

Heng Yang
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
E-mail: hankyang@mit.edu

Ling Liang
Department of Mathematics
National University of Singapore
E-mail: liang.ling@u.nus.edu

Luca Carlone
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
E-mail: lcarlone@mit.edu

Kim-Chuan Toh
Department of Mathematics and Institute of Operations Research and Analytics
National University of Singapore
E-mail: mattohkc@nus.edu.sg

The first two authors contributed equally.
Code available: <https://github.com/MIT-SPARK/STRIDE>

modified L-BFGS that does not require the objective function to be twice continuously differentiable. We conduct numerical experiments for solving second-order SDP relaxations arising from a diverse set of POPs. Our framework demonstrates state-of-the-art efficiency, scalability, and robustness in solving degenerate rank-one SDPs to high accuracy, even in the presence of millions of equality constraints.

Keywords Semidefinite programming · Polynomial optimization · Inexact projected gradient method · Rank-one solutions · Nonlinear programming · Degeneracy

Mathematics Subject Classification (2010) 90C06 · 90C22 · 90C23 · 90C55

1 Introduction

Let $p, h_1, \dots, h_l \in \mathbb{R}[x]$ be real-valued multivariate polynomials in $x \in \mathbb{R}^d$, we consider the following equality constrained polynomial optimization problem¹

$$\min_{x \in \mathbb{R}^d} \{p(x) \mid h_i(x) = 0, i = 1, \dots, l\}. \quad (\text{POP})$$

Assuming problem (POP) is feasible and bounded below, we denote $-\infty < p^* < \infty$ as the global minimum and x^* as a global minimizer. Finding (p^*, x^*) has applications in various fields of engineering and science [38]. It is, however, well recognized that problem (POP) is NP-hard in general (*e.g.*, it can model the binary constraint $x \in \{+1, -1\}$ via $x^2 = 1$). As a result, solving convex relaxations of (POP), especially semidefinite programming (SDP) relaxations, has been the predominant method for finding (p^*, x^*) . In this paper, we are interested in designing efficient, robust, and scalable algorithms for computing (p^*, x^*) through SDP relaxations.

SDP Relaxations for Polynomial Optimization Problems. Towards this goal, let us first give an overview of the celebrated Lasserre’s moment/sums-of-squares (SOS) semidefinite relaxation hierarchy [37, 55]. Let $\kappa \geq 1$ be an integer such that 2κ is equal or greater than the maximum degree of $p, h_i, i = 1, \dots, l$, and denote $v := [x]_\kappa$ as the standard vector of monomials in x of degree up to κ . We build the *moment matrix* $X := [x]_\kappa [x]_\kappa^\top$ that contains the standard monomials in x of degree up to 2κ . As a result, p and h_i ’s can be written as linear functions in X . For example, choosing $x = (x_1, x_2) \in \mathbb{R}^2$ and $\kappa = 2$ leads to the following moment matrix that contains all monomials of x up to degree 4:

$$X = \begin{bmatrix} 1 & x_1 & x_2 & x_1^2 & x_1x_2 & x_2^2 \\ x_1 & x_1^2 & x_1x_2 & x_1^3 & x_1^2x_2 & x_1x_2^2 \\ x_2 & x_1x_2 & x_2^2 & x_1^2x_2 & x_1x_2^2 & x_2^3 \\ x_1^2 & x_1^3 & x_1^2x_2 & x_1^4 & x_1^3x_2 & x_1^2x_2^2 \\ x_1x_2 & x_1^2x_2 & x_1x_2^2 & x_1^3x_2 & x_1^2x_2^2 & x_1x_2^3 \\ x_2^2 & x_1x_2^2 & x_2^3 & x_1^2x_2^2 & x_1x_2^3 & x_2^4 \end{bmatrix}. \quad (1)$$

We then consider two types of necessary linear constraints that can be imposed on X : (i) the entries of X are not linearly independent, in fact the same monomial could appear in multiple locations of X (*e.g.*, the monomial x_1x_2 in (1)); (ii)

¹ Although our framework can be generalized to problems with inequality constraints, here we focus on equality constraints to simplify the presentation.

each constraint h_i in the original (POP) generates a set of equalities on X of the form $h_i(x)[x]_{2\kappa-\deg(h_i)} = 0$, where $\deg(h_i)$ is the degree of h_i (i.e., if $h_i = 0$, then $h_i \cdot x_1 = 0$, $h_i \cdot (x_1x_2) = 0$ and so on). Since X is positive semidefinite by construction, one can therefore write the resulting semidefinite relaxation in standard *primal* form as

$$\min_{X \in \mathbb{S}^n} \{ \langle C, X \rangle \mid \mathcal{A}(X) = b, X \succeq 0 \}, \quad (\text{P})$$

where $n := \bar{d}_\kappa \triangleq \binom{d+\kappa}{\kappa}$ is the dimension of $[x]_\kappa$, $b \in \mathbb{R}^m$, $\mathcal{A} : \mathbb{S}^n \rightarrow \mathbb{R}^m$ is a linear map $\mathcal{A}(X) \triangleq (\langle A_i, X \rangle)_{i=1}^m$ with $A_i \in \mathbb{S}^n$, that is onto and collects all independent linear constraints generated by the relaxation of $X = [x]_\kappa [x]_\kappa^\top$ to $X \succeq 0$. Let $\mathcal{A}^* : \mathbb{R}^m \rightarrow \mathbb{S}^n$ be the adjoint of \mathcal{A} defined as $\mathcal{A}^*y \triangleq \sum_{i=1}^m y_i A_i$, then the Lagrangian dual of problem (P) reads

$$\max_{y \in \mathbb{R}^m, S \in \mathbb{S}^n} \{ \langle b, y \rangle \mid \mathcal{A}^*y + S = C, S \succeq 0 \}, \quad (\text{D})$$

and admits an interpretation using sums-of-squares polynomials [55]. SDP relaxations (P)-(D) correspond to the so-called *dense* hierarchy, while many *sparse* variants have been proposed to exploit the sparsity structure of p and h_i 's to generate SDP relaxations with multiple blocks and smaller sizes, see [67, 69, 70] and references therein. We remark that the algorithms developed in this paper can be extended to multiple blocks in a straightforward way to solve sparse relaxations.

Throughout this paper, we assume that strong duality holds for (P) and (D), and both (P) and (D) admit at least one solution.² We denote X^* and (y^*, S^*) as an optimal solution for (P) and (D), respectively. Then, the KKT condition for (P) and (D) given as

$$\mathcal{A}(X) - b = 0, \quad \mathcal{A}^*y + S - C = 0, \quad \langle X, S \rangle = 0, \quad X, S \succeq 0, \quad (2)$$

has (X^*, y^*, S^*) as one of its solutions. Moreover, it follows that $p^* \geq f_P^* = f_D^*$, where f_P^* and f_D^* are the optimal values of (P) and (D), respectively. The SDP relaxation is said to be *tight* if $p^* = f_P^*$. In such cases, for any global minimizer x^* of (POP), the rank one lifting $X = [x^*]_\kappa [x^*]_\kappa^\top$ is a minimizer of (P), and any rank one optimal solution X^* of (P) corresponds to a global minimizer of (POP). A special case of the relaxation hierarchy is Shor's relaxation ($\kappa = 1$) for quadratically constrained quadratical programming problems [63, 42], of which applications and analyses have been extensively studied [18, 1, 15, 61, 58, 21, 68, 17]. Although $\kappa = 1$ is certainly an interesting case, it cannot handle p and h_i 's with degrees above 2, and it is known to be not tight for many problems such as MAXCUT [29]. Moreover, it typically leads to SDPs with small m that can often be handled well by existing SDP solvers [66]. Therefore, in this paper, we mainly focus on solving *high-order* ($\kappa \geq 2$) relaxations that are more powerful in attaining tightness, as we describe below.

In the seminal work [37], Lasserre proved that f_P^* *asymptotically* approaches p^* as κ increases to infinity [37, Theorem 4.2]. Later on, several authors showed that tightness indeed happens for a finite κ [39, 52, 53]. Notably, Nie proved that, under the archimedean condition, if constraint qualification, strict complementarity and

² Particularly, if a redundant ball constraint is included in the relaxation, then dual Slater and strong duality hold for Lasserre's moment relaxations of all orders [34, Theorem 1].

second-order sufficient condition hold at every global minimizer of (POP), then the hierarchy converges at a finite κ [53, Theorem 1.1]. What is even more encouraging is that, empirically, for many important (POP) instances, tightness can be attained at a very low relaxation order such as $\kappa = 2, 3, 4$. For instance, Lasserre [36] showed that $\kappa = 2$ attains tightness for a set of 50 randomly generated MAXCUT problems; In the experiments of Henrion and Lasserre [31], tightness holds at a small κ for most of the (POP) problems in the literature; Doherty, Parrilo and Spedalieri [25] demonstrated that the second-order SOS relaxation is sufficient to decide quantum separability in all bound entangled states found in the literature of dimensions up to 6 by 6; Cifuentes [20] designed and proved that a sparse variant of the second-order moment relaxation is guaranteed to be tight for the structured total least squares problem under a low noise assumption; Yang and Carlone applied a sparse second-order moment relaxation to globally solve a broad family of nonconvex computer vision problems [74, 76, 75].

Computational Challenges in Solving High-order Relaxations. The surging applications make it imperative to design computational tools for solving high-order (tight) SDP relaxations. However, high-order relaxations pose notorious challenges to existing SDP solvers. For this reason, in all applications above, the experiments were performed on solving (POP) problems of very small size, *e.g.*, the original (POP) has d up to 20 for dense relaxations. The computational challenges mainly come from two aspects. In the first place, in stark contrast to Shor’s semidefinite relaxation, high order relaxations lead to SDPs with a large number of linear constraints, even if the dimension of the original (POP) is small. Taking the binary quadratic programming as an example (*i.e.*, $x_i \in \{+1, -1\}$ and p is quadratic in (POP)), when $d = 60$, the number of equality constraints for $\kappa = 2$ is $m = 1,266,971$. In the second place, the resulting SDP is not only large, but also highly degenerate. Specifically, when the relaxation (P) is tight and admits rank-one optimal solutions, it necessarily fails the primal nondegeneracy condition [4], which is crucial for ensuring numerical stability and fast convergence of existing algorithms [5, 82]. Due to these two challenges, our experiments show that *no existing solver* can consistently solve semidefinite relaxations with rank-one solutions to a desired accuracy when m is larger than 500,000. In particular, interior point methods (IPM), such as SDPT3 [66] and MOSEK [7], can only handle up to $m = 50,000$ before they run out of memory on an ordinary workstation.³ First-order methods based on ADMM and conditional gradient method, such as CDCS [83] and SketchyCGAL [79], converge very slowly and cannot attain even modest accuracy for challenging instances. The best performing existing solver is SDPNAL+ [78], which employs an augmented Lagrangian framework where the inner problem is solved inexactly by a semi-smooth Newton method with a conjugate gradient method (SSNCG), and hence has very good scalability. The problem with SDPNAL+ is that, in the presence of large m and degeneracy, solving the inexact SSN step involves solving a large and singular linear system of size $m \times m$, and CG becomes incapable of computing the search direction with sufficient accuracy. A tempting alternative approach, since (P) has rank-one optimal solutions, is to apply the low-rank factorization method of Burer and Monteiro [16], *i.e.*, solving

³ When chordal sparsity exists in the SDP such that a large positive semidefinite constraint can be decomposed into multiple smaller ones, IPMs can be more scalable [27, 81]. However, for problems considered in this paper, there is no chordal sparsity.

the nonlinear optimization $\min_{V \in \mathbb{R}^{n \times r}} \{\langle C, VV^\top \rangle \mid \mathcal{A}(VV^\top) = b\}$ for some small r . We note that since $\nabla_V \langle A_i, VV^\top \rangle = 2A_i V$, if $\text{rank}(V^*) = 1$ at the optimal solution V^* , then $\text{rank}(2[A_1 V^*, \dots, A_m V^*]) \leq n \ll m$ no matter how large r is. Therefore, V^* fails the linear independence constraint qualification (LICQ) and it may not be a KKT point, making it pessimistic for nonlinear programming solvers to find. In fact, successful applications of B-M factorization require $m \approx n$ so that the dual multipliers can be obtained in closed-form from nonlinear programming solutions [58, 57, 14]. Nevertheless, exploiting low rankness and nonlinear programming tools is a great source of inspiration, and as we will show, our solver also exploits similar insights but in a fashion that is not subject to the large m and degeneracy of problem (P).

Our Contribution. In this paper, we contribute the first solver that can solve high-order SDP relaxations of (POP) to high accuracy in spite of large m and degeneracy. Our solver employs a globally convergent *inexact projected gradient method* (iPGM) as the backbone for solving the convex SDP (P). Despite having favorable global convergence, the iPGM steps are typically *short* and lead to slow convergence, particularly in the presence of degeneracy. Therefore, we blend short iPGM steps with *long* rank-one steps generated by fast nonlinear programming (NLP) algorithms. The intuition is that, once we have achieved a sufficient amount of descent via iPGM, we can switch to NLP to perform local refinement to “jump” to a nearby rank-one point for rapid convergence. Due to this reason, we name our framework *SpecTrahedRon. Inexact projected gradient Descent along vErtices* (STRIDE), which essentially strides along the rank-one vertices of the spectrahedron until the global minimum of the SDP is reached.

Before presenting the mathematical details of STRIDE in a general setup, let us first introduce an informal version of STRIDE by running an illustrative example.

1.1 A Univariate Example

Consider minimizing a quartic polynomial subject to a single quartic equality constraint

$$\min_{x \in \mathbb{R}} \left\{ p(x) := x^4 + \frac{2}{3}x^3 - 8x^2 - 8x \mid h(x) := (x^2 - 4)(x^2 - 1) = 0 \right\}, \quad (3)$$

where it is obvious that x can only take four real values $\{+2, -2, +1, -1\}$ due to the constraint $h(x)$, among which $x^* = 2$ attains the global minimum $p^* = -\frac{80}{3} \approx -26.67$. A plot of $p(x)$ is shown in Fig. 1(a). However, let us discard our “global” view of $p(x)$ and design an algorithm to numerically obtain (x^*, p^*) through its SDP relaxation. Towards this, according to our brief introduction of Lasserre’s hierarchy, let us denote $v(x) := [1, x, x^2]^\top$ to be the vector of monomials in x of degree up to 2, and build the moment matrix $X := vv^\top$. X has monomials of x with degree up to 4, which we denote as $z_i := x^i, i = 1, \dots, 4$. With $h(x) = 0$, we further have $x^4 = 5x^2 - 4$, i.e., $z_4 = 5z_2 - 4$. Now we can write the cost function $p(x)$ as a linear function of z : $f(z) = z_4 + \frac{2}{3}z_3 - 8z_2 - 8z_1 = \frac{2}{3}z_3 - 3z_2 - 8z_1 - 4$. With this construction, the second-order relaxation of (3) reads

$$\min_{z \in \mathbb{R}^3} \left\{ f(z) := \frac{2}{3}z_3 - 3z_2 - 8z_1 - 4 \mid X(z) := \begin{bmatrix} 1 & z_1 & z_2 \\ z_1 & z_2 & z_3 \\ z_2 & z_3 & 5z_2 - 4 \end{bmatrix} \succeq 0 \right\}. \quad (4)$$

Problem (4) allows us to explicitly visualize its feasible set, which is in general called a *spectrahedron*. Using the fact that $X(z) \succeq 0$ if and only if the coefficients of its characteristic polynomial weakly alternate in sign [12, Proposition A.1], we plot the spectrahedron in Fig. 1(b) using Mathematica [32]. The four rank-one vertices are annotated with associated coordinates (z_1, z_2, z_3) , where $z_1 = x \in \{+2, -2, +1, -1\}$ corresponds to the four values that x can take. The negative gradient direction $-\nabla f(z) = [8, 3, -\frac{2}{3}]^\top$ is annotated by blue solid arrows. The hyperplane defined by $f(z)$ taking a constant value is shown as a green solid line in Fig. 1(b). Clearly, the rank-one vertex $z^* = (2, 4, 8)$ (“★” in (b)) attains the minimum of the SDP (4), which corresponds to $x^* = 2$ (“★” in (a)) in POP (3).

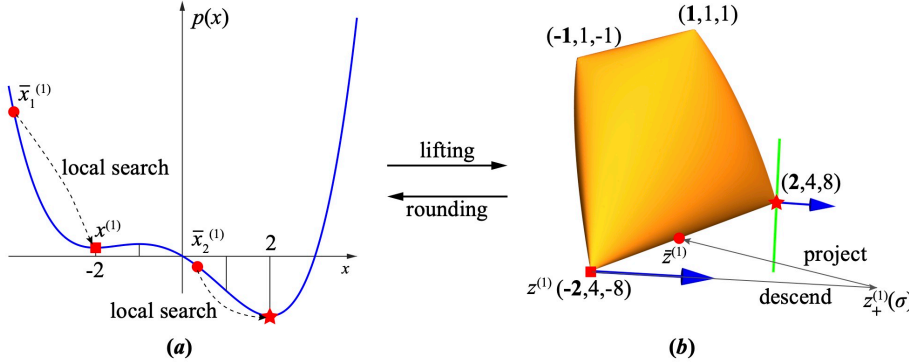


Fig. 1 Overview of our algorithm using a univariate example. (a) Plot of the univariate polynomial in (3). (b) Visualization of the feasible set of the SDP (4) arising from a second-order moment relaxation of (3). “■”: suboptimal local minimum of the POP and its corresponding (lifted) vertex in the SDP. “★”: globally optimal solutions of POP and SDP. “●”: A descent point in SDP and corresponding rounded hypotheses in POP. Blue arrows in (b): direction along the negative gradient of the SDP objective function. Green line in (b): a hyperplane of constant objective going through the optimal solution.

We now state our numerical algorithm that finds both x^* and z^* .

Initialization. We use a standard nonlinear programming (NLP) solver (*e.g.*, an interior point method interfaced via `fmincon` in Matlab) to solve problem (3). We initialize at $x^{(0)} = -10$, and NLP converges to $x^{(1)} = -2$, with a cost $p(x^{(1)}) \approx -5.33$ (“■” in Fig. 1(a)). Note that $x^{(1)} = -2$ is a strict local minimum because choosing a Lagrangian multiplier of “0” for the constraint h leads to $\nabla p(-2) = 0$, $\nabla^2 p(-2) = 24$, thereby satisfying first-order optimality and second-order sufficiency [54].

Lift and Descend. The initialization step converges to a suboptimal solution. Unfortunately, it is very challenging, if not impossible, for NLP to be aware of this suboptimality, not to mention about escaping it. The SDP relaxation now comes into play. By lifting $x^{(1)}$ to $X = v(x^{(1)})v(x^{(1)})^\top$ according to the SDP relaxation, we obtain the suboptimal rank-one vertex “■” in Fig. 1(b) with $z^{(1)} = [-2, 4, -8]^\top$. Although it is challenging to descend from $x^{(1)}$ in POP (3), descending from $z^{(1)}$ in the convex SDP (4) is relatively easy. We first take a step along the negative direction of the gradient to arrive at $z_+^{(1)}(\sigma) := z^{(1)} - \sigma \nabla f(z^{(1)}) = [-2 + 8\sigma, 4 +$

$3\sigma, -8 - \frac{2}{3}\sigma]^\top$ for a given step size $\sigma > 0$, and then project $z_+^{(1)}$ back to the spectrahedron in Fig. 1(b), denoted as $\bar{z}^{(1)}$. Note that the projection step solves the problem $\min_{z \in \mathbb{R}^3} \{\|X(z) - X(z_+^{(1)})\|^2 \mid X(z) \succeq 0\}$, which is itself a quadratic SDP. For this small-scale example, we can compute the projection exactly (using *e.g.*, IPMs). Taking $\sigma = 5$ yields $z_+^{(1)} = [38, 19, -\frac{34}{3}]^\top$ and $\bar{z}^{(1)} = [-0.35, 3.99, -1.67]^\top$, and the projection $\bar{z}^{(1)}$ is plotted in Fig. 1(b) as “•”.

Rounding and Local Search. The step we just performed is called *projected gradient descent*, and it is well known that iteratively performing this step guarantees convergence to the SDP optimal solution [11]. However, the drawback is that it typically requires many iterations for the convergence to happen. Now we will show that, by leveraging local search back in POP (3), we can quickly arrive at the rank-one optimal point. The intuition is, from Fig. 1(b), we can observe that $\bar{z}^{(1)}$ has moved “closer” to z^* , and hence maybe $\bar{z}^{(1)}$ already contains useful local information about z^* . Particularly, let us perform a spectral decomposition of the top-left 2×2 block of $X(\bar{z}^{(1)})$:

$$\begin{bmatrix} 1 & -0.35 \\ -0.35 & 3.99 \end{bmatrix} = \lambda_1 v_1 v_1^\top + \lambda_2 v_2 v_2^\top, \quad \text{with} \\ v_1 = 0.12 \begin{bmatrix} 1 \\ -8.59 \end{bmatrix}, \quad v_2 = 0.99 \begin{bmatrix} 1 \\ 0.12 \end{bmatrix}, \quad \lambda_{1,2} = (4.04, 0.96),$$

where we have written $v_{1,2}$ by normalizing their leading entries as 1, and ordered them such that $\lambda_1 \geq \lambda_2$. We then generate two *hypotheses* $\bar{x}_{1,2}^{(1)} = (-8.59, 0.12)$ (“•” in Fig. 1(a)) and use NLP to perform local search starting from the hypotheses $\bar{x}_{1,2}^{(1)}$, respectively. Surprisingly, although $\bar{x}_1^{(1)}$ still converges to $x^{(1)} = -2$, $\bar{x}_2^{(1)}$ converges to $x^* = 2$. We call this step *rounding and local search*.

Lift and Certify. Now that NLP has visited both $x^{(1)}$ and x^* , with x^* attaining a lower cost, we are certain that x^* is at least a better local minimum. To *certify* the global optimality of x^* , we need the SDP relaxation again. We lift x^* to $z^* = [2, 4, 8]^\top$, and try to descend from z^* by performing another step of projected gradient descent. However, because z^* is optimal, this step ends up back to z^* again, *i.e.*, no more descent is possible [11]. Therefore, we conclude that z^* is optimal for SDP (4). Because $X(z^*)$ is rank one, and $p(x^*) = f(z^*)$, we can say that x^* is indeed globally optimal for the nonconvex POP (3) (*i.e.*, a numerical *certificate* of global optimality is obtained for x^*).

Despite the simplicity of this example, it clearly demonstrates the stark contrast between our algorithm STRIDE and existing approaches. While all existing methods solve the SDP relaxation *independently* from the original POP (*i.e.*, they relax the POP into an SDP and solve the SDP without *computationally* revisiting the POP), our method alternates between *local search* in POP and *global descent* in SDP, and the connection between the POP and the SDP is established by *lifting* and *rounding*. The advantages for doing so are threefold. (i) *Scalability*: Although the SDP may have millions of variables (n and m can grow rapidly), the original POP has only hundreds or thousands of variables (recall that d is moderate) and an NLP solver can perform local search quickly (typically in a negligible amount of time compared to solving the SDP); (ii) *Degeneracy*: When the optimal SDP solution is rank one and degenerate, convergence of existing scalable SDP solvers (such as first-order methods) is typically very slow or even unachievable

(cf. results in Section 5). However, the original POP is much less sensitive to the degeneracy. Once the SDP iterate gets close to being optimal, NLP can quickly arrive at the global optimal solution. (iii) *Warm-starting*: Unlike IPMs, for which designing good initialization is relatively challenging, our algorithm easily benefits from warm-starting. As we will show in Section 5, in many practical engineering applications, there exist powerful heuristics that find the globally optimal POP solution with high probability of success. In this case, our algorithm only needs to perform the last step of lifting and certification.

However, several questions need to be answered in order for our algorithm to be general and practical for large-scale problems.

Question 1 Since the projection onto a spectrahedron is itself an SDP and typically cannot be done exactly when n, m are large, how can one design a *globally convergent* solver for the original SDP (P) that can tolerate *inexactness* of the projection subproblem?

Question 2 With rounding and local search in the loop, is it still possible for the algorithm to be globally convergent?

Question 3 How can one design a *scalable and efficient* numerical algorithm to handle millions of constraints (i.e., m) when computing the (inexact) projection onto the spectrahedron?

Question 4 While our univariate example shows that this framework works on the simple example above, will it work on more complicated problems arising in real world applications?

Section 2 answers Question 1, where we design an *inexact* projected gradient method (iPGM) for solving a generic SDP pair (P)-(D), prove its global convergence, and provide complexity analysis. Although our results are inspired by [33], several nontrivial extensions are made. In answering Question 2 (Section 3), we incorporate rounding and local search into iPGM and formally present STRIDE. In a nutshell, STRIDE follows the globally convergent trajectory driven by iPGM, but simultaneously probes long, but safeguarded, rank-one vertices of the spectrahedron generated from solutions of NLP, to seek rapid descent and convergence. Notably, we prove that, even with rounding and local search in the loop, STRIDE is guaranteed to converge to the optimal solution of (P)-(D). In Section 4, we focus on solving the critical subproblem of projecting a given symmetric matrix onto the feasible set of (P) and provide an efficient answer to Question 3. We propose a two-phase algorithm where phase one uses a *symmetric Gauss-Seidel based accelerated proximal gradient method* (sGS-APG) to generate a good initial point, and phase two applies a modified limited-memory BFGS (L-BFGS) method to compute an accurate solution. This two-phase algorithm is simple, robust, easy to implement, and can scale to very large SDP problems with millions of constraints. Additionally, for the modified L-BFGS algorithm, we establish a novel convergence result where the objective function does not need to be at least twice continuously differentiable. Finally, we answer Question 4 by providing extensive numerical results in Section 5. We apply STRIDE to solve (dense and sparse) second-order moment relaxations arising from a diverse set of POP problems and demonstrate its superior performance. We observe that STRIDE is the only solver that can consistently

solve rank-one semidefinite relaxations to high accuracy (*e.g.*, KKT residuals below $1\text{e-}9$) and it is up to 1-2 orders of magnitude faster than existing SDP solvers.

Notation. We use \mathcal{X} and \mathcal{Y} to denote finite dimensional Euclidean spaces. Let \mathbb{S}^n be the space of real symmetric $n \times n$ matrices, and \mathbb{S}_+^n (resp. \mathbb{S}_{++}^n) be the set of positive semidefinite (resp. definite) matrices. We also write $X \succeq 0$ (resp. $X \succ 0$) to indicate that X is positive semidefinite (resp. definite) when the dimension of X is clear. We use $\mathcal{S}^{d-1} := \{x \in \mathbb{R}^d \mid \|x\| = 1\}$ to denote the d -dimensional unit sphere. For $x \in \mathbb{R}^n$, $\|x\| = \sqrt{x^\top x}$ is the standard ℓ_2 norm. For $X \in \mathbb{R}^{m \times n}$, $\|X\| = \sqrt{\text{tr}(X^\top X)}$ denotes the Frobenius norm. $\|x\|_{\mathcal{H}} := \sqrt{\langle x, \mathcal{H}x \rangle}$ denotes the \mathcal{H} -weighted norm for a positive semidefinite mapping \mathcal{H} . We use \mathcal{I} to denote the identity map from $\mathbb{S}^n \rightarrow \mathbb{S}^n$. We use $\delta_{\mathcal{C}}(\cdot)$ to denote the indicator function of set \mathcal{C} . For a positive integer d , we use $\mathfrak{t}(n) := \frac{d(d+1)}{2}$ to denote the d -th triangle number, and we use $\bar{d}_\kappa := \binom{n+\kappa}{\kappa}$ for some positive integer κ , which is particularly helpful when describing the number of monomials in $x \in \mathbb{R}^d$ of degree up to κ . For $x \in \mathbb{R}^d$, we use $[x]_\kappa \in \mathbb{R}^{\bar{d}_\kappa}$ to denote the full set of standard monomials of degree up to κ (as already used in the introductory paragraphs).

2 An inexact Projected Gradient Method for SDPs

In this section, we describe an inexact Projected Gradient Method (iPGM) for solving the SDP (P) and analyze its global convergence properties. Accelerating iPGM via rounding and lifting by nonlinear programming will be presented in Section 3.

Algorithm 1 presents the pseudocode for iPGM. Denoting the feasible set of the primal SDP (P) as

$$\mathcal{F}_P := \{X \in \mathbb{S}^n \mid \mathcal{A}(X) = b, X \in \mathbb{S}_+^n\},$$

the k -th iteration of iPGM first moves along the direction of the negative gradient (*i.e.*, $-C$) with step size $\sigma_k > 0$, and then performs an inexact projection of the trial point $X^{k-1} - \sigma_k C$ onto the primal feasible set \mathcal{F}_P (*cf.* (5)), with inexactness conditions given in (6). In Section 4, we will present efficient algorithms that can fulfill the inexactness conditions in (6).

The following theorem states the convergence properties of Algorithm 1 (iPGM).

Theorem 1 *Let $\{(X^k, y^k, S^k)\}$ be any sequence generated by Algorithm 1, and suppose that there exists a constant $M > 0$ such that $\|y^k\| \leq M$ for all $k \geq 0$. Then, for all $k \geq 1$, it holds that*

$$\begin{cases} -\frac{\|y^*\| k \varepsilon_k}{k} \leq \langle C, X^k - X^* \rangle \leq \frac{1}{k} \left(\frac{1}{2\sigma_0} \|X^0 - X^*\|^2 + 2M \sum_{i=1}^k i \varepsilon_i \right), \\ \|\mathcal{A}(X^k) - b\| \leq \varepsilon_k \leq O\left(\frac{1}{k}\right), \\ \|\mathcal{A}^* y^k + S^k - C\| \leq O\left(\frac{1}{\sqrt{k}\sigma_k}\right), \end{cases}$$

where X^* is an optimal solution of the primal problem (P) and (y^*, S^*) is an optimal solution of the dual problem (D). In particular, if one chooses $\sigma_k \geq O(\sqrt{k}) > 0$

Algorithm 1 An inexact projected gradient method (iPGM) for SDP (P).

Input: Initial points $(X^0, y^0, S^0) \in \mathbb{S}_+^n \times \mathbb{R}^m \times \mathbb{S}_+^n$, a nondecreasing positive sequence $\{\sigma_k\}$ and a nonnegative sequence $\{\varepsilon_k\}$ such that $\{k\varepsilon_k\}$ is summable.

Repeat For $k \geq 1$:

 Compute

$$X^k \approx \Pi_{\mathcal{F}_P} \left(X^{k-1} - \sigma_k C \right), \quad X^k \in \mathbb{S}_+^n \quad (5)$$

with associated dual pair $(y^k, S^k) \in \mathbb{R}^m \times \mathbb{S}_+^n$ such that the following conditions hold:

$$\begin{aligned} \left\| \mathcal{A}(X^k) - b \right\| &\leq \varepsilon_k, \\ \mathcal{A}^* y^k + S^k - C - \frac{1}{\sigma_k} (X^k - X^{k-1}) &= 0, \\ \langle X^k, S^k \rangle &= 0. \end{aligned} \quad (6)$$

Until: Termination conditions are met (cf. Section 5).

Output: (X^k, y^k, S^k) .

for all $k \geq 1$, then we have

$$\max \left\{ \left| \langle C, X^k - X^* \rangle \right|, \left\| \mathcal{A}(X^k) - b \right\|, \left\| \mathcal{A}^* y^k + S^k - C \right\| \right\} \leq O \left(\frac{1}{k} \right).$$

The proof of Theorem 1 is given in Appendix A. Although an accelerated version of Algorithm 1 and its convergence analysis have been studied in [33], Theorem 1 and its proof are new, to the best of our knowledge. By the presented results, if one chooses ε_k to be sufficiently small and σ_k to be sufficiently large, then the convergence of primal and dual infeasibilities can be as fast as $O(1/k)$. However, a small ε_k or a large σ_k will make the projection problem in (5) more difficult to solve. Hence, for better overall efficiency, one may choose ε_k and σ_k dynamically to balance the convergence speed of Algorithm 1 and the efficiency of the (inexact) projection onto \mathcal{F}_P .

3 STRIDE: Accelerating iPGM by Nonlinear Programming

Despite being globally convergent, it may take many iterations for Algorithm 1 to converge to a solution of high accuracy, especially when the SDP (P) has large scale and the optimal solution X^* is low-rank and degenerate. Therefore, in this section, we propose to accelerate Algorithm 1 by rounding and local search, just as what we have shown in the simple numerical example in Fig. 1 of Section 1. The intuition here is simple: when the SDP relaxation (P) is exact, there exist rank-one optimal solutions and they may be computed much more efficiently by performing local search in the low-dimensional (POP) with proper initialization.

With this intuition, we now develop the details of the acceleration scheme. At each iteration with current iterate $X^{k-1} \in \mathbb{S}_+^n$, we first follow (5) in Algorithm 1 and compute the projection

$$\bar{X}^k \approx \Pi_{\mathcal{F}_P} \left(X^{k-1} - \sigma_k C \right)$$

under the inexactness conditions described in (6), where $\sigma_k > 0$ is a given step size. Then by the analysis in Theorem 1, \bar{X}^k is guaranteed to make certain progress towards optimality. However, the point \bar{X}^k may not be a promising candidate for the next iteration since it may not attain rapid convergence. Motivated by the success of rounding and local search in the simple example in Fig. 1, we propose to compute a potentially better candidate based on \bar{X}^k via the following steps:

1. **(Rounding)**. Let $\bar{X}^k = \sum_{i=1}^n \lambda_i v_i v_i^\top$ be the spectral decomposition of \bar{X}^k with $\lambda_1 \geq \dots \geq \lambda_n$ in nonincreasing order. Compute $r \geq 1$ hypotheses for the (POP) from the leading r eigenvectors v_1, \dots, v_r

$$\bar{x}_i^k = \text{rounding}(v_i), \quad i = 1, \dots, r, \quad (7)$$

where the function **rounding** can be problem-dependent and we provide examples in the numerical experiments in Section 5. Generally, if the SDP (P) comes from a dense relaxation, and the feasible set of the original (POP), denoted as \mathcal{F}_{POP} , is simple to project, then we can design **rounding** to be

$$v_i \leftarrow \frac{v_i}{v_i[1]}, \quad \bar{x}_i^k \leftarrow \Pi_{\mathcal{F}_{\text{POP}}}(v_i[x]), \quad (8)$$

where one first normalizes v_i such that its leading entry $v_i[1]$ is equal to 1, and then projects its entries corresponding to order-one monomials, *i.e.*, $v_i[x]$, to the feasible set of (POP). Note that the rationale for designing this rounding method is that practical POP applications often involve simple constraints such as binary [29], unit sphere [74], and orthogonality [76, 15], whose projection maps are simple (*e.g.*, if x is binary, then $\Pi_{\mathcal{F}_{\text{POP}}}(v_i[x]) = \text{sgn}(v_i[x])$ just takes the sign of each entry of $v_i[x]$). If \mathcal{F}_{POP} is not easy to project, then we omit the projection, in which case the local search will start from an infeasible initialization.

2. **(Local Search)**. Apply a local search method for the POP (as an NLP) with the initial point chosen as \bar{x}_i^k for each hypothesis $i = 1, \dots, r$. Denote the solution of each local search as \hat{x}_i^k , with associated objective value $p(\hat{x}_i^k)$, choose the best local solution with minimum objective value. Formally, we have

$$\hat{x}_i^k = \text{nlp}(\bar{x}_i^k), \quad i = 1, \dots, r, \quad \hat{x}^k = \arg \min_{\hat{x}_i^k, i=1, \dots, r} p(\hat{x}_i^k). \quad (9)$$

3. **(Lifting)**. Perform a rank-one lifting of the best local solution according to the SDP relaxation scheme

$$\hat{X}^k = v(\hat{x}^k) v(\hat{x}^k)^\top, \quad (10)$$

where $v : \mathbb{R}^d \rightarrow \mathbb{R}^n$ is a dense or sparse monomial lifting (*e.g.*, $v(x) = [x]_\kappa$ is the full set of monomials up to degree κ in the case of dense Lasserre's hierarchy at order κ).

Now, we are given two candidates for the next iteration, namely \bar{X}^k (generated by computing the projection of $X^{k-1} - \sigma_k C$ onto the feasible set \mathcal{F}_P inexactly) and \hat{X}^k (obtained by rounding, local search and lifting described just now), the follow-up question is: which one should we choose to be the next iterate X^k such that the entire sequence $\{X^k\}$ is globally convergent for the SDP (P)?

The answer to this question is quite natural –we accept \hat{X}^k if and only if it attains a strictly lower cost than \bar{X}^k – and guarantees that the algorithm visits a sequence of rank-one vertices (local minima via NLP) with descending costs. With this insight, we now introduce our algorithm STRIDE in Algorithm 2. STRIDE is a combination of Algorithm 1 and the acceleration techniques in items 1-3, but with a judicious safeguarding policy (11) that ensures (i) the rank-one iterate \hat{X}^k is feasible (i.e., $\hat{X}^k \in \mathcal{F}_P$), and (ii) the rank-one iterate \hat{X}^k attains a strictly lower cost than \bar{X}^k and previously accepted rank-one iterates. Notice that requiring \hat{X}^k to attain a lower cost than previous rank-one iterates can prevent the algorithm from revisiting the same vertex.

Algorithm 2 An inexact projected gradient method accelerated via rounding and lifting by NLP for solving tight SDP relaxations of POPs (STRIDE).

Input: Initial points $(X^0, y^0, S^0) \in \mathbb{S}_+^n \times \mathbb{R}^m \times \mathbb{S}_+^n$, a nondecreasing positive sequences $\{\sigma_k\}$, a nonnegative sequence $\{\varepsilon_k\}$ such that $\{k\varepsilon_k\}$ is summable, a stopping criterion $\eta : \mathbb{S}_+^n \times \mathbb{R}^m \times \mathbb{S}_+^n \rightarrow \mathbb{R}_+$ with a tolerance $\text{To1} > 0$. Initialize $\mathcal{V} = \{X^0\}$ if $X^0 \in \mathcal{F}_P$ and $\mathcal{V} = \emptyset$ otherwise. Choose a positive integer $r \in [1, n]$, and a positive constant $\epsilon > 0$.

Iterate the following steps for $k = 1, \dots$:

Step 1 (Projection). Compute (\bar{X}^k, y^k, S^k) satisfying (5) and (6) (see Section 4).

Step 2 (Certification). If $\eta(\bar{X}^k, y^k, S^k) < \text{To1}$, **output** (\bar{X}^k, y^k, S^k) and **stop**.

Step 3 (Acceleration). Compute \hat{X}^k as in items 1-3:

$$\begin{aligned} \bar{X}^k &= \sum_{i=1}^n \lambda_i v_i v_i^\top, \quad \bar{x}_i^k = \text{rounding}(v_i), \quad i = 1, \dots, r, \\ \hat{x}_i^k &= \text{nlp}(\bar{x}_i^k), \quad i = 1, \dots, r, \quad \hat{x}^k = \arg \min_{\hat{x}_i^k, i=1, \dots, r} p(\hat{x}_i^k), \\ \hat{X}^k &= v(\hat{x}^k) v(\hat{x}^k)^\top. \end{aligned}$$

Step 4 (Policy). Update X^k according to

$$X^k = \begin{cases} \hat{X}^k & \text{if } \langle C, \hat{X}^k \rangle < \min \left(\langle C, \bar{X}^k \rangle, \min_{X \in \mathcal{V}} \{ \langle C, X \rangle \mid X \in \mathcal{V} \} \right) - \epsilon, \hat{X}^k \in \mathcal{F}_P \\ \bar{X}^k & \text{otherwise} \end{cases}. \quad (11)$$

If $X^k = \hat{X}^k$, set $\mathcal{V} \leftarrow \mathcal{V} \cup \{\hat{X}^k\}$.

We now state the convergence result for STRIDE.

Theorem 2 Let $\{(\bar{X}^k, y^k, S^k)\}$ be any sequence generated by Algorithm 2, and suppose that there exists a constant $M > 0$ such that $\|y^k\| \leq M$ for all $k \geq 0$. If one chooses $\sigma_k \geq O(\sqrt{k})$ for all $k \geq 1$, then, $\{\langle C, \bar{X}^k \rangle\}$ converges to the optimal value of the SDP problem (P).

Proof Let us only consider the case with $X^0 \in \mathcal{F}_P$ such that \mathcal{V} is initialized as $\{X^0\}$ (the case with \mathcal{V} initialized as \emptyset can be argued in a similar manner). By the construction of the set \mathcal{V} , it contains a sequence of feasible points of (P), and the objective value along this sequence is *strictly* decreasing according to the acceptance policy (11). Hence, we have the following relation:

$$\inf \{ \langle C, X \rangle \mid X \in \mathcal{V} \} \leq \langle C, X^0 \rangle - (|\mathcal{V}| - 1)\epsilon,$$

If $|\mathcal{V}| = \infty$, then there exists an infinite sequence of feasible points to problem (P) whose objective function value converges to $-\infty$. In this case, problem (P) is unbounded, a contradiction (recall that we assumed strong duality throughout this paper in Section 1). Hence, $|\mathcal{V}| < \infty$. If $|\mathcal{V}| = 1$, then none of the \hat{X}^k generated by **Step 3** has been accepted, in which case Algorithm 2 reduces to Algorithm 1 iPGM. Therefore, by Theorem 1, Algorithm 2 converges as $O(1/k)$ if one chooses $\sigma_k \geq O(\sqrt{k})$ for all $k \geq 1$. If $|\mathcal{V}| > 1$, then some of the \hat{X}^k generated by **Step 3** have been accepted. Denote the last element of \mathcal{V} as $\hat{X}_{\mathcal{V}}$ with associated dual variables $(y_{\mathcal{V}}, S_{\mathcal{V}})$ (from **Step 1**), then Algorithm 2 essentially reduces to iPGM with a new initial point at $(\hat{X}_{\mathcal{V}}, y_{\mathcal{V}}, S_{\mathcal{V}})$. Again, by Theorem 1, Algorithm 2 converges as $O(1/k)$. This completes the proof. \square

We now make a few remarks about Algorithm 2 (STRIDE).

The first remark is on the local search algorithm. In general, one can use any nonlinear programming method to perform local search on the original (POP) starting from an initial guess. A good choice is to use a solver that is based on a primal-dual interior point method, for which the global and local convergence properties are well studied [54, Chapter 19]. Notice that even if the NLP solver can fail to converge to a feasible point of (POP), the safeguarding policy (11) ensures that all the rank-one points in \mathcal{V} are feasible for the SDP (P) (*i.e.*, a failed NLP solution will not be accepted). In many POPs arising from practical engineering applications, the constraint set typically defines a smooth manifold (see examples in Section 5). In such cases, we prefer to use unconstrained optimization algorithms on the manifold as the local search method, for example the Riemannian trust region method that admits favorable global and local convergence properties [2, Chapter 7]. A general interior point method for NLP is available through `fmincon` in Matlab (see [54, Section 19.9] for other available software), while the Riemannian trust region method is available through `Manopt` [13]. We emphasize here that the idea of using local search algorithms for accelerating iPGM is heuristic and only supported by numerical experiments in Section 5. In other words, while we guarantee global convergence of STRIDE for solving the SDP (P) (under mild technical assumptions), the amount of acceleration gained by nonlinear programming may be problem dependent and is mostly observed empirically. In the present paper, we are not able to establish conditions under which the local search algorithms can provide provably better rank-one candidates than the iterates generated by iPGM. We think this aspect deserves deeper future research. Nevertheless, the nice property of STRIDE is that, even if we reject all the candidates provided by local search, global convergence is still guaranteed by taking the safeguarded iPGM steps.

Next, we comment on the number of eigenvectors to round. In STRIDE, the hyperparameter r decides how many eigenvectors to round and how many hypotheses to generate at each iteration. Since the NLP solver performs local search very quickly, it is affordable to choose a large r . However, we empirically observed that choosing r between 2 and 5 is sufficient for finding the global optimal solution.

Lastly, we provide a possible extension to STRIDE. The careful reader may have noticed that Algorithm 2 uses the general Algorithm iPGM *without* acceleration. We emphasize here that the accelerated version studied in [33] can also be used as the backbone of STRIDE with the same global convergence property. The reason we only implement the unaccelerated version is because empirically, with a proper warmstart (described in Section 3.1), we observe that Algorithm 2 converges in 1

or 2 iterations, just like what we have shown in the simple univariate example in Section 1. Therefore, the acceleration scheme would not have major improvement on the efficiency of Algorithm 2.

3.1 Initialization

STRIDE requires an initial guess $(X^0, y^0, S^0) \in \mathbb{S}_+^n \times \mathbb{R}^m \times \mathbb{S}_+^n$ as described in Algorithm 2. Although one can choose an arbitrary initial point and the algorithm is guaranteed to converge, a good initialization can significantly promote fast convergence. We first mention that, for many POPs arising from engineering applications, practitioners often have designed powerful heuristics for solving the POPs based on their specific domain knowledge. By heuristics we mean algorithms that can find the *globally optimal* POP solutions with very high probability of success, but *cannot* provide a certificate of global optimality (or suboptimality). Therefore, when such heuristics exist, we could use them to generate x^0 for the (POP) and perform a rank-one lifting of x^0 to form X^0 , *i.e.*, $X^0 = v(x^0)v(x^0)^\top$, where v is the lifting monomials. In Section 5.3, we demonstrate the effectiveness of one such heuristic in speeding up STRIDE on a computer vision application.

Here we describe a general initialization scheme based on a convergent semi-proximal alternating direction method of multipliers (sPADMM) [64]. sPADMM can be used to generate the dual initialization (y^0, S^0) when POP heuristic exists, or to generate both X^0 and (y^0, S^0) when no POP heuristics exist. We begin by reformulating problem (D) as follows:

$$\min_{y \in \mathbb{R}^m, S \in \mathbb{S}^n} \left\{ \delta_{\mathbb{S}_+^n}^*(-S) - \langle b, y \rangle \mid \mathcal{A}^*y + S = C \right\}. \quad (12)$$

Given $\sigma > 0$, the augmented Lagrangian associated with (12) is given as

$$\mathcal{L}_\sigma(y, S; X) = \delta_{\mathbb{S}_+^n}^*(-S) - \langle b, y \rangle + \langle X, \mathcal{A}^*y + S - C \rangle + \frac{\sigma}{2} \|\mathcal{A}^*y + S - C\|^2,$$

for $(X, y, S) \in \mathbb{S}^n \times \mathbb{R}^m \times \mathbb{S}^n$. The sPADMM method for solving problem (12) is described in Algorithm 3, and its convergence is well studied in [64]. An implementation of Algorithm 3 is included in the software package SDPNAL+ [78], namely the `admmplus` subroutine. Let the output of sPADMM be (X, y, S) , we use $(\Pi_{\mathbb{S}_+^n}(X), y, S)$ to be the initial point for Algorithm 2 (recall that STRIDE requires $X^0 \in \mathbb{S}_+^n$, so we project X to be positive semidefinite).

We note that the steplength γ in Algorithm 3 can take values in the interval $(0, 2)$ instead of the usual interval of $(0, (1 + \sqrt{5})/2)$. In Step 1 and 3 of the algorithm, \hat{y}^{k+1} and y^{k+1} are computed by using the sparse Cholesky factorization of $\mathcal{A}\mathcal{A}^*$, which is computed once at the beginning of the algorithm. One can also compute \hat{y}^{k+1} and y^{k+1} inexactly by using a preconditioned conjugate gradient method without affecting the convergence of the overall algorithm as long as the residual norms of the inexact solutions are bounded by a summable error sequence; we refer the reader to [19] for the details. We further add that for the computation of S^{k+1} in Step 2, one can make use of the expected low-rank property of $\Pi_{\mathbb{S}_+^n}(X^k + \sigma(\mathcal{A}^*\hat{y}^{k+1} - C))$. In our implementation we use the subroutine `dsyevx` in LAPACK to compute only the positive eigen-pairs of $X^k + \sigma(\mathcal{A}^*\hat{y}^{k+1} - C)$.

Algorithm 3 A semi-proximal ADMM for solving (12).

Input: Initial points $X^0 = S^0 = 0 \in \mathbb{S}^n$ and $\gamma \in (0, 2)$.

Iterate the following steps for $k = 1, \dots$:

Step 1. Compute

$$\hat{y}^{k+1} = \arg \min_y \mathcal{L}_\sigma(y, S^k; X^k) = (\mathcal{A}\mathcal{A}^*)^{-1} \left(\frac{1}{\sigma} b - \mathcal{A} \left(\frac{1}{\sigma} X^k + S^k - C \right) \right).$$

Step 2. Compute

$$\begin{aligned} S^{k+1} &= \arg \min_S \mathcal{L}_\sigma(\hat{y}^{k+1}, S; X^k) \\ &= \frac{1}{\sigma} \left(\Pi_{\mathbb{S}_+^n} \left(X^k + \sigma(\mathcal{A}^* \hat{y}^{k+1} - C) \right) - (X^k + \sigma(\mathcal{A}^* \hat{y}^{k+1} - C)) \right). \end{aligned}$$

Step 3. Compute

$$y^{k+1} = \arg \min_y \mathcal{L}_\sigma(y, S^{k+1}; X^k) = (\mathcal{A}\mathcal{A}^*)^{-1} \left(\frac{1}{\sigma} b - \mathcal{A} \left(\frac{1}{\sigma} X^k + S^{k+1} - C \right) \right).$$

Step 4. Compute

$$X^{k+1} = X^k + \gamma \sigma(S^{k+1} + \mathcal{A}^* y^{k+1} - C).$$

Until termination conditions are met.

Output: $(X^{k+1}, y^{k+1}, S^{k+1})$.

4 Solving the Projection Subproblem

Recall that the feasible set of (P) is $\mathcal{F}_P = \{X \in \mathbb{S}^n \mid \mathcal{A}(X) = b, X \succeq 0\}$. In this section, we aim at computing the projection onto \mathcal{F}_P efficiently since it is required at each iteration in STRIDE (Algorithm 2). Note that since there are m linear equality constraints defining \mathcal{F}_P with m possibly as large as a few millions, the projection algorithm has to be scalable. To this end, we propose a two-phase algorithm. In phase one (Section 4.1), we use an sGS-based accelerated proximal gradient method to generate a reasonably good initial point for the purpose of warm starting. Then in phase two (Section 4.2), we apply a modified version of the classical limited-memory BFGS (L-BFGS) method to compute a highly accurate solution.

Formally, given a point $Z \in \mathbb{S}^n$, the projection problem is described as finding the closest point in \mathcal{F}_P with respect to Z

$$\min_{X \in \mathbb{S}^n} \left\{ \frac{1}{2} \|X - Z\|^2 \mid X \in \mathcal{F}_P \right\}. \quad (13)$$

Notice that the feasible set \mathcal{F}_P is a spectrahedron defined by the intersection of two convex sets, namely the hyperplane $\{X \in \mathbb{S}^n : \mathcal{A}(X) = b\}$ and the PSD cone \mathbb{S}_+^n . Therefore, a natural idea is to apply Dykstra's projection (see *e.g.*, [22]) for generating an approximate solution to (13) by alternating the projection onto the hyperplane and the projection onto the PSD cone, both of which are easy to compute. However, Dykstra's projection is known to have slow convergence rate and it may take too many iterations until a satisfactory approximate projection is found. In this paper, instead of solving (13) directly, we consider its dual problem for which we can handle more efficiently.

It is not difficult to write down the Lagrangian dual problem (ignoring the constant term $-\frac{1}{2}\|Z\|^2$ and converting “max” to “min”) of (13) as:

$$\min_{W, \xi} \left\{ \frac{1}{2} \|W + \mathcal{A}^* \xi + Z\|^2 - \langle b, \xi \rangle \mid \xi \in \mathbb{R}^m, W \in \mathbb{S}_+^n \right\}. \quad (14)$$

For the rest of this section, we assume that the KKT system for (13)-(14) given as

$$\mathcal{A}(X) = b, \quad \mathcal{A}^* \xi + W = X - Z, \quad X, W \succeq 0, \quad \langle X, W \rangle = 0, \quad (15)$$

admits at least one solution, which is satisfied if problem (13) (or (P)) satisfies the Slater’s condition, *i.e.*, there exists an $X^\dagger \in \mathbb{S}^n$ such that

$$\mathcal{A}(X^\dagger) - b = 0, \quad X^\dagger \succ 0.$$

Before presenting our two-phase algorithm, let us briefly discuss about stopping conditions for solving the pair of projection SDPs (13) and (14). To this end, let $(W, \xi) \in \mathbb{S}_+^n \times \mathbb{R}^m$ be the output of any algorithm that solves the dual problem (14). From the KKT conditions (15), we deem

$$X(W, \xi) := \mathcal{A}^* \xi + W + Z \in \mathbb{S}^n \quad (16)$$

as an approximate solution for the primal problem (13). Given a tolerance parameter tol , we accept $X(W, \xi)$ as an approximate projection point if the relative KKT residue is below tol :

$$\eta_{\text{proj}}(W, \xi) := \max \left\{ \frac{\|\mathcal{A}(X(W, \xi)) - b\|}{1 + \|b\|}, \frac{\|X(W, \xi) - \Pi_{\mathbb{S}_+^n}(\mathcal{A}^* \xi + Z)\|}{1 + \|X(W, \xi)\| + \|W\|} \right\} \leq \text{tol}. \quad (17)$$

The first-order optimality conditions for problem (14), *i.e.*,

$$0 \in \begin{pmatrix} W + \mathcal{A}^* \xi + Z + \partial \delta_{\mathbb{S}_+^n}(W) \\ \mathcal{A}(W + \mathcal{A}^* \xi + Z) - b \end{pmatrix} = \begin{pmatrix} X(W, \xi) + \partial \delta_{\mathbb{S}_+^n}(W) \\ \mathcal{A}(X(W, \xi)) - b \end{pmatrix},$$

ensures that $\eta_{\text{proj}}(W, \xi)$ is small if (W, ξ) is sufficiently accurate, since $0 \in X(W, \xi) + \partial \delta_{\mathbb{S}_+^n}(W)$ implies that $X(W, \xi) - \Pi_{\mathbb{S}_+^n}(X(W, \xi) - W) = X(W, \xi) - \Pi_{\mathbb{S}_+^n}(\mathcal{A}^* \xi + Z) = 0$.

Now let us present our two-phase algorithm for solving the dual (14) for the remaining part of this section.

4.1 Phase One: An sGS-based Accelerated Proximal Gradient Method

Problem (14) is in the form of a convex composite minimization problem

$$\min_{W \in \mathbb{S}^n, \xi \in \mathbb{R}^m} \{F(W, \xi) := f(W, \xi) + g(W, \xi)\}, \quad (18)$$

where $g(W, \xi) := \delta_{\mathbb{S}_+^n}(W)$ denotes the indicator function for \mathbb{S}_+^n which is convex but non-smooth, and $f(W, \xi)$ is the following convex quadratic function

$$f(W, \xi) := \frac{1}{2} \left\langle \begin{pmatrix} W \\ \xi \end{pmatrix}, \mathcal{Q} \begin{pmatrix} W \\ \xi \end{pmatrix} \right\rangle + \left\langle \begin{pmatrix} Z \\ \mathcal{A}(Z) - b \end{pmatrix}, \begin{pmatrix} W \\ \xi \end{pmatrix} \right\rangle$$

with \mathcal{Q} written as

$$\mathcal{Q} := \begin{pmatrix} \mathcal{I} & \mathcal{A}^* \\ \mathcal{A} & \mathcal{A}\mathcal{A}^* \end{pmatrix}.$$

It is desirable to apply a proximal-type method for problem (18) whose objective is the sum of a smooth and a non-smooth functions. The key idea in a proximal-type method is to approximate the smooth part $f(\cdot)$ by a wisely chosen convex quadratic function, namely $q_k(\cdot)$, such that (i) $q_k + g$ is a good approximation to $f + g$ near the current iterate, and (ii) minimizing $q_k + g$ can be solved efficiently. In particular, one needs to choose an appropriate positive definite mapping \mathcal{H}_k ($k \geq 1$) for approximating the function $f(\cdot)$ by $q_k(\cdot)$ that is defined as follows:

$$q_k(W, \xi) := f(\widetilde{W}^k, \tilde{\xi}^k) + \left\langle \nabla f(\widetilde{W}^k, \tilde{\xi}^k), \begin{pmatrix} W - \widetilde{W}^k \\ \xi - \tilde{\xi}^k \end{pmatrix} \right\rangle + \frac{1}{2} \left\| \begin{pmatrix} W - \widetilde{W}^k \\ \xi - \tilde{\xi}^k \end{pmatrix} \right\|_{\mathcal{H}_k}^2 \quad (19)$$

for a given $(\widetilde{W}^k, \tilde{\xi}^k) \in \mathbb{S}_+^n \times \mathbb{R}^m$. There are many possible choices for \mathcal{H}_k . For example, \mathcal{H}_k can be chosen as \mathcal{Q} since the function $f(\cdot)$ is itself a quadratic function. However, since the variables W and ξ are coupled, minimizing the approximate function $q_k(W, \xi) + g(W, \xi)$, for $\mathcal{H}_k = \mathcal{Q}$, is as difficult as solving the original projection problem (18). Another possibility is to choose $\mathcal{H}_k = L_k \text{diag}(\mathcal{I}, \mathcal{I}_m) - \mathcal{Q}$ where L_k is a given positive scalar. However, the value of L_k could be difficult to choose. Indeed, choosing L_k too small or too large can both harm the efficiency and convergence. Even though one could apply certain back-tracking techniques (see *e.g.*, [10, Section 4]) for estimating L_k at each iteration, the computational cost may be expensive. To decouple the variables W and ξ and to achieve fast convergence, we will choose the operator \mathcal{H}_k by applying the symmetric block Gauss-Seidel (sGS) decomposition technique in [40].

To present the sGS decomposition method, we first need to introduce some useful notation. Let \mathcal{Q} have the following decomposition

$$\mathcal{Q} = \mathcal{U} + \mathcal{D} + \mathcal{U}^*, \quad (20)$$

where

$$\mathcal{U} = \begin{pmatrix} 0 & \mathcal{A}^* \\ 0 & 0 \end{pmatrix}, \quad \mathcal{D} = \begin{pmatrix} \mathcal{I} & 0 \\ 0 & \mathcal{A}\mathcal{A}^* \end{pmatrix},$$

and $\mathcal{D} \succ 0$ because \mathcal{A} is onto. Define the linear operator $\mathcal{T}_{\mathcal{Q}} := \mathcal{U}\mathcal{D}^{-1}\mathcal{U}^*$. Clearly, $\mathcal{T}_{\mathcal{Q}}$ is positive semidefinite. In sGS decomposition, we choose

$$\mathcal{H}_k := \mathcal{Q} + \mathcal{T}_{\mathcal{Q}}, \quad \forall k. \quad (21)$$

Let us first verify that \mathcal{H}_k is positive definite. To show this, we write

$$\mathcal{Q} + \mathcal{T}_{\mathcal{Q}} = (\mathcal{D} + \mathcal{U} + \mathcal{U}^*) + \mathcal{U}\mathcal{D}^{-1}\mathcal{U}^* = (\mathcal{D} + \mathcal{U})\mathcal{D}^{-1}(\mathcal{D} + \mathcal{U}^*), \quad (22)$$

and easily see that $\mathcal{H}_k \succ 0$ due to $\mathcal{D} \succ 0$ and $\mathcal{D} + \mathcal{U}$ being nonsingular. Inserting this choice of $\mathcal{H}_k = \mathcal{Q} + \mathcal{T}_{\mathcal{Q}}$ into the expression q_k in (19), we have

$$q_k(W, \xi) = f(W, \xi) + \frac{1}{2} \left\langle \begin{pmatrix} W - \widetilde{W}^k \\ \xi - \tilde{\xi}^k \end{pmatrix}, \mathcal{T}_{\mathcal{Q}} \begin{pmatrix} W - \widetilde{W}^k \\ \xi - \tilde{\xi}^k \end{pmatrix} \right\rangle, \quad (23)$$

and the subproblem at the k -th iteration to be solved is

$$\min_{W \in \mathbb{S}^n, \xi \in \mathbb{R}^m} \{q_k(W, \xi) + g(W, \xi)\}. \quad (24)$$

The nice property of choosing \mathcal{H}_k as in (21) and arriving at the subproblem (24) is that an optimal solution for (24) can be computed efficiently by solving three simpler problems, as stated in the following theorem.

Theorem 3 *Assume $\mathcal{A}\mathcal{A}^*$ is nonsingular, then the optimal solution $(W^+, \xi^+) \in \mathbb{S}_+^n \times \mathbb{R}^m$ for (24) can be computed exactly via the following steps:*

$$\begin{cases} \xi^\dagger := \arg \min_{\xi \in \mathbb{R}^m} \{f(\widetilde{W}^k, \xi)\}, \\ W^+ := \arg \min_{W \in \mathbb{S}_+^n} \{\delta_{\mathbb{S}_+^n}(W) + f(W, \xi^\dagger)\}, \\ y^+ := \arg \min_{\xi \in \mathbb{R}^m} \{f(W^+, \xi)\}. \end{cases} \quad (25)$$

The proof of Theorem 3 is given in [40, Theorem 1] and is omitted here. The reason that procedure (25) is simple is because (i) with W fixed as \widetilde{W}^k or W^+ (first and third line in (25)), optimizing ξ boils down to solving a linear system and can be solved exactly due to the sparsity of \mathcal{A} ; (ii) with ξ fixed as ξ^\dagger (second line in (25)), optimizing W resorts to performing a projection onto \mathbb{S}_+^n and can be done in closed form via computing eigenvalue decompositions.

With these preparations, we now formally present the sGS-based accelerated proximal gradient method (sGS-APG) for solving (14) in Algorithm 4, where we also give the closed-form solution of the procedure (25) in (26). Note that in Algorithm 4, the computation of ξ^k in **Step 1** is only needed for checking terminations.

Algorithm 4 An sGS-based accelerated proximal gradient method for (14).

Input: Initial points $\widetilde{W}^1 = W^0 \in \mathbb{S}_+^n$, termination tolerance $\text{tol} > 0$, and $t_1 = 1$.

Iterate the following steps for $k = 1, \dots$:

Step 1 (sGS update). Compute

$$\begin{aligned} \tilde{\xi}^k &= (\mathcal{A}\mathcal{A}^*)^{-1} (b - \mathcal{A}(Z) - \mathcal{A}(\widetilde{W}^k)), \\ W^k &= \Pi_{\mathbb{S}_+^n} (-\mathcal{A}^* \tilde{\xi}^k - Z) = \Pi_{\mathbb{S}_+^n} (\mathcal{A}^* \tilde{\xi}^k + Z) - (\mathcal{A}^* \tilde{\xi}^k + Z), \\ \xi^k &= (\mathcal{A}\mathcal{A}^*)^{-1} (b - \mathcal{A}(Z) - \mathcal{A}(W^k)). \end{aligned} \quad (26)$$

Step 2. Compute $t_{k+1} = \frac{1 + \sqrt{1 + 4t_k^2}}{2}$, and $\widetilde{W}^{k+1} = W^k + \frac{t_k - 1}{t_{k+1}} (W^k - W^{k-1})$.

Until: $\eta_{\text{proj}}(W^k, \xi^k) \leq \text{tol}$.

Output: (W^k, ξ^k) .

The convergence of sGS-APG is stated as follows.

Theorem 4 *Suppose that $(W^*, \xi^*) \in \mathbb{S}_+^n \times \mathbb{R}^m$ is an optimal solution of problem (14) and that $\mathcal{A}\mathcal{A}^*$ is nonsingular. Let $\{(W^k, \xi^k)\}$ be the sequence generated by Algorithm 4. Then there exists a constant $c \geq 0$ such that*

$$0 \leq f(W^k, \xi^k) - f(W^*, \xi^*) \leq \frac{c}{(k+1)^2}.$$

The proof of Theorem 4 can be taken directly from [40, Proposition 2]. Note that the proposed method in this subsection is in fact a direct application of the one proposed in [40, Section 4] which is designed for a more general convex composite quadratic programming model with multiple blocks. Moreover, [40, Section 4] also investigates the inexact computation for the corresponding subproblem. Thus, when $\mathcal{A}\mathcal{A}^*$ cannot be factorized efficiently, one may apply an iterative solver for solving the linear systems in (26) of Algorithm 4. We omit the details here since for the applications studied in this paper, factorizing $\mathcal{A}\mathcal{A}^*$ can always be done efficiently due to the sparsity of \mathcal{A} . We further remark that for the computation of W^k in (26), we can make use of the expected low-rank property of the projection $\Pi_{\mathbb{S}_+^n}(\mathcal{A}^*\xi^k + Z)$ corresponding to the primal variable. Again, we use the subroutine `dsyevx` in LAPACK to compute only the positive eigen-pairs of $\mathcal{A}^*\xi^k + Z$.

4.2 Phase Two: A Modified Limited-memory BFGS Method

Even though the sGS-APG method, as presented in Algorithm 4, converges as $O(1/k^2)$ in terms of objective value, the constant c depends on the distance between the initial point and the optimal solution, which can be quite large if the initial point is not well chosen. Therefore, it may require many iterations to reach a satisfactory solution when the initial point is far way from the optimal solution set. Moreover, the sGS-APG method may not have a fast local linear convergent property since the problem (14) is not strongly convex. To speed up the computation of the dual projection problem (14), we now propose a modified version of the classical limited-memory BFGS (L-BFGS) method that is warm-started by the solution from Algorithm 4.

To proceed, we observe that, fixing the unconstrained ξ , the dual projection problem (14) becomes finding the closest $W \in \mathbb{S}_+^n$ to the matrix $-(\mathcal{A}^*\xi + Z)$ and admits a closed-form expression

$$W = \Pi_{\mathbb{S}_+^n}(-\mathcal{A}^*\xi - Z). \quad (27)$$

As a result, problem (14) can be further simplified, after inserting (27), as

$$\min_{\xi \in \mathbb{R}^m} \phi(\xi) := \frac{1}{2} \left\| \Pi_{\mathbb{S}_+^n}(\mathcal{A}^*\xi + Z) \right\|^2 - \langle b, \xi \rangle, \quad (28)$$

with the gradient of $\phi(\xi)$ given as

$$\nabla \phi(\xi) = \mathcal{A} \Pi_{\mathbb{S}_+^n}(\mathcal{A}^*\xi + Z) - b. \quad (29)$$

We have used the following equality in getting (28)

$$\Pi_{\mathbb{S}_+^n}(-\mathcal{A}^*\xi - Z) + \mathcal{A}^*\xi + Z = \Pi_{\mathbb{S}_+^n}(\mathcal{A}^*\xi + Z)$$

due to the Moreau identity [45, Theorem 2.2]. Thus, if ξ^* is an optimal solution for problem (28), then we can recover W^* from (27). Formulating the dual projection problem as (28) has appeared multiple times; see, for instances [82, 44]. We note that the function $\phi(\cdot)$ is smooth but not twice continuously differentiable, and it is convex but not strongly convex.

Now that (28) is an unconstrained convex minimization problem in $\xi \in \mathbb{R}^m$ with gradient given in (29), many efficient algorithms are available for solving the problem, such as (accelerated) gradient descent methods [51], nonlinear conjugate gradient methods [23], quasi-Newton methods [54] and the semi-smooth Newton method [82]. For this paper, we decide to propose a modified limited-memory BFGS (L-BFGS) method because L-BFGS is easy to implement, can handle very large unconstrained optimization problems, and is typically the “the algorithm of choice” for large-scale problems. Empirically, we observed that our proposed L-BFGS warm-started by sGS-APG is efficient and robust for various class of applications (shown in Section 5). This observation also supports the discussions made in [54, Chapter 7]. To the best of our knowledge, this is the first work that demonstrates the effectiveness of L-BFGS, or in general quasi-Newton methods, on solving large and challenging SDPs.⁴

The template for the modified L-BFGS method is presented in Algorithm 5.

Algorithm 5 A modified Limited-memory BFGS method for (28).

Input: Starting point ξ^0 , integer $\text{mem} > 0$, a positive constant K , termination tolerance $\text{tol} > 0$, $\mu \in (0, 1/2)$, $\mu \leq \mu' < 1$, $\rho \in (0, 1)$ and $\tau_1, \tau_2 \in (0, \infty)$.

Iterate the following steps for $k = 1, \dots$:

Step 1 (Search direction). Choose $Q_k^0 \succ 0$, $\beta_k := \tau_1 \|\nabla\phi(\xi^k)\|^{\tau_2}$ and compute $d^k = -\beta_k \nabla\phi(\xi^k) - g^k$ where $g^k := Q_k \nabla\phi(\xi^k)$ with $Q_k \succeq 0$, is obtained via the two-loop recursion as in [54, Algorithm 7.4]. If $\|d^k\| \geq K$, then choose $d^k = -\beta_k \nabla\phi(\xi^k)$ (i.e., set $Q_k = 0$).

Step 2 (Line search). Set $\alpha_k = \rho^{m_k}$, where m_k is the smallest nonnegative integer m such that

$$\phi(\xi^k + \rho^k d^k) \leq \phi(\xi^k) + \mu \rho^m \langle \nabla\phi(\xi^k), d^k \rangle.$$

Compute $\xi^{k+1} = \xi^k + \alpha_k d^k$.

Step 3 (Update memory). Compute and save $u^k = \xi^{k+1} - \xi^k$ and $w^k = \nabla\phi(\xi^{k+1}) - \nabla\phi(\xi^k)$. If $k > \text{mem}$, discard the vector $\{u^{k-\text{mem}}, w^{k-\text{mem}}\}$ from storage.

Until: $\eta_{\text{proj}}(W^{k+1}, \xi^{k+1}) \leq \text{tol}$ with $W^{k+1} = \Pi_{\mathbb{S}_+^n}(-\mathcal{A}^* \xi^{k+1} - Z)$.

Output: (W^{k+1}, ξ^{k+1}) .

In Algorithm 5, we always choose $Q_k^0 = I_m$ for all $k \geq 0$ which implies that $Q_k \succeq 0$ (see e.g., [54, eq. (7.19)]). Therefore, the matrix $Q_k + \beta_k I_m \succ 0$ when $\nabla\phi(\xi^k)$ is not zero, and d^k computed in **Step 1** in Algorithm 5 can be shown to be a descent direction (i.e., the algorithm is well-defined). We state the convergence property of Algorithm 5 in the following theorem, whose proof is presented in Appendix B.

Theorem 5 Suppose that $\mathcal{A}\mathcal{A}^*$ is nonsingular and the Slater condition holds for (P). Then Algorithm 5 is well defined. Let the sequence $\{\xi^k\}$ be generated by Al-

⁴ From our extensive numerical trials, we found that accelerated gradient descent, nonlinear conjugate gradient, and semismooth Newton with CG iterative solver fail to give satisfactory performance, especially when m is very large.

gorithm 5 such that $\nabla\phi(\xi^k) \neq 0, \forall k \geq 0$. Then, the sequence $\{\xi^k\}$ is bounded and any accumulation point ξ of this sequence is an optimal solution of (28).

4.3 Connection between the Projection SDP and the Original SDP

Recall from Algorithm 2 that, $Z = X^{k-1} - \sigma_k C$ at iteration k for $k \geq 1$. Also recall from (16) that $X(W, \xi) = \mathcal{A}^* \xi + Z + W$ with $W = \Pi_{\mathbb{S}_+^n}(-\mathcal{A}^* \xi - Z) \in \mathbb{S}_+^n$. Combining these equations, we have $X(W, \xi) = \Pi_{\mathbb{S}_+^n}(\mathcal{A}^* \xi + Z) \in \mathbb{S}_+^n$ and hence, $\langle X(W, \xi), W \rangle = 0$. Furthermore, $X(W, \xi) = \mathcal{A}^* \xi + W + Z$ implies

$$\frac{1}{\sigma_k} \left(X(W, \xi) - X^{k-1} \right) = \mathcal{A}^* \left(\frac{1}{\sigma_k} \xi \right) + \frac{1}{\sigma_k} W - C, \quad (30)$$

which means that when $X(W, \xi)$ is close to X^{k-1} (which is the case when X^{k-1} is approximately optimal), then $\left(\frac{1}{\sigma_k} W, \frac{1}{\sigma_k} \xi \right)$ is an approximate solution for the original dual problem (D). Therefore, from the solution ξ of L-BFGS, we output

$$W = \Pi_{\mathbb{S}_+^n}(-\mathcal{A}^* \xi - Z) \quad (31)$$

$$X(W, \xi) = \mathcal{A}^* \xi + Z + W = \Pi_{\mathbb{S}_+^n}(\mathcal{A}^* \xi + Z) \quad (32)$$

$$(\bar{X}^k, y^k, S^k) = \left(X(W, \xi), \frac{1}{\sigma_k} \xi, \frac{1}{\sigma_k} W \right) \quad (33)$$

for **Step 1** of the k -th iteration of Algorithm 2. In particular, if

$$\|\nabla\phi(\xi)\| = \|\mathcal{A}(\bar{X}^k) - b\| \leq \varepsilon_k$$

for given $\varepsilon_k \geq 0$, then (\bar{X}^k, y^k, S^k) satisfies the inexact conditions (6).

5 Applications and Numerical Experiments

In this section, we conduct numerical experiments by using STRIDE to solve SDP relaxations for several important classes of POPs. In Section 5.1 and 5.2, we solve *dense* second-order moment relaxation of random *binary quadratic programming* problems, and random *quartic optimization problems over the unit square*, both with increasing number of variables d . We demonstrate, for the first time, that the relaxation is always tight for the instances we have generated up to $d = 60$. In Section 5.3 and 5.4, we solve *sparse* second-order moment relaxations coming from two engineering applications, namely an *outlier-robust Wahba problem* that underpins many computer vision applications, and a *nearest structured rank deficient matrix problem* that finds extensive applications in control, statistics, computer algebra, among others. We demonstrate that, with a sparse relaxation scheme, we can globally solve POP problems with d up to 1000, far beyond the reach of SDP solvers used in the corresponding engineering literature. When solving the outlier-robust Wahba problem in Section 5.3, we additionally show that (i) leveraging domain-specific POP heuristics for primal initialization can further speed up STRIDE by 2-3 times, and (ii) STRIDE can *certifiably optimally* solve two real applications of the outlier-robust Wahba problem, namely image stitching and scan

matching. Our experiments show that STRIDE is the *only* solver that can consistently solve rank-one tight semidefinite relaxations to high accuracy (*e.g.*, KKT residuals below $1e-9$), in the presence of millions of equality constraints.

Before presenting each application, let us describe the details about implementation and experimental setup.

Implementation. We implement the STRIDE Algorithm 2 in MATLAB R2020a, with core subroutines, such as projection onto the PSD cone, implemented in C for efficiency. Our implementation is available at

<https://github.com/MIT-SPARK/STRIDE>

and also supports SDP problems with multiple PSD blocks (hence, it can also handle relaxations of (POP) problems with inequality constraints).

Note that different from generic SDP solvers, we also implement the rounding, local search and lifting procedures required in Step 2 of Algorithm 2, for each POP. Since these procedures are problem dependent, we will describe them in the corresponding subsections.

Stopping Conditions. To measure the feasibility and optimality at a given approximate solution $(X, y, S) \in \mathbb{S}_+^n \times \mathbb{R}^m \times \mathbb{S}_+^n$, we define the following standard relative KKT residues:

$$\eta_p = \frac{\|\mathcal{A}(X) - b\|}{1 + \|b\|}, \quad \eta_d = \frac{\|\mathcal{A}^*y + S - C\|}{1 + \|C\|}, \quad \eta_g = \frac{|\langle C, X \rangle - \langle b, y \rangle|}{1 + |\langle C, X \rangle| + |\langle b, y \rangle|}. \quad (34)$$

For a given tolerance $\text{tol} > 0$, we terminate STRIDE when $\max\{\eta_p, \eta_d, \eta_g\} \leq \text{tol}$, and we choose $\text{tol} = 1e-8$ for all our experiments. Because our goal is to obtain a solution of the original (POP) with an optimality or suboptimality certificate, we also compute a relative suboptimality gap from the SDP solution (X, y, S) as

$$\eta_s = \frac{|p(\hat{x}) - (\langle b, y \rangle + M_b \lambda_{\min}(C - \mathcal{A}^*y))|}{1 + |p(\hat{x})| + |\langle b, y \rangle + M_b \lambda_{\min}(C - \mathcal{A}^*y)|}, \quad (35)$$

where $\hat{x} \in \mathcal{F}_{\text{POP}}$ is a feasible approximate solution to the (POP) that is rounded from the leading eigenvector of X ,⁵ λ_{\min} denotes the minimum eigenvalue, and $M_b \geq \text{tr}(X)$ is a bound on the trace of X when X is generated by a rank-one lifting. We easily have $p(\hat{x}) \geq p^* \geq \langle b, y \rangle + M_b \lambda_{\min}(C - \mathcal{A}^*y)$, and hence $\eta_s = 0$ (*i.e.*, the upper bound and the lower bound coincide) certifies that $p(\hat{x}) = p^*$ and \hat{x} is a global minimizer of the nonconvex (POP).

Baseline Solvers. We compare STRIDE with a diverse set of existing SDP solvers. We choose SDPT3 [66] and MOSEK [7] as representative interior point methods; CDCS [83] and SketchyCGAL [79] as representative first-order methods; and SDPNAL+ [78] as a representative method that combines first-order and second-order Newton-type methods. For SDPT3 and MOSEK, we use default parameters. For CDCS, we use the `sos` solver with maximum 20,000 iterations instead of the default homogeneous self-dual embedding solver because we found `sos` to typically perform better. For SketchyCGAL, we use the default parameters with sketching size 10. We set the maximum runtime to be 10,000 seconds and maximum number of iterations to be 20,000. For SDPNAL+, we use $1e-8$ as the tolerance, and we run it for maximum 20,000 iterations and 10,000 seconds. For very large problems

⁵ For a generic POP, even finding a feasible $x \in \mathcal{F}_{\text{POP}}$ is NP-hard. However, in all our numerical examples, rounding a feasible point from the approximate SDP solution is easy.

(*e.g.*, m above one million), we increase the maximum runtime of SketchyCGAL and SDPNAL+, which will be described in relevant subsections.

Hardware. All experiments are performed on a Linux PC with 12-core Intel i9-7920X CPU@2.90GHz and 128GB RAM.

5.1 Binary quadratic programming

Consider minimizing a *quadratic* polynomial over the d -dimensional binary cube

$$\min_{x_i \in \{+1, -1\}, i=1, \dots, d} \langle c, [x]_2 \rangle \quad (\text{BQP})$$

where $[x]_2 : \mathbb{R}^d \rightarrow \mathbb{R}^{\bar{d}_2}$ is the vector of monomials in x of degree up to 2, and $c \in \mathbb{R}^{\bar{d}_2}$ contains the coefficients of all monomials. Problem (BQP) is a classical NP-hard combinatorial problem with examples such as the maximum cut (MAXCUT) problem [29], the 0-1 knapsack problem [30], the number partitioning problem [50, 28], and the linear quadratic regulator control problem with binary inputs [71]. It is well known that the standard Shor's semidefinite relaxation for (BQP) is typically not tight, *e.g.*, in MAXCUT problems, and hence a globally optimal solution cannot be obtained with an optimality certificate (albeit a lower bound can be obtained).

We consider the *second-order* dense moment relaxation for (BQP), which creates a *positive semidefinite* moment matrix $X := [x]_2 [x]_2^\top$, using which the cost function in (BQP) can be written as $\langle C, X \rangle$. The moment matrix $X \in \mathbb{S}^{\bar{d}_2}$ contains all the monomials in x of degree up to 4, hence is a linear subspace of $\mathbb{S}^{\bar{d}_2}$ with dimension \bar{d}_4 . As a result, X must satisfy $m_{\text{mom}} = \mathbf{t}(\bar{d}_2) - \bar{d}_4$ linearly independent equality constraints, referred to as the *moment* constraints. In addition, X must satisfy *redundant* equality constraints, obtained by the fact that since $x_i^2 - 1 = 0$, it also holds $[x]_2(x_i^2 - 1) = 0$ for each $i = 1, \dots, d$. This leads to a total of $m_{\text{loc}} = d \times \bar{d}_2$ linear equality constraints. Last but not least, the top-left entry of X is equal to 1 due to the leading element of $[x]_2$ is 1 (the zero-order monomial).⁶ Therefore, the second-order relaxation for (BQP) leads to an SDP with size

$$n = \bar{d}_2, \quad m = m_{\text{mom}} + m_{\text{loc}} + 1 = \mathbf{t}(\bar{d}_2) - \bar{d}_4 + d \times \bar{d}_2 + 1, \quad (36)$$

which grows rapidly with d . Lasserre [36] showed that the second-order relaxation is empirically tight on a set of 50 randomly generated MAXCUT problems. However, due to the limitation of interior point methods back then, the experiments were performed on problems with small size $d = 10$.

In this paper, we aim to solve (BQP) instances with much larger d . We generate random instances of (BQP) by sampling the coefficients vector c from the standard zero-mean Gaussian distribution, *i.e.*, $c_i \sim \mathcal{N}(0, 1), i = 1, \dots, \bar{d}_2$. At $d \in \{10, 20, 30, 40, 50, 60\}$, we randomly generate three instances each and solve the second-order moment relaxation using SDPT3, MOSEK, SDPNAL+, CDCS, SketchyCGAL, and STRIDE. For STRIDE, we use the standard `fmincon` interface in Matlab with an interior point solver (supplied with analytical objective and constraint

⁶ The reader can refer to [43] or our code for details about generating SDP data \mathcal{A}, b, C .

gradients) as the **nlp** method. For **rounding** hypotheses from the moment matrix, we follow

$$X = \sum_{i=1}^n \lambda_i v_i v_i^\top, \quad v_i \leftarrow v_i / v_i[1], \quad \bar{x}_i = \text{sgn}(v_i[x]), \quad (37)$$

which first performs a spectral decomposition of X with $\lambda_1 \geq \dots \geq \lambda_n$ in non-increasing order, then normalizes the i -th eigenvector v_i so that its leading element is 1, and finally generates \bar{x}_i by taking the sign of the elements of v_i that correspond to the order-one monomials. We round $r = 5$ hypotheses using the first 5 eigenvectors ($i = 1, \dots, 5$). In order to compute η_s , we set $M_b = n$ because the diagonal entries of X are all equal to 1.

Table 1 gives the numerical results of different solvers. We make the following observations. (i) We first look at the performance of interior point methods (IPMs, SDPT3 and MOSEK). For small-scale problems ($d = 10$), IPMs can solve the SDPs efficiently to high accuracy (*e.g.*, around 1 second). For medium-scale problems ($d = 20$), although IPMs can still obtain solutions with high accuracy, the computational time starts to grow significantly (*e.g.*, 300-400 seconds). Moreover, for large-scale problems ($d \geq 30$), IPMs cannot be executed on ordinary workstations due to intensive memory consumption. The fundamental challenge of IPMs lies in solving large and dense linear systems (*i.e.*, the $m \times m$ Schur complement system) at each iteration. Although it is possible to use iterative solvers to solve the linear system [65], they are known to suffer from slow convergence as interior-point iterates approach optimality. (ii) First-order solvers (CDCS and SketchyCGAL) can solve problems to medium or low accuracy for $d \leq 20$, but their runtime can be worse than IPMs (although CDCS is faster than MOSEK for $d = 20$, its accuracy is orders of magnitude worse than MOSEK). However, first-order methods are indeed advantageous in terms of memory consumption and they can still be executed for problems with up to $d = 60$. Nevertheless, they are not able to compute POP solutions of certified global optimality within reasonable time, as shown by the nonzero η_s in Table 1 (*i.e.*, they are not able to show that the relaxation is indeed tight). This phenomenon further stresses the challenge for solving degenerate rank-one SDP relaxations. We also observe that, between CDCS and SketchyCGAL, CDCS seems to perform much better for such problems. This suggests that sketching may not be the best choice for degenerate SDPs with large m . (iii) SDPNAL+ has the best performance among existing solvers for (BQP). It can solve (BQP) instances to certified global optimality for up to $d = 40$, and it is over 10 times faster than MOSEK and CDCS when $d = 20$. However, when $d = 50$ and $d = 60$, SDPNAL+ cannot solve the SDPs to sufficient accuracy (within 10000 seconds), and hence the POP solution cannot be certified as globally optimal (*cf.* the nonzero η_s). (iv) Finally, we look at the performance of our solver STRIDE. We observe that STRIDE solved all the SDPs to high accuracy, certified the global optimality of the POP solutions, and demonstrated the tightness of the SDP relaxations (*cf.* the numerically zero η_s). For small and medium problems ($d = 10$ and 20), STRIDE attains accuracy that is comparable to MOSEK, while being about 30 times faster when $d = 20$. For large problems ($d \geq 30$), STRIDE attains accuracy that is superior to SDPNAL+, while being 5-10 times faster when $d = 50$. At $d = 60$ with m over a million, STRIDE becomes the only solver that can obtain solutions of high accuracy.

Table 1 Results on solving second-order relaxation of random (BQP) instances. “**” indicates solver out of memory.

Dimension	Run	Metric	SDPT3 [66]	MOSEK [7]	CDCS [83]	SketchyCGAL [79]	SDPNAL+ [78]	STRIDE
$d: 10$ $n: 66$ $m: 1871$	#1	η_p	1.4e-10	2.3e-9	9.7e-9	0.45	2.4e-9	8.2e-16
		η_d	4.9e-11	1.0e-9	8.4e-9	4.1e-6	7.8e-9	7.9e-16
		η_g	1.5e-8	2.8e-11	8.0e-11	0.0112	8.8e-8	5.0e-13
		η_s	1.2e-12	5.6e-14	1.6e-12	3.0e-9	1.6e-12	3.1e-16
		time	2.2	1.3	8.8	50.4	0.85	1.1
	#2	η_p	8.1e-11	7.4e-11	8.6e-9	0.022	6.7e-9	8.2e-16
		η_d	1.2e-11	3.0e-11	9.9e-9	1.8e-6	6.0e-10	8.8e-16
		η_g	6.0e-10	5.7e-13	1.1e-11	0.0037	6.2e-9	6.1e-13
		η_s	5.9e-14	2.5e-16	2.4e-13	1.7e-9	2.5e-14	0.0
		time	1.8	1.3	3.0	49.6	0.63	1.03
	#3	η_p	2.4e-10	5.5e-12	9.2e-9	1.37	2.89e-9	8.2e-16
		η_d	2.8e-11	5.8e-12	9.3e-9	1.3e-14	5.3e-10	6.3e-14
		η_g	3.5e-8	5.1e-15	3.0e-11	0.384	5.6e-9	5.1e-13
		η_s	2.3e-10	7.8e-17	4.4e-12	0.0033	2.1e-13	0.0
		time	1.9	1.4	10.3	49.8	0.89	0.96
$d: 20$ $n: 231$ $m: 20,791$	#1	η_p	3.0e-9	1.1e-8	1.1e-5	7.34	6.7e-9	1.6e-15
		η_d	3.9e-9	3.0e-9	2.0e-4	0.046	9.6e-11	5.1e-14
		η_g	1.8e-5	1.9e-11	9.6e-9	0.124	3.3e-9	8.1e-13
		η_s	9.9e-8	5.7e-13	2.3e-4	0.0076	1.8e-16	1.2e-16
		time	351.8	246.9	119.6	405.4	12.55	10.67
	#2	η_p	2.4e-10	1.8e-12	2.9e-5	4.64	3.7e-11	1.61e-15
		η_d	7.8e-9	4.3e-13	0.0013	0.0438	1.4e-8	8.2e-14
		η_g	2.1e-6	6.9e-15	3.9e-7	0.1529	2.6e-7	1.3e-12
		η_s	1.5e-12	5.0e-16	0.0013	0.0047	5.0e-16	9.9e-17
		time	341.5	250.0	119.5	403.2	19.6	10.4
	#3	η_p	1.6e-11	4.0e-11	3.9e-5	2.0963	2.2e-10	1.6e-15
		η_d	6.7e-10	1.0e-11	2.3e-5	0.0409	2.8e-11	3.3e-14
		η_g	1.3e-7	5.5e-14	3.5e-8	0.0391	4.2e-10	5.4e-13
		η_s	4.6e-13	3.4e-16	1.8e-5	2.9e-4	1.1e-16	1.1e-16
		time	342.8	223.5	117.4	404.3	13.0	10.0
$d: 30$ $n: 496$ $m: 91,761$	#1	η_p			5.0e-5	7.86	2.4e-10	2.4e-15
		η_d	**	**	0.0025	0.1199	5.5e-9	1.7e-15
		η_g			4.8e-7	0.1850	8.1e-7	2.6e-12
		η_s			0.0448	0.0208	5.4e-15	4.5e-16
		time			403.6	2028	61.9	40.1
	#2	η_p			1.0e-4	3.93	2.8e-9	2.4e-15
		η_d	**	**	0.0023	0.1311	7.0e-10	6.8e-14
		η_g			2.8e-7	0.0624	1.4e-8	1.6e-12
		η_s			0.0082	0.0049	0	2.3e-16
		time			398.0	2051	94.2	56.7
	#3	η_p			1.1e-4	5.62	5.0e-11	2.4e-15
		η_d	**	**	0.0032	0.0995	1.6e-8	1.0e-13
		η_g			1.0e-6	0.2739	1.4e-6	2.9e-12
		η_s			0.0115	0.0084	6.3e-15	6.4e-17
		time			405.5	2026	95.5	56.3
$d: 40$ $n: 861$ $m: 269,781$	#1	η_p			1.2e-4	8.79	3.4e-9	3.2e-15
		η_d	**	**	0.0030	0.1386	1.5e-9	1.5e-14
		η_g			7.8e-7	0.4433	5.2e-8	5.7e-13
		η_s			0.0602	0.0828	3.1e-16	2.3e-16
		time			1836	8962	1384	461.6
	#2	η_p			9.5e-5	8.60	3.66e-10	3.2e-15
		η_d	**	**	0.0027	0.1340	2.1e-9	5.8e-14
		η_g			6.8e-7	0.3349	8.1e-8	2.2e-12
		η_s			0.0395	0.1746	4.0e-16	4.0e-16
		time			1752	8781	1295	460.6
	#3	η_p			2.1e-4	7.23	9.1e-9	3.2e-15
		η_d	**	**	0.0043	0.1331	2.8e-9	1.7e-13
		η_g			7.8e-7	0.2462	1.0e-7	5.7e-12
		η_s			0.0444	0.0295	2.9e-16	7.2e-17
		time			1780	8845	414.6	385.8
$d: 50$ $n: 1,326$ $m: 629,851$	#1	η_p			5.9e-5	18.07	6.9e-12	4.0e-15
		η_d	**	**	0.0020	0.5285	0.0012	7.7e-14
		η_g			3.6e-7	0.9033	0.0133	4.0e-12
		η_s			0.0446	0.9990	0.0490	4.1e-16
		time			5232	7654	10000	1180
	#2	η_p			4.2e-5	18.06	9.2e-10	4.0e-15
		η_d	**	**	0.0016	0.5526	3.5e-7	3.0e-14
		η_g			2.3e-7	0.7747	0.0013	1.6e-12
		η_s			0.0532	0.9394	6.2e-8	5.7e-17
		time			5203	7656	5805	1440
	#3	η_p			6.8e-5	17.98	1.3e-7	4.0e-15
		η_d	**	**	0.0024	0.5607	5.3e-8	1.6e-15
		η_g			3.6e-7	0.7953	2.5e-6	8.2e-13
		η_s			0.0749	0.9023	7.4e-14	4.4e-16
		time			5312	7562	9884	1368
$d: 60$ $n: 1,891$ $m: 1,266,971$	#1	η_p			5.8e-5	25.91	8.3e-13	4.7e-15
		η_d	**	**	0.0026	2.17	4.7e-4	4.5e-13
		η_g			2.8e-7	0.9721	0.0209	3.0e-11
		η_s			0.0891	0.9756	0.0015	1.2e-15
		time			9731	8387	10000	4083
	#2	η_p			6.8e-5	25.87	1.4e-11	4.7e-15
		η_d	**	**	0.0028	2.16	3.5e-5	1.6e-13
		η_g			3.2e-7	0.8612	0.0019	1.1e-11
		η_s			0.0411	0.9998	9.1e-5	1.0e-15
		time			9732	8361	10000	3868
	#3	η_p			6.0e-5	25.87	0.5488	4.7e-15
		η_d	**	**	0.0026	2.23	2.3e-7	2.6e-13
		η_g			2.9e-7	0.8746	0.9969	1.7e-11
		η_s			0.0973	0.9993	0.7549	1.7e-16
		time			10567	8320	10000	3804

5.2 Quartic programming on the sphere

Consider minimizing a *quartic* polynomial over the d -dimension unit sphere

$$\min_{x \in S^{d-1}} \langle c, [x]_4 \rangle \quad (\text{Q4S})$$

where $[x]_4 : \mathbb{R}^d \rightarrow \mathbb{R}^{\bar{d}_4}$ is the vector of monomials in x of degree up to 4, and $c \in \mathbb{R}^{\bar{d}_4}$ is the vector of known coefficients. Problem (Q4S) is known to be NP-hard with important examples such as computing the largest stable set of a graph [24, Theorem 3.4], computing the $2 \rightarrow 4$ norm of a matrix [8], and the best separable state problem in quantum information theory [25]. See [26, 41] and references therein for a thorough discussion about problem (Q4S).

Here we consider the dense second-order (also the lowest order) moment relaxation of (Q4S) and numerically show that they are indeed tight and admit rank-one solutions. By following the same relaxation scheme as in Section 5.1 (*i.e.*, build the moment matrix $X = [x]_2[x]_2^\top$ and add equality constraints), we can count the size of the SDP relaxation to be

$$n = \bar{d}_2, \quad m = \mathfrak{t}(\bar{d}_2) - \bar{d}_4 + \bar{d}_2 + 1. \quad (38)$$

At each $d \in \{10, 20, 30, 40, 50, 60\}$, we generate three random instances of (Q4S) by drawing c from the standard normal distribution. We solve the resulting SDP using SDPT3, MOSEK, CDCS, SketchyCGAL, SDPNAL+, and STRIDE. For STRIDE, we exploit the manifold structure of the sphere constraint and adopt Manopt [13] with a trust region solver as the `nlp` method. One can also treat (Q4S) as a standard nonlinear programming and solve it with `fmincon`, but we found that Manopt is faster and more robust for this problem. To generate hypotheses for `nlp` from the moment matrix, we follow

$$X = \sum_{i=1}^n \lambda_i v_i v_i^\top, \quad v_i \leftarrow v_i / v_i[1], \quad \bar{x}_i = \frac{v_i[x]}{\|v_i[x]\|}, \quad (39)$$

which first performs a spectral decomposition of X with $\lambda_1 \geq \dots \geq \lambda_n$ in non-increasing order, then normalizes the i -th eigenvector v_i so that its leading element is 1, and finally generates \bar{x}_i by projecting the elements of v_i that correspond to the order-one monomials onto the unit sphere. We generate $r = 5$ hypotheses using the first 5 eigenvectors. A valid upper bound M_b on $\text{tr}(X)$ can be obtained as

$$\begin{aligned} \text{tr}(X) &= \text{tr}([x]_2[x]_2^\top) = [x]_2^\top[x]_2 \\ &= 1 + \sum_{i=1}^d x_i^2 + \sum_{1 \leq i < j \leq d} (x_i x_j)^2 \leq 2 + \left(\sum_{i=1}^d x_i^2 \right)^2 = 3 := M_b. \end{aligned}$$

Table 2 gives the numerical results of different solvers. We make the following observations. (i) Similar to Table 1 for the (BQP) problem, IPMs can solve small and medium problems ($d = 10$ and 20) to high accuracy, although the runtime grows significantly from $d = 10$ (about 1 second) to $d = 20$ (about 100 seconds). (ii) Both CDCS and SDPNAL+ are able to solve all SDPs to high accuracy and certify the tightness of the second-order relaxation (despite that CDCS only attained medium accuracy for #2 at $d = 40$). However, SDPNAL+ is significantly faster than

Table 2 Results on solving second-order relaxation of random (Q4S) instances. “**” indicates solver out of memory.

Dimension	Run	Metric	SDPT3 [66]	MOSEK [7]	CDCS [83]	SketchyCGAL [79]	SDPNAL+ [78]	STRIDE
$d: 10$ $n: 66$ $m: 1277$	#1	η_p	4.5e-9	3.3e-8	1.2e-11	0.1388	4.6e-12	1.5e-16
		η_d	5.1e-11	5.1e-9	9.7e-11	8.7e-4	1.9e-11	3.5e-15
		η_g	5.7e-9	3.4e-9	9.4e-13	0.1823	6.9e-12	6.8e-10
		η_s	3.6e-12	3.2e-11	1.4e-12	0.0116	2.1e-11	2.0e-11
		time	1.2	0.6	2.5	58.0	0.6	1.2
	#2	η_p	1.7e-9	1.3e-8	8.8e-11	0.0291	1.8e-11	1.6e-16
		η_d	1.4e-10	2.0e-9	2.4e-11	6.0e-4	1.7e-11	1.0e-11
		η_g	6.5e-9	2.1e-10	8.2e-11	0.0244	1.7e-11	5.0e-11
		η_s	1.3e-12	6.5e-10	4.0e-12	1.8e-4	2.4e-13	2.9e-12
		time	1.0	0.8	0.7	57.7	0.6	1.0
	#3	η_p	2.8e-9	1.6e-12	9.8e-11	0.0427	2.5e-11	1.4e-16
		η_d	1.5e-11	9.1e-11	3.2e-11	6.4e-4	3.6e-11	5.5e-12
		η_g	1.8e-10	1.1e-13	8.5e-11	0.0384	3.4e-11	2.1e-11
		η_s	2.2e-12	3.9e-15	3.1e-12	2.3e-4	1.4e-12	1.9e-12
		time	1.0	0.6	0.6	58.2	0.6	1.1
$d: 20$ $n: 231$ $m: 16,402$	#1	η_p	3.0e-11	2.2e-9	3.6e-12	0.0800	6.1e-11	1.0e-16
		η_d	1.3e-10	1.3e-10	9.9e-11	4.9e-4	2.9e-11	4.6e-15
		η_g	3.4e-6	4.3e-11	2.3e-11	0.0875	6.4e-11	2.7e-11
		η_s	4.1e-9	4.7e-13	2.5e-12	4.3e-4	1.6e-12	2.5e-12
		time	173.8	99.6	11.2	566.3	2.9	4.4
	#2	η_p	1.5e-11	2.1e-8	2.2e-12	0.1975	1.7e-11	1.2e-16
		η_d	7.2e-11	1.2e-9	1.0e-10	0.0010	7.1e-11	7.9e-14
		η_g	2.9e-6	1.1e-9	2.2e-12	0.2571	1.5e-11	5.5e-13
		η_s	1.0e-8	1.5e-10	4.6e-13	0.0207	7.4e-13	8.6e-14
		time	192.0	98.7	30.7	553.3	3.0	14.7
	#3	η_p	2.2e-10	7.7e-9	9.4e-11	0.0373	8.8e-11	2.7e-16
		η_d	1.2e-10	4.5e-10	4.5e-11	6.3e-4	1.9e-11	4.5e-13
		η_g	2.2e-6	9.3e-11	9.1e-11	0.0323	1.3e-10	4.5e-13
		η_s	2.6e-9	5.1e-11	1.7e-11	3.8e-4	1.5e-12	2.7e-13
		time	174.8	90.3	9.0	551.7	2.7	4.5
$d: 30$ $n: 496$ $m: 77,377$	#1	η_p			8.8e-13	0.2407	4.4e-13	2.1e-16
		η_d	**	**	1.0e-10	7.3e-4	1.0e-12	5.0e-15
		η_g			1.2e-11	0.3913	9.6e-10	7.9e-12
		η_s			2.1e-12	0.0037	9.5e-13	1.3e-12
		time			202.3	1092	15.1	17.3
	#2	η_p			8.5e-13	0.2198	5.8e-11	1.8e-16
		η_d	**	**	1.0e-10	7.3e-4	7.4e-11	4.4e-14
		η_g			4.7e-13	0.3468	7.5e-11	5.7e-13
		η_s			1.0e-13	0.0070	2.6e-13	1.0e-13
		time			98.0	1052	10.6	18.5
	#3	η_p			8.8e-13	0.2637	1.5e-12	1.5e-16
		η_d	**	**	1.0e-10	1.3e-4	8.3e-11	1.1e-12
		η_g			5.2e-12	0.4770	1.3e-11	1.5e-11
		η_s			9.7e-13	0.0196	2.4e-12	2.7e-12
		time			224.0	1046	12.3	16.1
$d: 40$ $n: 861$ $m: 236,202$	#1	η_p			1.3e-12	0.3876	2.4e-12	2.9e-16
		η_d	**	**	1.0e-10	1.8e-14	4.4e-11	5.2e-15
		η_g			5.0e-12	0.7867	7.4e-12	3.9e-12
		η_s			1.0e-12	0.0804	1.1e-12	6.8e-13
		time			413.0	3005	32.7	46.8
	#2	η_p			5.1e-7	0.3401	9.4e-11	2.1e-16
		η_d	**	**	8.3e-6	2.0e-14	7.2e-11	5.5e-13
		η_g			1.0e-5	0.7936	5.9e-10	4.2e-12
		η_s			0.0063	0.1052	7.5e-12	7.2e-12
		time			1690	3034	50.5	76.7
	#3	η_p			9.5e-13	0.3397	1.5e-11	2.1e-16
		η_d	**	**	1.0e-10	2.1e-14	5.4e-11	5.2e-15
		η_g			6.1e-12	0.8206	3.1e-11	6.0e-11
		η_s			1.1e-12	0.1558	3.1e-12	9.8e-12
		time			622.8	3045	32.1	73.4
$d: 50$ $n: 1,326$ $m: 564,877$	#1	η_p			8.0e-11	0.3188	2.6e-13	2.0e-16
		η_d	**	**	1.4e-8	3.2e-14	3.4e-10	5.8e-15
		η_g			8.4e-10	0.8868	1.6e-9	2.7e-11
		η_s			4.0e-8	0.3756	7.2e-11	4.8e-12
		time			4201	5364	156.5	240.2
	#2	η_p			3.9e-12	0.3341	1.1e-12	2.6e-16
		η_d	**	**	1.0e-10	3.3e-14	6.1e-11	5.8e-15
		η_g			3.0e-11	0.8910	6.9e-12	1.7e-12
		η_s			6.4e-12	0.4208	1.2e-12	3.1e-13
		time			1015	5356	75.9	138.5
	#3	η_p			9.4e-13	0.3199	5.3e-14	1.8e-16
		η_d	**	**	1.0e-10	3.1e-14	5.7e-11	5.9e-15
		η_g			2.4e-11	0.9205	6.5e-9	1.9e-12
		η_s			4.2e-12	0.4937	1.4e-10	3.2e-13
		time			1284	5356	113.3	145.9
$d: 60$ $n: 1,891$ $m: 1,155,402$	#1	η_p			2.1e-13	0.3276	3.4e-13	2.1e-16
		η_d	**	**	1.0e-10	3.8e-14	1.8e-10	6.2e-15
		η_g			4.1e-12	0.9610	5.2e-9	3.0e-12
		η_s			7.9e-13	0.6813	8.2e-10	5.7e-13
		time			5059	6895	321.9	283.4
	#2	η_p			3.1e-13	0.3301	4.0e-14	4.2e-16
		η_d	**	**	1.0e-10	5.2e-14	9.9e-11	6.3e-15
		η_g			3.1e-12	0.9281	1.6e-10	5.5e-14
		η_s			7.1e-10	0.5856	3.5e-11	8.8e-15
		time			7684	6905	286.3	303.5
	#3	η_p			6.2e-13	0.3133	7.9e-15	4.5e-16
		η_d	**	**	1.0e-10	4.0e-14	6.1e-11	6.1e-15
		η_g			3.5e-12	0.9572	4.3e-9	3.6e-11
		η_s			6.9e-13	0.5980	2.6e-10	6.6e-12
		time			3669	6909	268.9	311.5

CDCS, in most cases 10-30 times faster. Compared to Table 1, these results suggest that the (Q4S) relaxation is easier to solve than the (BQP) relaxation, perhaps because (Q4S) only has a single unit-norm constraint. (iii) SketchyCGAL, however, failed to solve most of the SDPs to high accuracy, despite taking more time than CDCS and SDPNAL+. (iv) Our solver STRIDE achieved similar performance compared to SDPNAL+. Although STRIDE can be slightly slower than SDPNAL+, it generally attained higher accuracy than SDPNAL+.

5.3 Outlier-robust Wahba problem

Consider the problem of finding the best 3D rotation to align two sets of 3D points while explicitly tolerating *outliers*

$$\min_{q \in \mathcal{S}^3} \sum_{i=1}^N \min \left\{ \frac{\|\tilde{z}_i - q \circ \tilde{w}_i \circ q^{-1}\|^2}{\beta_i^2}, 1 \right\} \quad (40)$$

where $q \in \mathcal{S}^3$ is the *unit quaternion* parametrization of a 3D rotation, $(z_i \in \mathbb{R}^3, w_i \in \mathbb{R}^3)_{i=1}^N$ are given N pairs of 3D points (often normalized to have unit norm), $\tilde{z} \triangleq [z^\top, 0]^\top \in \mathbb{R}^4$ denotes the zero-homogenization of a 3D vector z , $q^{-1} \triangleq [-q_1, -q_2, -q_3, q_4]^\top$ is the inverse quaternion, “ \circ ” denotes the quaternion product defined as

$$q \circ p \triangleq \begin{bmatrix} q_4 & -q_3 & q_2 & q_1 \\ q_3 & q_4 & -q_1 & q_2 \\ -q_2 & q_1 & q_4 & q_3 \\ -q_1 & -q_2 & -q_3 & q_4 \end{bmatrix} p, \quad \forall q, p \in \mathbb{R}^4, \quad (41)$$

$\beta_i > 0$ is a given threshold that determines the maximum *inlier* residual, and $\min \{\cdot, \cdot\}$ realizes the so-called *truncated least squares* (TLS) cost function in robust estimation [6]. Intuitively, the term $q \circ \tilde{w}_i \circ q^{-1}$ is the rotated copy of w_i , and the ℓ_2 norm in (40) measures the Euclidean distance between z_i and w_i after rotation (a metric for the goodness of fit). Problem (40) therefore seeks to find the best 3D rotation that minimizes the sum of (normalized) squared Euclidean distances between z_i and w_i while preventing outliers from damaging the estimation via the usage of the TLS cost function, which assigns a constant value to those pairs of points that cannot be aligned well (*i.e.*, outliers). A pictorial description of the outlier-robust Wahba problem is presented in Fig. 2. Problem (40) is nonsmooth, but can be equivalently reformulated as

$$\min_{\substack{q \in \mathcal{S}^3, \\ \theta_i \in \{+1, -1\}, i=1, \dots, N}} \sum_{i=1}^N \frac{1 + \theta_i}{2} \frac{\|\tilde{z}_i - q \circ \tilde{w}_i \circ q^{-1}\|^2}{\beta_i^2} + \frac{1 - \theta_i}{2} \quad (\text{Wahba})$$

by introducing N binary variables $\{\theta_i\}$ that expose the combinatorial nature. Each θ_i acts as the selection variable for determining whether the i -th pair of 3D points (z_i, w_i) is an inlier or an outlier. Problem (Wahba) is a fundamental problem in aerospace, robotics and computer vision, and is the rotation subproblem in point cloud registration [77, 73].

To solve (Wahba) to global optimality, Yang and Carlone [74] proposed the following semidefinite relaxation that was empirically shown to be always tight. Let

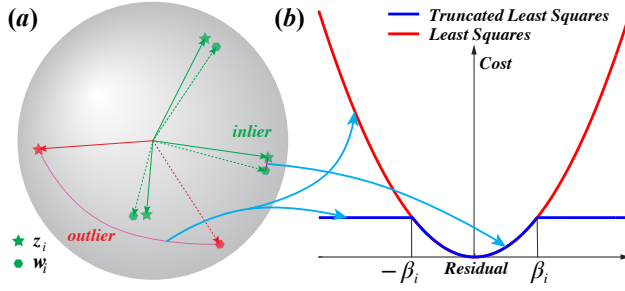


Fig. 2 An example of the outlier-robust (**Wahba**) problem. (a) Four pairs of 3D points (z_i, w_i) lying around a unit sphere, with one of the pairs being an outlier that cannot be aligned well by a 3D rotation. (b) A truncated least squares (TLS) cost function compared with a least squares cost function. The TLS cost function prevents the outlier from contaminating the estimation problem by assigning a constant cost to the outlier. Adapted from [74].

$x = [q^\top, \theta_1, \dots, \theta_N]^\top \in \mathbb{R}^d, d = 4 + N$, be the variable of the nonlinear programming problem (**Wahba**), construct

$$[x]_s = [q^\top, \theta_1 q^\top, \dots, \theta_N q^\top]^\top \in \mathbb{R}^n, \quad n = 4N + 4 \quad (42)$$

as the sparse set of monomials in x of degree up to 2 (a technique that was dubbed *binary cloning*), and then build $X = [x]_s [x]_s^\top$ as the *sparse* moment matrix. Because of the binary constraint $\theta_i^2 = 1$, it can be easily seen that: (i) the diagonal 4×4 blocks of X are all identical ($\theta_i^2 q q^\top = q q^\top$), and (ii) the off-diagonal 4×4 blocks are symmetric ($\theta_i \theta_j q q^\top \in \mathbb{S}^4$). Because of the unit quaternion constraint, X satisfies $\text{tr}(X) = N + 1$. Therefore, this leads to a semidefinite relaxation of size

$$n = 4N + 4, \quad m = 10N + 1 + 3N(N + 1). \quad (43)$$

Compared to the dense second-order moment relaxation in Section 5.1 and 5.2, this sparse second-order relaxation is much more manageable, and the largest N whose relaxation was successfully solved by interior point method was $N = 100$ [74]. This sparse second-order relaxation scheme has been shown as a general framework for certifiable outlier-robust machine perception [76].

Here we show the scalability of our solver by solving instances of (**Wahba**) up to $N = 1000$ and obtaining the globally optimal solution. At each $N = 50, 100, 200, 500, 1000$, we generate three random instances of the (**Wahba**) problem as follows. (i) We draw a random 3D rotation $R \in \text{SO}(3)$ (a rotation matrix can be converted from and to a unit quaternion easily); (ii) we simulate N 3D unit vectors $w_i, i = 1, \dots, N$ uniformly on the unit sphere; (iii) we generate

$$z_i = R w_i + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, 0.01^2), i = 1, \dots, N \quad (44)$$

by rotating w_i and adding Gaussian noise; (iv) we replace 50% of the z_i 's by random unit vectors on the sphere so that they do not follow the generative model (44) and are considered as outliers. We then use SDPT3, MOSEK, CDCS, Sketchy-CGAL, SDPNAL+, and STRIDE to solve the SDP relaxations. For STRIDE, we use Manopt with a trust region solver as the **nlp** method for solving the nonlinear programming (**Wahba**). Specifically, $q \in \mathcal{S}^3$ is modeled as a sphere manifold, and

$\theta \in \{+1, -1\}^N$ is modeled as an oblique manifold of size $1 \times N$ (an oblique manifold of size $n \times m$ is the set of matrices of size $n \times m$ with unit-norm columns), and the problem is treated as an unconstrained problem on the product of two manifolds. To round hypotheses from a moment matrix X , we follow

$$X = \sum_{i=1}^n \lambda_i v_i v_i^\top, \quad q_i = \frac{v_i[q]}{\|v_i[q]\|}, \quad \theta_j^i = \text{sgn}(q_i^\top v_i[\theta_j q]), \quad j = 1, \dots, N, \quad (45)$$

where we first perform spectral decomposition of X with eigenvalues in nonincreasing order, then round q_i by normalizing the corresponding entries of v_i to have unit norm, and finally generate θ_j^i by taking the sign of the dot product between the rounded q_i and the entries of v_i corresponding to each $\theta_j q$ block (the rationale for using this **rounding** method is easily seen from (42) where θ_j^i is identified with $q_i^\top (\theta_j q_i)$ for the rounded q_i with unit norm). We generate $r = 5$ hypotheses by rounding 5 eigenvectors from the moment matrix. We set $M_b = N + 1 = \text{tr}(X)$ to compute η_s as in (35).

Table 3 gives the numerical results for different solvers. Notice that at $N = 1000$, we increased the maximum runtime of SketchyCGAL and SDPNAL+ to be 50000 seconds for a fair comparison with STRIDE. We make the following observations. (i) IPMs can solve small and medium problems ($N = 50$ and 100) to high accuracy and certify global optimality of the POP solutions. However, their runtime grows quickly and they cannot scale to problems with $N \geq 200$. (ii) CDCS, SketchyCGAL and SDPNAL+ perform poorly on this problem. Notably, CDCS and SketchyCGAL failed on all instances and they cannot certify global optimality and tightness of the relaxation. SDPNAL+ succeeded on problems with $N = 50$ but failed to attain high accuracy for all other problems. Comparing Table 3 with Tables 1-2, the degraded performance of CDCS and SDPNAL+ seems to suggest that sparse relaxations are more challenging to solve than dense relaxations. (iii) STRIDE was able to solve all SDP instances to high accuracy. Particularly, for $N = 50$ and $N = 100$, STRIDE achieved similar accuracy compared to MOSEK, while being 3 times faster at $N = 50$ and 20 times faster at $N = 100$. For $N \geq 200$, STRIDE is the only solver that can attain high accuracy and certify global optimality and tightness (despite taking much less time than the other solvers).

STRIDE with Domain-Specific Primal Initialization. In Section 3.1, we mentioned that STRIDE can benefit from domain-specific primal initialization. We now use the (Wahba) problem to support our claim. Although the (Wahba) problem is a combinatorial problem with binary variables, Yang *et al.* [72] have designed a heuristic method called *graduated non-convexity* (GNC) that can solve (Wahba) to global optimality with high probability of success (note that GNC only outputs a solution without optimality certificate). Therefore, we first use GNC to solve the combinatorial (Wahba) problem and then use its solution as a primal initialization for STRIDE. Particularly, let $(\hat{q}, \hat{\theta})$ be the output of GNC, we input a rank-one point $X^0 = [\hat{x}]_s [\hat{x}]_s^\top$ to STRIDE, where $[\hat{x}]_s$ is computed from (42) with $\hat{x} = (\hat{q}, \hat{\theta})$. The last column of Table 3 shows the numerical results for STRIDE with primal initialization supplied by GNC. We can see that the GNC primal initialization gives STRIDE an additional 2-3 times speedup.

Outlier-Robust Wahba Problem on Real Data: To show the practical usefulness of STRIDE, we test it on two applications of the Wahba problem on real data. The first application is image stitching on PASSTA [49] shown in Fig. 3(a).

Table 3 Results on solving sparse second-order relaxation of random (Wahba) instances. “**” indicates solver out of memory. The last column shows results for STRIDE with primal initialization using graduated non-convexity (GNC) [72].

Dimension	Run	Metric	SDPT3 [66]	MOSEK [7]	CDCS [83]	SketchyCGAL [79]	SDPNAL+ [78]	STRIDE	w/ GNC
<i>N</i> : 50 <i>n</i> : 204 <i>m</i> : 8,151	#1	η_p	4.9e-11	1.6e-13	2.4e-6	0.8356	4.6e-14	1.4e-16	1.2e-15
		η_d	1.0e-12	3.7e-15	3.3e-5	0.0354	5.0e-10	2.4e-15	4.6e-15
		η_g	1.2e-8	1.3e-10	0.0020	0.9996	5.2e-6	1.0e-9	4.0e-10
		η_s	9.5e-13	8.9e-15	0.1974	0.9990	2.1e-12	5.4e-14	1.1e-13
		time	65.8	32.7	90.2	302.3	48.8	13.7	4.0
	#2	η_p	5.0e-9	2.2e-12	2.4e-6	0.8372	3.4e-11	2.6e-15	2.6e-15
		η_d	1.3e-12	8.0e-11	3.4e-5	0.0346	3.5e-11	3.6e-15	5.6e-13
		η_g	1.8e-7	8.6e-10	0.0021	0.9996	3.8e-7	8.7e-12	5.9e-9
		η_s	3.6e-12	1.9e-13	0.0887	0.9996	9.6e-14	9.6e-14	1.3e-13
		time	67.4	33.2	87.2	303.8	39.8	13.3	4.2
	#3	η_p	1.3e-9	3.0e-12	2.2e-6	0.8352	4.5e-13	1.4e-16	0.0
		η_d	2.2e-12	3.5e-15	3.2e-5	0.0348	1.7e-10	4.0e-15	3.3e-13
		η_g	3.4e-8	1.3e-9	0.0019	0.9996	1.7e-6	3.9e-9	3.3e-9
		η_s	4.0e-15	5.8e-13	0.0717	0.9996	3.9e-14	3.8e-12	3.4e-14
		time	64.6	34.7	86.7	301.2	44.7	13.5	3.9
<i>N</i> : 100 <i>n</i> : 404 <i>m</i> : 31,301	#1	η_p		1.3e-11	2.4e-7	0.8761	7.1e-15	5.3e-15	5.3e-15
		η_d	**	4.5e-15	2.6e-5	0.0513	9.6e-6	1.6e-13	5.4e-15
		η_g	**	4.1e-11	1.5e-4	0.9999	0.0750	2.3e-9	4.7e-9
		η_s	**	7.0e-14	0.1338	0.9996	0.0219	1.9e-11	1.9e-15
		time	**	974.9	326.8	830.9	526.0	50.7	32.7
	#2	η_p		3.8e-12	3.6e-7	0.8622	1.1e-14	0.0	4.5e-13
		η_d	**	5.6e-15	2.5e-5	0.0518	1.0e-5	5.5e-15	4.9e-15
		η_g	**	5.7e-11	2.3e-4	0.9999	0.0797	8.4e-9	8.3e-10
		η_s	**	7.8e-14	0.3539	0.9999	0.0212	5.5e-14	1.3e-13
		time	**	849.6	321.4	831.7	488.2	46.4	17.7
	#3	η_p		6.7e-12	2.2e-7	0.8570	2.2e-14	1.5e-15	1.7e-15
		η_d	**	4.9e-15	2.5e-5	0.0514	1.1e-5	2.5e-13	5.9e-15
		η_g	**	3.9e-9	1.4e-4	0.9999	0.0882	3.7e-9	2.4e-9
		η_s	**	1.3e-12	0.2188	0.9997	0.0186	1.4e-13	1.4e-13
		time	**	880.7	323.9	830.5	454.1	44.0	16.1
<i>N</i> : 200 <i>n</i> : 804 <i>m</i> : 122,601	#1	η_p			1.0e-6	0.8931	8.9e-15	0.0	9.1e-15
		η_d	**	**	1.6e-5	0.0904	1.2e-5	2.6e-13	2.9e-13
		η_g	**	**	4.8e-4	1.0	0.1086	5.3e-9	1.8e-9
		η_s	**	**	0.3140	0.9999	0.0029	3.9e-11	4.3e-9
		time	**	**	1208	1723	1941	291.9	174.6
	#2	η_p			1.0e-6	0.8913	4.5e-8	5.9e-15	4.5e-15
		η_d	**	**	1.8e-5	0.0919	1.3e-5	2.1e-13	2.6e-13
		η_g	**	**	4.8e-4	1.0	0.0671	1.2e-9	5.2e-9
		η_s	**	**	0.2921	0.9999	0.0136	5.5e-9	4.7e-13
		time	**	**	1206	1646	2349	311.5	185.2
	#3	η_p			9.7e-7	0.8892	8.8e-6	2.6e-13	2.9e-13
		η_d	**	**	1.8e-5	0.0821	1.3e-5	8.2e-15	1.2e-14
		η_g	**	**	4.6e-4	1.0	0.0693	1.7e-9	1.5e-9
		η_s	**	**	0.3027	0.9999	0.0398	4.5e-13	1.8e-11
		time	**	**	1204	1712	2392	272.6	171.7
<i>N</i> : 500 <i>n</i> : 2004 <i>m</i> : 756,501	#1	η_p			4.1e-7	0.8985	3.1e-14	4.3e-14	3.2e-13
		η_d	**	**	1.3e-5	0.1318	1.1e-5	1.5e-14	1.4e-13
		η_g	**	**	1.7e-4	1.0	0.0662	3.8e-11	4.3e-9
		η_s	**	**	0.3385	1.0	0.2480	2.3e-13	1.7e-12
		time	**	**	7659	6035	10001	4389	1819
	#2	η_p			1.8e-7	0.8984	2.6e-4	7.7e-13	1.1e-12
		η_d	**	**	1.5e-5	0.1295	1.1e-5	2.7e-14	5.0e-14
		η_g	**	**	7.4e-5	1.0	0.0780	7.8e-10	1.5e-9
		η_s	**	**	0.3813	1.0	0.1154	2.6e-14	2.1e-13
		time	**	**	7718	6023	10001	2848	1335
	#3	η_p			5.7e-7	0.8984	2.5e-4	2.6e-14	5.5e-14
		η_d	**	**	1.3e-5	0.1142	1.1e-5	1.4e-14	1.7e-14
		η_g	**	**	2.3e-4	1.0	0.0745	1.1e-10	1.2e-11
		η_s	**	**	0.3326	1.0	0.0644	4.4e-15	1.0e-13
		time	**	**	7849	5994	10219	3743	2316
<i>N</i> : 1000 <i>n</i> : 4004 <i>m</i> : 3,013,001	#1	η_p			1.5e-6	0.8993	4.0e-14	1.0e-13	1.4e-14
		η_d	**	**	1.7e-5	0.1747	8.4e-6	3.1e-14	3.9e-14
		η_g	**	**	4.0e-3	1.0	6.5e-3	1.3e-10	3.3e-13
		η_s	**	**	0.5787	1.0	0.3102	1.2e-13	1.2e-9
		time	**	**	62900	42546	50004	30269	26894
	#2	η_p			1.5e-6	0.8994	3.7e-8	8.0e-13	1.1e-16
		η_d	**	**	1.8e-5	0.1316	8.4e-6	3.1e-14	3.3e-14
		η_g	**	**	4.0e-3	1.0	1.2e-3	1.4e-10	1.2e-9
		η_s	**	**	0.5907	1.0	0.3118	2.2e-13	9.8e-13
		time	**	**	60554	42504	44929	50423	41647
	#3	η_p			1.5e-6	0.8994	1.3e-13	2.4e-13	3.4e-15
		η_d	**	**	1.8e-5	0.1212	8.7e-6	2.6e-14	3.6e-14
		η_g	**	**	4.0e-3	1.0	0.0174	2.0e-11	6.2e-10
		η_s	**	**	0.5846	1.0	0.3265	7.3e-14	2.5e-13
		time	**	**	64392	42498	50003	48846	45100

Given two images taken by the same camera with an unknown relative rotation, we first use SURF [9] to establish $N = 70$ putative keypoint matches, and then use STRIDE to solve the SDP relaxation of the Wahba problem ($n = 284, m = 15, 611$) to estimate the relative rotation and stitch the two images. STRIDE obtains the globally optimal solution ($\max\{\eta_p, \eta_d, \eta_g\} = 6.2\text{e-}13, \eta_s = 7.8\text{e-}14$) in 79 seconds. The second application is point cloud registration on 3DMatch [80] shown in Fig. 3(b). Given two point clouds with an unknown relative rotation, we first use FPFH [60] and ROBIN [62] to establish $N = 108$ keypoint matches, and then use STRIDE to solve the SDP relaxation ($n = 436, m = 36, 397$) to estimate the rotation and register the point clouds. STRIDE obtains the globally optimal solution ($\max\{\eta_p, \eta_d, \eta_g\} = 1.6\text{e-}10, \eta_s = 1.8\text{e-}13$) in 53 seconds.

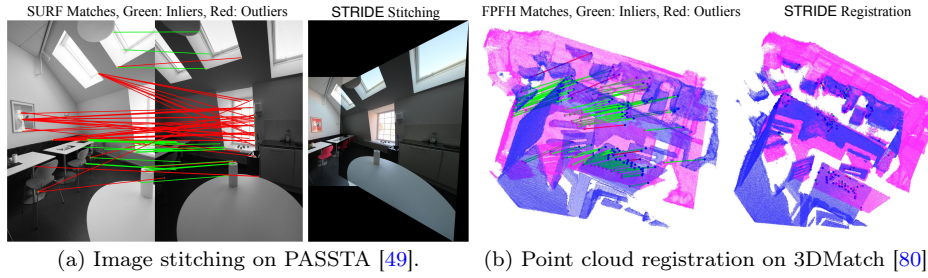


Fig. 3 STRIDE solves outlier-robust Wahba problems on real datasets to global optimality.

5.4 Nearest structured rank deficient matrices

Let N, N_1, N_2 be positive integers with $N_1 \leq N_2$, and let $\mathcal{L} : \mathbb{R}^N \rightarrow \mathbb{R}^{N_1 \times N_2}$ be an affine map. Consider finding the nearest structured rank deficient matrix problem

$$\min_{u \in \mathbb{R}^N} \left\{ \|u - \theta\|^2 \mid \mathcal{L}(u) \text{ is rank deficient} \right\}, \quad (46)$$

where $\theta \in \mathbb{R}^N$ is a given point. Problem (46) is commonly known as the *structured total least squares* (STLS) problem [59, 48], and has numerous applications in control, systems theory, statistics [46], approximate greatest common divisor [35], camera triangulation [3], among others [47]. Problem (46) can be reformulated as the following polynomial optimization problem

$$\min_{z \in \mathcal{S}^{N_1-1}, u \in \mathbb{R}^N} \left\{ \|u - \theta\|^2 \mid z^T \left(L_0 + \sum_{i=1}^N u_i L_i \right) = 0 \right\}, \quad (\text{STLS})$$

where $L_i \in \mathbb{R}^{N_1 \times N_2}, i = 0, \dots, N$, are the set of independent bases of the affine map \mathcal{L} , $z \in \mathcal{S}^{N_1-1}$ is a unit vector in the left kernel of $\mathcal{L}(u)$ and acts as a witness of rank deficiency. Problem (STLS) is easily seen to be nonconvex and the best algorithm in practice is based on local nonlinear programming [47].

Recently, Cifuentes [20] devised a semidefinite relaxation for (STLS) and proved that the relaxation is guaranteed to be tight under a low-noise assumption [21].

This semidefinite relaxation is very similar to the relaxation for (Wahba) presented in Section 5.3 and is also a sparse second-order moment relaxation. Let $x = [z^\top, u^\top]^\top \in \mathbb{R}^d$, $d = N_1 + N$ be the vector of unknowns in (STLS), construct the sparse monomial vector of degree up to 2 in x as

$$[x]_s = [z^\top, u_1 z^\top, \dots, u_N z^\top]^\top \in \mathbb{R}^n, \quad n = (N+1)N_1 \quad (47)$$

and build the moment matrix $X = [x]_s [x]_s^\top$. It can be easily checked that all the off-diagonal $N_1 \times N_1$ blocks, $u_i u_j z z^\top$, are symmetric by construction. Using $[x]_s$, the N_2 equality constraints in (STLS) can be conveniently written as $[x]_s^\top a_i = 0$, $i = 1, \dots, N_2$, for constant vectors $a_i \in \mathbb{R}^n$. In addition, each of the equality constraint also gives rise to n redundant constraints of the form $([x]_s^\top a_i)[x]_s = 0$. Finally, the unit sphere constraint $z \in \mathcal{S}^{N_1-1}$ implies the trace of the leading $N_1 \times N_1$ block of X is equal to 1. The construction above leads to a semidefinite relaxation with size

$$n = (N+1)N_1, \quad m = 1 + nN_2 + \mathfrak{t}(N_1 - 1) \times \mathfrak{t}(N). \quad (48)$$

Due to the limitation of interior point methods, Cifuentes [20] was only able to numerically verify the tightness of the relaxation for very small problems (*e.g.*, $N_1 \leq N_2 \leq 10$, $N < 20$).

We aim to compute globally optimal solutions of (STLS) with much larger dimensions. We perform experiments on random instances of (STLS) where the affine map \mathcal{L} is structured to be a square Hankel matrix such that $N = N_1 + N_2 - 1$, $N_1 = N_2$. We set $N_1 \in \{10, 20, 30, 40\}$, and at each level we randomly generate three problem instances by drawing $\theta \sim \mathcal{N}(0, I_N)$ from the standard Gaussian distribution. We then solve the sparse semidefinite relaxation using SDPT3, MOSEK, CDCS, SketchyCGAL, SDPNAL+, and STRIDE. For STRIDE, since the nonlinear programming (STLS) does not admit any manifold structure, we use `fmincon` with an interior point method as the `nlp` method. To generate hypotheses for `nlp` from the moment matrix X , we follow

$$X = \sum_{i=1}^n \lambda_i v_i v_i^\top, \quad z_i = \frac{v_i[z]}{\|v_i[z]\|}, \quad u_j^i = z_i^\top v_i[u_j z], \quad j = 1, \dots, N, \quad (49)$$

where we first perform a spectral decomposition, then round z_i by projecting the entries of v_i corresponding to block z onto the unit sphere, and round u_j^i by computing the inner product between z_i and the entries of v_i corresponding to block $u_j z$ (again, the rationale for this rounding comes from the lifting (47)). We generate $r = 5$ hypotheses by rounding 5 eigenvectors. In order to set M_b for computing η_s , we make the assumption that the search variable u of (STLS) contains random vectors that follow $\mathcal{N}(0, I_N)$, and its squared norm follows a chi-square distribution of degree N . As a result, we choose the quantile corresponding to a 99.9% probability, denoted as \overline{M} , as the bound on $\|u\|^2$, such that $M_b = \overline{M} + 1$ can upper bound the trace of X .

Table 4 gives the numerical results of different solvers. Notice that at $N_1 = 40$, we have increased the maximum runtime of CDCS and SDPNAL+ to 20000 seconds to make a fair comparison with STRIDE. Generally, Table 4 has similar results as Table 3. (i) IPMs can solve small problems to high accuracy, but cannot handle large problems due to memory issues. (ii) First-order solvers (CDCS and SketchyCGAL) are

Table 4 Results on solving sparse second-order relaxation of random (STLS) instances. “**” indicates solver out of memory.

Dimension	Run	Metric	SDPT3 [66]	MOSEK [7]	CDCS [83]	SketchyCGAL [79]	SDPNAL+ [78]	STRIDE
$N_1: 10$ $n: 200$ $m: 10,551$	#1	η_p	4.8e-12	1.6e-11	3.7e-9	0.3602	1.0e-10	6.2e-16
		η_d	1.2e-10	2.2e-15	1.5e-5	4.3e-14	1.2e-11	6.3e-15
		η_g	5.5e-6	1.1e-12	9.4e-9	9.1e-4	5.3e-9	4.8e-12
		η_s	5.7e-8	8.7e-14	0.0078	0.8808	1.2e-8	1.1e-11
		time	94.0	48.2	138.6	339.7	13.3	14.1
	#2	η_p	1.2e-10	1.8e-10	4.5e-6	0.3787	1.7e-11	1.5e-15
		η_d	1.7e-10	4.1e-11	3.8e-4	2.6e-5	1.0e-10	2.5e-14
		η_g	5.5e-5	5.2e-9	1.9e-4	0.2894	6.3e-10	2.1e-11
		η_s	6.4e-6	2.4e-9	0.1179	0.8544	1.5e-8	2.5e-11
		time	92.2	43.9	140.9	344.2	48.6	18.4
	#3	η_p	1.6e-10	9.0e-12	8.3e-7	0.3943	7.5e-11	8.2e-16
		η_d	2.3e-11	9.7e-9	1.4e-4	1.6e-5	8.9e-11	1.0e-13
		η_g	1.9e-6	1.7e-10	5.3e-5	0.0164	7.3e-11	8.0e-11
		η_s	2.4e-8	1.5e-11	0.1406	0.8993	3.9e-13	1.9e-10
		time	99.2	44.4	139.8	338.1	38.9	17.4
$N_1: 20$ $n: 800$ $m: 164,201$	#1	η_p			2.7e-6	0.9608	3.2e-13	1.9e-15
		η_d			3.4e-4	0.0325	7.8e-10	3.8e-14
		η_g	**	**	2.2e-4	0.1207	2.8e-7	1.8e-10
		η_s			0.3945	0.9863	1.2e-8	2.8e-10
		time			4206	2171	1189	292.5
	#2	η_p			2.5e-6	0.9326	5.6e-13	1.5e-15
		η_d			2.7e-4	0.0109	1.7e-10	2.4e-15
		η_g	**	**	1.6e-4	0.4961	2.0e-10	2.8e-10
		η_s			0.4805	0.9714	9.1e-8	7.2e-10
		time			4187	2163	870.6	257.7
	#3	η_p			4.3e-6	0.9450	5.3e-11	1.6e-15
		η_d			3.4e-4	0.0200	7.9e-11	2.2e-15
		η_g	**	**	3.1e-4	0.1536	1.3e-8	1.2e-10
		η_s			0.4493	0.9798	4.1e-9	2.5e-10
		time			4224	2178	1169	324.4
$N_1: 30$ $n: 1800$ $m: 823,951$	#1	η_p			1.4e-5	1.3862	1.7e-12	5.4e-15
		η_d			1.0e-3	1.3752	2.4e-5	4.9e-14
		η_g	**	**	1.7e-4	0.5175	0.1060	1.6e-9
		η_s			0.87	0.9997	0.2083	6.8e-10
		time			12234	6047	10000	1679
	#2	η_p			1.6e-5	1.3826	2.0e-14	1.6e-8
		η_d			0.0011	1.3544	1.2e-5	2.3e-11
		η_g	**	**	3.3e-4	0.4171	0.1077	2.8e-7
		η_s			0.8248	0.9996	0.0454	2.7e-7
		time			12324	5994	10001	1813
	#3	η_p			1.4e-5	1.3819	2.2e-12	4.2e-15
		η_d			0.0010	1.3376	5.8e-6	2.3e-15
		η_g	**	**	3.7e-4	0.3730	0.0569	6.9e-10
		η_s			0.8191	0.9996	0.0640	4.6e-10
		time			12360	5976	10000	2050
$N_1: 40$ $n: 3200$ $m: 2,592,801$	#1	η_p				1.6887	9.7e-14	5.6e-15
		η_d	**	**	**	2.8283	2.3e-5	2.2e-12
		η_g				0.4648	0.0439	4.3e-10
		η_s				1.0	0.2345	1.3e-7
		time				17222	20003	14541
	#2	η_p				1.6865	2.3e-5	4.9e-15
		η_d	**	**	**	2.8528	2.6e-5	1.0e-13
		η_g				0.4391	0.0750	2.1e-9
		η_s				1.0	0.3296	3.5e-9
		time				17270	20008	9994
	#3	η_p				1.6875	3.5e-4	7.9e-15
		η_d	**	**	**	2.9918	1.4e-4	5.3e-12
		η_g				0.5329	0.1743	2.0e-8
		η_s				1.0	0.6361	2.8e-7
		time				17227	20345	10120

memory efficient but exhibit slow convergence and cannot attain high accuracy to certify global optimality and tightness of the relaxation. SDPNAL+ can solve problems with $N_1 = 10$ and 20 to high accuracy, but failed for $N_1 = 30$ and 40. The degraded performance of CDCS, SketchyCGAL, and SDPNAL+ compared to Tables 1-2 again suggests the difficulty in solving sparse relaxations compared to dense relaxations. (iii) STRIDE computed solutions of high accuracy for all test instances and it is about 5 times faster than SDPNAL+.

6 Concluding Remarks

In this paper, we have designed an efficient algorithmic framework, STRIDE, to solve large-scale tight semidefinite programming (SDP) relaxations of polynomial optimization problems (POPs). STRIDE employs an inexact projected gradient method (iPGM) as its backbone, but leverages fast nonlinear programming (NLP) methods to seek rapid primal acceleration. By conducting a novel convergence analysis for the iPGM and taking safeguarded steps, we have proved the global convergence of STRIDE. For solving the projection subproblem in every iPGM step, we have designed a modified limited-memory BFGS method and proved its global convergence. In addition to the contributions in algorithmic design and convergence analysis, we have studied several important classes of POPs and their dense or sparse (tight) relaxations. In our extensive numerical experiments, STRIDE globally solved all the POPs and the corresponding SDP relaxations to very high accuracy, and it offered state-of-the-art efficiency and robustness. We hope the practical performance of STRIDE can encourage more theoretical study towards investigating when and why the NLP iterates can produce effective acceleration to the iPGM backbone.

Acknowledgements The authors would like to thank Diego Cifuentes for useful discussions about the nearest structured rank deficient matrix problem, and Jie Wang for clarifications about sparse semidefinite relaxations. H. Yang and L. Carlone were partially funded by ARL DCIST CRA W911NF-17-2-0181 and NSF CAREER award “Certifiable Perception for Autonomous Cyber-Physical Systems”.

A Proof of Theorem 1

Let us first prove the following useful lemma for later convenience.

Lemma 1 *Under the inexactness conditions in (6), for any $X \in \mathbb{S}_+^n$, it holds that*

$$\langle C, X - X^k \rangle \geq \frac{1}{\sigma_k} \|X^k - X^{k-1}\|^2 + \frac{1}{\sigma_k} \langle X^{k-1} - X, X^k - X^{k-1} \rangle + \langle \mathcal{A}^* y^k, X - X^k \rangle.$$

Proof From the last two conditions in (6), $\mathcal{A}^* y^k + S^k - C = \frac{1}{\sigma_k} (X^k - X^{k-1})$ and $\langle X^k, S^k \rangle = 0$, we have that

$$\begin{aligned} \langle C, X - X^k \rangle &= \left\langle \mathcal{A}^* y^k + S^k - \frac{1}{\sigma_k} (X^k - X^{k-1}), X - X^k \right\rangle \\ &= \frac{1}{\sigma_k} \langle X^k - X^{k-1}, X^k - X \rangle + \langle \mathcal{A}^* y^k, X - X^k \rangle + \underbrace{\langle S^k, X \rangle}_{\geq 0} - \underbrace{\langle S^k, X^k \rangle}_{=0} \\ &\geq \frac{1}{\sigma_k} \|X^k - X^{k-1}\|^2 + \frac{1}{\sigma_k} \langle X^{k-1} - X, X^k - X^{k-1} \rangle + \langle \mathcal{A}^* y^k, X - X^k \rangle, \end{aligned}$$

where the last inequality is due to the following fact:

$$\begin{aligned} \frac{1}{\sigma_k} \langle X^k - X^{k-1}, X^k - X \rangle &= \frac{1}{\sigma_k} \langle X^k - X^{k-1}, X^k - X^{k-1} + X^{k-1} - X \rangle \\ &= \frac{1}{\sigma_k} \|X^k - X^{k-1}\|^2 + \frac{1}{\sigma_k} \langle X^{k-1} - X, X^k - X^{k-1} \rangle. \end{aligned}$$

Thus, the proof is completed. \square

Then, we shall conduct our proof of Theorem 1 as follows.

Proof For notational simplicity, denote $r^k = \mathcal{A}(X^k) - b$ and $e^k = X^k - X^*$ for all $k \geq 0$. Then, $\mathcal{A}(e^k) = r^k$ since $\mathcal{A}(X^*) = b$. Moreover, we have $\|r^k\| \leq \varepsilon_k$ from the first condition in (6). Therefore, the convergence of $\{\|r^k\|\}$ is obvious, since $\|r^k\| \leq \varepsilon_k$ and $\{k\varepsilon_k\}$ is summable. Recall that X^* is an optimal solution to problem (P). Let $X = X^{k-1}$ in Lemma 1, we get

$$\langle C, X^{k-1} - X^* \rangle \geq \frac{1}{\sigma_k} \|e^k - e^{k-1}\|^2 + \langle \mathcal{A}^* y^k, e^{k-1} - e^k \rangle \quad (50)$$

On the other hand, let $X = X^*$ in Lemma 1, we obtain

$$\langle C, X^* - X^k \rangle \geq \frac{1}{\sigma_k} \|e^k - e^{k-1}\|^2 + \frac{1}{\sigma_k} \langle e^{k-1}, e^k - e^{k-1} \rangle - \langle \mathcal{A}^* y^k, e^k \rangle. \quad (51)$$

Multiplying $k-1$ to (50) and add the resulting inequality to (51) yields

$$\begin{aligned} & \langle C, (k-1)(X^{k-1} - X^*) - k(X^k - X^*) \rangle \\ & \geq \frac{k}{\sigma_k} \|e^k - e^{k-1}\|^2 + \frac{1}{\sigma_k} \langle e^{k-1}, e^k - e^{k-1} \rangle + \langle \mathcal{A}^* y^k, (k-1)e^{k-1} - ke^k \rangle \\ & = \frac{1}{\sigma_k} \left(k\|e^k\|^2 + (k-1)\|e^{k-1}\|^2 - (2k-1)\langle e^k, e^{k-1} \rangle \right) + \langle y^k, (k-1)r^{k-1} - kr^k \rangle \\ & \geq \frac{1}{\sigma_k} \left(k\|e^k\|^2 + (k-1)\|e^{k-1}\|^2 - \frac{2k-1}{2} (\|e^k\|^2 + \|e^{k-1}\|^2) \right) + \langle y^k, (k-1)r^{k-1} - kr^k \rangle \\ & = \frac{1}{2\sigma_k} \|e^k\|^2 - \frac{1}{2\sigma_k} \|e^{k-1}\|^2 + \langle y^k, (k-1)r^{k-1} - kr^k \rangle \\ & \geq \frac{1}{2\sigma_k} \|e^k\|^2 - \frac{1}{2\sigma_{k-1}} \|e^{k-1}\|^2 + \langle y^k, (k-1)r^{k-1} - kr^k \rangle, \end{aligned} \quad (52)$$

where in the last inequality, we use the fact that $\sigma_k \geq \sigma_{k-1} > 0$ for all $k \geq 1$. By summing the inequality in (52) for $k = 1, \dots, k$, we get (recall that $\|y^k\| \leq M$ for all $k \geq 0$)

$$\begin{aligned} -k \langle C, X^k - X^* \rangle & \geq \frac{1}{2\sigma_k} \|e^k\|^2 - \frac{1}{2\sigma_0} \|e_0\|^2 + \sum_{i=1}^k \langle y^i, (i-1)r^{i-1} - ir^i \rangle \\ & \geq \frac{1}{2\sigma_k} \|e^k\|^2 - \frac{1}{2\sigma_0} \|e_0\|^2 - 2M \sum_{i=1}^k i\varepsilon_i. \end{aligned}$$

The above inequality leads to the following upper bound on $\langle C, X^k - X^* \rangle$:

$$\langle C, X^k - X^* \rangle \leq \frac{1}{k} \left(\frac{1}{2\sigma_0} \|e_0\|^2 + 2M \sum_{i=1}^k i\varepsilon_i \right). \quad (53)$$

Next, we shall provide a lower bound for $\langle C, X^k - X^* \rangle$ for $k \geq 1$. To this end, let us define the set \mathcal{F}_k , which is an enlargement of the primal feasible set \mathcal{F}_P with respect to $\varepsilon_k \geq 0$, as follows:

$$\mathcal{F}_k := \{X \in \mathbb{S}^n \mid \|\mathcal{A}(X) - b\| \leq \varepsilon_k, X \in \mathbb{S}_+^n\}, \quad \forall k \geq 1.$$

Moreover, denote

$$X_k^* := \arg \min \{ \langle C, X \rangle \mid X \in \mathcal{F}_k \}, \quad \forall k \geq 1.$$

Since $(y^*, S^*) \in \mathbb{R}^m \times \mathbb{S}_+^n$ is a dual optimal solution, by [33, Lemma 3.4], we have that

$$0 \leq \langle C, X^* - X_k^* \rangle \leq \|y^*\| \varepsilon_k, \quad \forall k \geq 1.$$

As a consequence, because $X^k \in \mathcal{F}_k$ and $\langle C, X^k \rangle \geq \langle C, X_k^* \rangle$, it holds that

$$\langle C, X^k - X^* \rangle \geq \langle C, X_k^* - X^* \rangle \geq -\|y^*\| \varepsilon_k = -\frac{\|y^*\| k \varepsilon_k}{k}. \quad (54)$$

Combining (53) and (54), we see that

$$-\frac{\|y^*\| k \varepsilon_k}{k} \leq \langle C, X^k - X^* \rangle \leq \frac{1}{k} \left(\frac{1}{2\sigma_0} \|e^0\|^2 + 2M \sum_{i=1}^k i \varepsilon_i \right).$$

Finally, let us estimate the dual infeasibility. To this end, we get from (50) that

$$\frac{1}{\sigma_k} \|X^k - X^{k-1}\|^2 \leq \langle C, X^{k-1} - X^* \rangle - \langle C, X^k - X^* \rangle + \|y^k\| (\|r^{k-1}\| + \|r^k\|)$$

which implies that

$$\begin{aligned} \|\mathcal{A}^* y^k + S^k - C\| &= \frac{1}{\sigma_k} \|X^k - X^{k-1}\| \\ &\leq \frac{1}{\sqrt{\sigma_k}} \sqrt{\frac{1}{k-1} \left(\frac{1}{2\sigma_0} \|e^0\|^2 + 2M \sum_{i=1}^{k-1} i \varepsilon_i \right) + \|y^*\| \varepsilon_k + M(\varepsilon_{k-1} + \varepsilon_k)} \\ &\leq O\left(\frac{1}{\sqrt{k\sigma_k}}\right), \end{aligned}$$

and establishes the convergence of the dual infeasibility. Therefore, the proof is completed. \square

B Proof for Theorem 5

From Step 1 of the modified L-BFGS algorithm, we have $d^k = (-\beta_k I + Q_k) \nabla \phi(\xi^k)$, with $Q_k = 0$ if $\|d^k\| \geq K$. This leads to $\langle \nabla \phi(\xi^k), d^k \rangle = \langle \nabla \phi(\xi^k), -(\beta_k I + Q_k) \nabla \phi(\xi^k) \rangle$, which implies that

$$\lambda_{\min}(\beta_k I + Q_k) \leq \frac{\langle -\nabla \phi(\xi^k), d^k \rangle}{\|\nabla \phi(\xi^k)\|^2}, \quad (55)$$

where $\lambda_{\min}(\cdot)$ denotes the minimum eigenvalue. This implies that d^k is a descent direction and the line search scheme is well-defined.

Let the sequence $\{\xi^k\}$ be generated by the algorithm. We next show the second part of the theorem, *i.e.*, $\{\xi^k\}$ is a bounded sequence. Since $\mathcal{A}\mathcal{A}^*$ is nonsingular and the Slater condition holds, by [56, Theorem 17' & 18'], the level set $\mathcal{L}_\phi := \{\xi \in \mathbb{R}^m : \phi(\xi) \leq \phi(\xi^0)\}$ is compact. Thus, the boundedness of the sequence is proven.

Let us now prove the third part of the theorem, which states that every accumulation point of $\{\xi^k\}$ is an optimal solution of problem (28). To this end, without loss of generality, we assume that $\{\xi^k\}$ is an infinite sequence (the case for a finite sequence is trivial and we omit it here) and $\xi^k \rightarrow \hat{\xi}$ as $k \rightarrow \infty$, and we shall prove that $\nabla \phi(\hat{\xi}) = 0$. We plan to prove the latter result by contradiction. Let us assume that $\nabla \phi(\hat{\xi}) \neq 0$. Under this assumption, we will deduce a contradiction in three steps.

Step A. We first claim that when $\nabla \phi(\hat{\xi}) \neq 0$, $\|d^k\|$ is bounded from below and above, *i.e.*, there exist two positive constants M_1 and M_2 such that $0 < M_1 \leq \|d^k\| \leq M_2$ for all k .

On the one hand, by Step 1 of the L-BFGS Algorithm 5, we know that either $d^k = -(\beta_k I + Q_k) \nabla \phi(\xi^k)$, in which case $\|d^k\| \leq K$ (otherwise the algorithm rejects it), or $d^k = -\beta_k \nabla \phi(\xi^k)$, in which case $\|d^k\| \leq \tau_1 M_0^{1+\tau_2}$ (due to $\beta_k = \tau_1 \|\nabla \phi(\xi^k)\|^{\tau_2}$). Therefore, we have $\|d^k\| \leq \max\{K, \tau_1 M_0^{1+\tau_2}\} := M_2 < \infty$, and $\|d^k\|$ is upper bounded.

On the other hand, if for some index set J , $\|d^{k_j}\| \rightarrow 0$ for $j \in J$, then we can deduce that

$$\|\nabla \phi(\xi^{k_j})\| = \|(\beta_{k_j} I + Q_{k_j})^{-1} d^{k_j}\| \leq \frac{1}{\beta_{k_j}} \|d^{k_j}\| = \frac{1}{\tau_1 \|\nabla \phi(\xi^{k_j})\|^{\tau_2}} \|d^{k_j}\|, \quad j \in J.$$

The above inequality implies that $\|\nabla \phi(\hat{\xi})\| = \lim_{j \in J} \|\nabla \phi(\xi^{k_j})\| = 0$, a contradiction.

Thus, $\|d^k\|$ must also be lower bounded, by some constant $M_1 > 0$.

Step B. We next claim that $\{\rho^{m_k}\}$ is bounded away from zero.

Suppose that, for some index set J , $\rho^{m_{k_j}} \rightarrow 0$, $j \in J$ (i.e., $m_{k_j} \rightarrow \infty$ due to $\rho \in (0, 1)$). Recall that m_{k_j} is the smallest nonnegative integer to satisfy the inequality in the line search scheme (so that $m_{k_j} - 1$ does not satisfy the inequality). Thus, it holds that

$$\phi(\xi^{k_j} + \rho^{m_{k_j}-1} d^{k_j}) > \phi(\xi^{k_j}) + \mu \rho^{m_{k_j}-1} \langle \nabla \phi(\xi^{k_j}), d^{k_j} \rangle. \quad (56)$$

Meanwhile, we can expand $\phi(\xi^{k_j} + \rho^{m_{k_j}-1} d^{k_j})$ as

$$\phi(\xi^{k_j} + \rho^{m_{k_j}-1} d^{k_j}) = \phi(\xi^{k_j}) + \int_0^1 \langle \nabla \phi(\xi^{k_j} + t \rho^{m_{k_j}-1} d^{k_j}), \rho^{m_{k_j}-1} d^{k_j} \rangle dt. \quad (57)$$

By substituting (57) into (56), we obtain the following inequality:

$$\int_0^1 \langle \nabla \phi(\xi^{k_j} + t \rho^{m_{k_j}-1} d^{k_j}), d^{k_j} \rangle dt > \mu \langle \nabla \phi(\xi^{k_j}), d^{k_j} \rangle, \quad j \in J. \quad (58)$$

Since the sequence $\{d^{k_j}\}$ is bounded from both below and above, by taking a subset of J if necessary, we assume that there exists \hat{d} such that $d^{k_j} \rightarrow \hat{d}$, $j \in J$. Taking limit with respect to $j \in J$ in (58) yields that $\langle \nabla \phi(\hat{\xi}), \hat{d} \rangle \geq \mu \langle \nabla \phi(\hat{\xi}), \hat{d} \rangle$, which further implies that $\langle \nabla \phi(\hat{\xi}), \hat{d} \rangle \geq 0$ by the fact that $\mu \in (0, 1/2)$. Also, since $\langle \nabla \phi(\hat{\xi}), \hat{d} \rangle \leq 0$, we derive $\langle \nabla \phi(\hat{\xi}), \hat{d} \rangle = 0$. However, denoting $\hat{\beta}$ as the limit of β_{k_j} as $k_j \rightarrow \infty$, $j \in J$, and noting that $Q_k \succeq 0$ for all k , the inequality in (55) shows that

$$0 = -\langle \nabla \phi(\hat{\xi}), \hat{d} \rangle \geq \hat{\beta} \|\nabla \phi(\hat{\xi})\|^2 \geq \tau_1 \|\nabla \phi(\hat{\xi})\|^{2+\tau_2},$$

which implies that $\nabla \phi(\hat{\xi}) = 0$. This is also a contradiction. Therefore, $\{\rho^{m_k}\}$ is bounded away from zero.

Step C. Now, using the fact that ϕ is bounded from below on the bounded sequence $\{\xi^k\}$, we have that $\{\phi(\xi^{k+1}) - \phi(\xi^k)\} \rightarrow 0$ (since the sequence $\{\phi(\xi^k)\}$ is nonincreasing and $\phi(\cdot)$ is continuous). This implies, by the line search rule again, that $\rho^{m_k} \langle \nabla \phi(\xi^k), d^k \rangle \rightarrow 0$, as $k \rightarrow \infty$. Therefore, it holds that $\langle \nabla \phi(\hat{\xi}), \hat{d} \rangle = 0$, since $\{\rho^{m_k}\}$ is bounded away from zero, which again implies that $\nabla \phi(\hat{\xi}) = 0$, and this contradicts our initial hypothesis that $\nabla \phi(\hat{\xi}) \neq 0$.

At this point, we have shown that $\nabla \phi(\hat{\xi}) = 0$, as desired. At last, by the convexity of $\phi(\cdot)$, vanishing of the gradient at $\hat{\xi}$ implies that $\hat{\xi}$ is an optimal solution of problem (28). Therefore, the proof is completed. \square

References

1. Abbe, E., Bandeira, A.S., Hall, G.: Exact recovery in the stochastic block model. *IEEE Transactions on Information Theory* **62**(1), 471–487 (2015) [3](#)
2. Absil, P.A., Mahony, R., Sepulchre, R.: *Optimization algorithms on matrix manifolds*. Princeton University Press (2009) [13](#)
3. Aholt, C., Agarwal, S., Thomas, R.: A QCQP approach to triangulation. In: *European Conference on Computer Vision*, pp. 654–667. Springer (2012) [32](#)
4. Alizadeh, F., Haeberly, J.P.A., Overton, M.L.: Complementarity and nondegeneracy in semidefinite programming. *Mathematical programming* **77**(1), 111–128 (1997) [4](#)
5. Alizadeh, F., Haeberly, J.P.A., Overton, M.L.: Primal-dual interior-point methods for semidefinite programming: convergence rates, stability and numerical results. *SIAM Journal on Optimization* **8**(3), 746–768 (1998) [4](#)
6. Antonante, P., Tzoumas, V., Yang, H., Carlone, L.: Outlier-robust estimation: Hardness, minimally-tuned algorithms, and applications. *IEEE Trans. Robotics* (2021) [28](#)

7. ApS, M.: The MOSEK optimization toolbox for MATLAB manual. Version 9.0. (2019). URL <http://docs.mosek.com/9.0/toolbox/index.html> 4, 22, 25, 27, 31, 34
8. Barak, B., Brandao, F.G., Harrow, A.W., Kelner, J., Steurer, D., Zhou, Y.: Hypercontractivity, sum-of-squares proofs, and their applications. In: Proceedings of the forty-fourth annual ACM symposium on Theory of computing, pp. 307–326 (2012) 26
9. Bay, H., Tuytelaars, T., Van Gool, L.: Surf: Speeded up robust features. In: European conference on computer vision, pp. 404–417. Springer (2006) 32
10. Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences* 2(1), 183–202 (2009) 17
11. Bertsekas, D.: *Nonlinear Programming*. Athena Scientific (1999) 7
12. Blekherman, G., Parrilo, P.A., Thomas, R.R.: *Semidefinite optimization and convex algebraic geometry*. SIAM (2012) 6
13. Boumal, N., Mishra, B., Absil, P.A., Sepulchre, R.: Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research* 15(42), 1455–1459 (2014). URL <https://www.manopt.org> 13, 26
14. Boumal, N., Voroninski, V., Bandeira, A.S.: The non-convex Burer-Monteiro approach works on smooth semidefinite programs. In: Conference on Neural Information Processing Systems (NeurIPS) (2016) 5
15. Briales, J., Gonzalez-Jimenez, J.: Convex global 3d registration with lagrangian duality. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4960–4969 (2017) 3, 11
16. Burer, S., Monteiro, R.D.: A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming* 95(2), 329–357 (2003) 4
17. Burer, S., Ye, Y.: Exact semidefinite formulations for a class of (random and non-random) nonconvex quadratic programs. *Mathematical Programming* pp. 1–17 (2019) 3
18. Candes, E.J., Eldar, Y.C., Strohmer, T., Voroninski, V.: Phase retrieval via matrix completion. *SIAM review* 57(2), 225–251 (2015) 3
19. Chen, L., Sun, D., Toh, K.: An efficient inexact symmetric gauss-seidel based majorized ADMM for high-dimensional convex composite conic programming. *Mathematical Programming* 161, 237–270 (2017) 14
20. Cifuentes, D.: A convex relaxation to compute the nearest structured rank deficient matrix. *SIAM Journal on Matrix Analysis and Applications* 42(2), 708–729 (2021) 4, 32, 33
21. Cifuentes, D., Agarwal, S., Parrilo, P.A., Thomas, R.R.: On the local stability of semidefinite relaxations. *Mathematical Programming* (2021) 3, 32
22. Combettes, P.L., Pesquet, J.C.: *Proximal splitting methods in signal processing*. In: Fixed-point algorithms for inverse problems in science and engineering, pp. 185–212. Springer (2011) 15
23. Dai, Y.H., Yuan, Y.: A nonlinear conjugate gradient method with a strong global convergence property. *SIAM Journal on optimization* 10(1), 177–182 (1999) 20
24. De Klerk, E.: The complexity of optimizing over a simplex, hypercube or sphere: a short survey. *Central European Journal of Operations Research* 16(2), 111–125 (2008) 26
25. Doherty, A.C., Parrilo, P.A., Spedalieri, F.M.: Complete family of separability criteria. *Physical Review A* 69(2), 022308 (2004) 4, 26
26. Fang, K., Fawzi, H.: The sum-of-squares hierarchy on the sphere and applications in quantum information theory. *Mathematical Programming* pp. 1–30 (2020) 26
27. Fukuda, M., Kojima, M., Murota, K., Nakata, K.: Exploiting sparsity in semidefinite programming via matrix completion i: General framework. *SIAM Journal on optimization* 11(3), 647–674 (2001) 4
28. Gamarnik, D., Kızıldağ, E.C.: Algorithmic obstructions in the random number partitioning problem. *arXiv preprint arXiv:2103.01369* (2021) 23
29. Goemans, M.X., Williamson, D.P.: Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)* 42(6), 1115–1145 (1995) 3, 11, 23
30. Helmberg, C., Rendl, F., Weismantel, R.: A semidefinite programming approach to the quadratic knapsack problem. *Journal of combinatorial optimization* 4(2), 197–215 (2000) 23
31. Henrion, D., Lasserre, J.B.: GloptiPoly: Global optimization over polynomials with Matlab and SeDuMi. *ACM Transactions on Mathematical Software (TOMS)* 29(2), 165–194 (2003) 4

32. Inc., W.R.: Mathematica, Version 12.2. URL <https://www.wolfram.com/mathematica>. Champaign, IL, 2020 [6](#)
33. Jiang, K., Sun, D., Toh, K.C.: An inexact accelerated proximal gradient method for large scale linearly constrained convex sdp. *SIAM Journal on Optimization* **22**(3), 1042–1064 (2012) [8](#), [10](#), [13](#), [36](#)
34. Jozs, C., Henrion, D.: Strong duality in Lasserre’s hierarchy for polynomial optimization. *Optimization Letters* **10**(1), 3–10 (2016) [3](#)
35. Kaltofen, E., Yang, Z., Zhi, L.: Approximate greatest common divisors of several polynomials with linearly constrained coefficients and singular polynomials. In: *Proceedings of the 2006 international symposium on Symbolic and algebraic computation*, pp. 169–176 (2006) [32](#)
36. Lasserre, J.B.: An explicit exact SDP relaxation for nonlinear 0-1 programs. In: *International Conference on Integer Programming and Combinatorial Optimization*, pp. 293–303. Springer (2001) [4](#), [23](#)
37. Lasserre, J.B.: Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization* **11**(3), 796–817 (2001) [2](#), [3](#)
38. Lasserre, J.B.: *Moments, positive polynomials and their applications*, vol. 1. World Scientific (2009) [2](#)
39. Laurent, M.: Semidefinite representations for finite varieties. *Mathematical programming* **109**(1), 1–26 (2007) [3](#)
40. Li, X., Sun, D., Toh, K.C.: A block symmetric Gauss–Seidel decomposition theorem for convex composite quadratic programming and its applications. *Mathematical Programming* **175**(1), 395–418 (2019) [17](#), [18](#), [19](#)
41. Ling, C., Nie, J., Qi, L., Ye, Y.: Biquadratic optimization over unit spheres and semidefinite programming relaxations. *SIAM Journal on Optimization* **20**(3), 1286–1310 (2010) [26](#)
42. Luo, Z.Q., Ma, W.K., So, A.M.C., Ye, Y., Zhang, S.: Semidefinite relaxation of quadratic optimization problems. *IEEE Signal Processing Magazine* **27**(3), 20–34 (2010) [3](#)
43. Mai, N.H.A., Magron, V., Lasserre, J.B.: A hierarchy of spectral relaxations for polynomial optimization. *arXiv preprint arXiv:2007.09027* (2020) [23](#)
44. Malick, J., Povh, J., Rendl, F., Wiegale, A.: Regularization methods for semidefinite programming. *SIAM Journal on Optimization* **20**(1), 336–356 (2009) [20](#)
45. Malick, J., Sendov, H.S.: Clarke generalized Jacobian of the projection onto the cone of positive semidefinite matrices. *Set-Valued Analysis* **14**(3), 273–293 (2006) [20](#)
46. Markovsky, I.: Structured low-rank approximation and its applications. *Automatica* **44**(4), 891–909 (2008) [32](#)
47. Markovsky, I., Usevich, K.: Software for weighted structured low-rank approximation. *Journal of Computational and Applied Mathematics* **256**, 278–292 (2014) [32](#)
48. Markovsky, I., Van Huffel, S.: Overview of total least-squares methods. *Signal processing* **87**(10), 2283–2302 (2007) [32](#)
49. Meneghetti, G., Danelljan, M., Felsberg, M., Nordberg, K.: Image alignment for panorama stitching in sparsely structured environments. In: *Scandinavian Conference on Image Analysis*, pp. 428–439. Springer (2015) [30](#), [32](#)
50. Mertens, S.: Number partitioning. *Computational Complexity and Statistical Physics* p. 125 (2006) [23](#)
51. Nesterov, Y.: *Lectures on convex optimization*, vol. 137. Springer (2018) [20](#)
52. Nie, J.: Polynomial optimization with real varieties. *SIAM Journal On Optimization* **23**(3), 1634–1646 (2013) [3](#)
53. Nie, J.: Optimality conditions and finite convergence of Lasserre’s hierarchy. *Mathematical programming* **146**(1), 97–121 (2014) [3](#), [4](#)
54. Nocedal, J., Wright, S.: *Numerical Optimization*. Springer Science & Business Media (2006) [6](#), [13](#), [20](#)
55. Parrilo, P.A.: Semidefinite programming relaxations for semialgebraic problems. *Mathematical programming* **96**(2), 293–320 (2003) [2](#), [3](#)
56. Rockafellar, R.T.: *Conjugate duality and optimization*. SIAM (1974) [37](#)
57. Rosen, D.M.: Scalable low-rank semidefinite programming for certifiably correct machine perception. In: *Intl. Workshop on the Algorithmic Foundations of Robotics (WAFR)*, vol. 3 (2020) [5](#)
58. Rosen, D.M., Carlone, L., Bandeira, A.S., Leonard, J.J.: SE-Sync: A certifiably correct algorithm for synchronization over the special Euclidean group. *The International Journal of Robotics Research* **38**(2-3), 95–125 (2019) [3](#), [5](#)

59. Rosen, J.B., Park, H., Glick, J.: Total least norm formulation and solution for structured problems. *SIAM Journal on matrix analysis and applications* **17**(1), 110–126 (1996) [32](#)
60. Rusu, R., Blodow, N., Beetz, M.: Fast point feature histograms (fpfh) for 3d registration. In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 3212–3217. Citeseer (2009) [32](#)
61. Shi, J., Yang, H., Carlone, L.: Optimal pose and shape estimation for category-level 3d object perception. In: *Robotics: Science and Systems (RSS)* (2021) [3](#)
62. Shi, J., Yang, H., Carlone, L.: ROBIN: a graph-theoretic approach to reject outliers in robust estimation using invariants. In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)* (2021) [32](#)
63. Shor, N.Z.: Dual quadratic estimates in polynomial and boolean programming. *Annals of Operations Research* **25**(1), 163–168 (1990) [3](#)
64. Sun, D., Toh, K.C., Yang, L.: A convergent 3-block semiproximal alternating direction method of multipliers for conic programming with 4-type constraints. *SIAM journal on Optimization* **25**(2), 882–915 (2015) [14](#)
65. Toh, K.C.: Solving large scale semidefinite programs via an iterative solver on the augmented systems. *SIAM Journal on Optimization* **14**(3), 670–698 (2004) [24](#)
66. Toh, K.C., Todd, M.J., Tütüncü, R.H.: SDPT3—a MATLAB software package for semidefinite programming, version 1.3. *Optimization methods and software* **11**(1-4), 545–581 (1999) [3](#), [4](#), [22](#), [25](#), [27](#), [31](#), [34](#)
67. Waki, H., Kim, S., Kojima, M., Muramatsu, M.: Sums of squares and semidefinite programming relaxations for polynomial optimization problems with structured sparsity. *SIAM J. Optimization* **17**, 218–242 (2006) [3](#)
68. Wang, A.L., Kılınç-Karzan, F.: On the tightness of SDP relaxations of QCQPs. *Mathematical Programming* pp. 1–41 (2021) [3](#)
69. Wang, J., Magron, V., Lasserre, J.B.: Chordal-TSSOS: a moment-SOS hierarchy that exploits term sparsity with chordal extension. *SIAM Journal on Optimization* **31**(1), 114–141 (2021) [3](#)
70. Wang, J., Magron, V., Lasserre, J.B.: TSSOS: A Moment-SOS hierarchy that exploits term sparsity. *SIAM Journal on Optimization* **31**(1), 30–58 (2021) [3](#)
71. Wu, G., Sun, J., Chen, J.: Optimal linear quadratic regulator of switched systems. *IEEE transactions on automatic control* **64**(7), 2898–2904 (2018) [23](#)
72. Yang, H., Antonante, P., Tzoumas, V., Carlone, L.: Graduated non-convexity for robust spatial perception: From non-minimal solvers to global outlier rejection. *IEEE Robotics and Automation Letters (RA-L)* **5**(2), 1127–1134 (2020) [30](#), [31](#)
73. Yang, H., Carlone, L.: A polynomial-time solution for robust registration with extreme outlier rates. In: *Robotics: Science and Systems (RSS)* (2019) [28](#)
74. Yang, H., Carlone, L.: A quaternion-based certifiably optimal solution to the Wahba problem with outliers. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1665–1674 (2019) [4](#), [11](#), [28](#), [29](#)
75. Yang, H., Carlone, L.: In perfect shape: Certifiably optimal 3d shape reconstruction from 2d landmarks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 621–630 (2020) [4](#)
76. Yang, H., Carlone, L.: One ring to rule them all: Certifiably robust geometric perception with outliers. In: *Conference on Neural Information Processing Systems (NeurIPS)* (2020) [4](#), [11](#), [29](#)
77. Yang, H., Shi, J., Carlone, L.: Teaser: Fast and certifiable point cloud registration. *IEEE Transactions on Robotics* (2020) [28](#)
78. Yang, L., Sun, D., Toh, K.C.: SDPNAL+: a majorized semismooth Newton-CG augmented Lagrangian method for semidefinite programming with nonnegative constraints. *Mathematical Programming Computation* **7**(3), 331–366 (2015) [4](#), [14](#), [22](#), [25](#), [27](#), [31](#), [34](#)
79. Yurtsever, A., Tropp, J.A., Fercoq, O., Udell, M., Cevher, V.: Scalable semidefinite programming. *SIAM Journal on Mathematics of Data Science* **3**(1), 171–200 (2021) [4](#), [22](#), [25](#), [27](#), [31](#), [34](#)
80. Zeng, A., Song, S., Nießner, M., Fisher, M., Xiao, J., Funkhouser, T.: 3dmatch: Learning the matching of local 3d geometry in range scans. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, vol. 1, p. 4 (2017) [32](#)
81. Zhang, R.Y., Lavaei, J.: Sparse semidefinite programs with guaranteed near-linear time complexity via dualized clique tree conversion. *Mathematical programming* **188**(1), 351–393 (2021) [4](#)

-
82. Zhao, X.Y., Sun, D., Toh, K.C.: A Newton-CG augmented Lagrangian method for semidefinite programming. *SIAM Journal on Optimization* **20**(4), 1737–1765 (2010) [4](#), [20](#)
 83. Zheng, Y., Fantuzzi, G., Papachristodoulou, A., Goulart, P., Wynn, A.: Chordal decomposition in operator-splitting methods for sparse semidefinite programs. *Mathematical Programming* **180**(1), 489–532 (2020) [4](#), [22](#), [25](#), [27](#), [31](#), [34](#)