

Outer Approximation With Conic Certificates For Mixed-Integer Convex Problems*

Chris Coey[§], Miles Lubin[§], and Juan Pablo Vielma[‡]

[§]MIT Operations Research Center

[‡]MIT Sloan School of Management

coey@mit.edu, miles.lubin@gmail.com, jvielma@mit.edu

August 17, 2018

Abstract

A mixed-integer convex (MI-convex) optimization problem is one that becomes convex when all integrality constraints are relaxed. We present a branch-and-bound LP outer approximation algorithm for an MI-convex problem transformed to *MI-conic* form. The polyhedral relaxations are refined with \mathcal{K}^* cuts derived from *conic certificates* for continuous primal-dual conic subproblems. Under the assumption that all subproblems are *well-posed*, the algorithm detects infeasibility or unboundedness or returns an optimal solution in finite time. Using properties of the conic certificates, we show that the \mathcal{K}^* cuts imply certain practically-relevant guarantees about the quality of the polyhedral relaxations, and demonstrate how to maintain helpful guarantees when the LP solver uses a positive feasibility tolerance. We discuss how to *disaggregate* \mathcal{K}^* cuts in order to tighten the polyhedral relaxations and thereby improve the speed of convergence, and propose fast heuristic methods of obtaining useful \mathcal{K}^* cuts. Our new open source MI-conic solver *Pajarito* (github.com/JuliaOpt/Pajarito.jl) uses an external mixed-integer linear (MILP) solver to manage the search tree and an external continuous conic solver for subproblems. Benchmarking on a library of mixed-integer second-order cone (MISOCP) problems, we find that Pajarito greatly outperforms Bonmin (the leading open source alternative) and is competitive with CPLEX's specialized MISOCP algorithm. We demonstrate the robustness of Pajarito by solving diverse MI-conic problems involving mixtures of positive semidefinite, second-order, and exponential cones, and provide evidence for the practical value of our analyses and enhancements of \mathcal{K}^* cuts.

*M. Lubin is now at Google. The authors thank MOSEK for access to alpha releases of MOSEK 9 and Russell Bent and Emre Yamangil for support for and early contributions to Pajarito. This work has been partially funded by the National Science Foundation under grant CMMI-1351619, the Department of Energy Computational Science Graduate Fellowship under grant DE-FG02-97ER25308, and the Office of Naval Research under grant N00014-18-1-2079.

Contents

1	Mixed-Integer Convex Optimization	3
1.1	Branch-And-Bound Algorithms	3
1.2	Mixed-Integer Conic Form	4
1.3	Overview And Contributions	5
2	A Branch-And-Bound LP Outer Approximation Algorithm	6
2.1	Continuous Subproblems And Conic Duality	7
2.2	Dynamic Polyhedral Relaxations	9
2.3	The Conic-Certificate-Based Algorithm	10
3	Polyhedral Relaxation Guarantees From Conic Certificates	13
3.1	Under An Exact LP Solver	13
3.2	Under An LP Solver With A Feasibility Tolerance	15
4	Tightening Polyhedral Relaxations	16
4.1	Extreme Ray Disaggregation	17
4.2	Initial Fixed Polyhedral Relaxations	18
4.3	Separation Of Infeasible Points	19
5	Pajarito Solver And Related Software	19
5.1	Julia And MathProgBase	20
5.2	Accessing Pajarito	20
5.3	Basic Algorithmic Implementations	23
5.4	Some Advanced Algorithmic Enhancements	27
6	Computational Experiments	28
6.1	Presentation Of Results	28
6.2	MISOCP Solver Performance Comparisons	30
6.3	Comparative Testing Of Algorithmic Variants	32
A	The Standard Primitive Nonpolyhedral Cones	38
A.1	Exponential Cone	38
A.2	Second-Order Cone	39
A.3	Positive Semidefinite Cone	40
B	The Second-Order Cone Extended Formulation	41
C	SOCQP Outer Approximation For PSD Cones	42
	Bibliography	44

1 Mixed-Integer Convex Optimization

A *mixed-integer convex* (MI-convex) problem is a finite-dimensional optimization problem that minimizes a convex objective function over convex constraints and integrality restrictions on a subset of the variables. Belotti et al. [2013] and Bonami et al. [2012] review MI-convex applications and Lubin et al. [2017b] characterize which nonconvex feasible regions are MI-convex-representable. Since an MI-convex problem without integrality restrictions is just a convex problem, MI-convex optimization generalizes both mixed-integer linear optimization (MILP) and convex optimization. This structure also leads to effective branch-and-bound (B&B) algorithms, which recursively partition the possible values of the integer variables in a search tree and obtain objective bounds and feasible solutions from efficiently-solvable subproblems.

1.1 Branch-And-Bound Algorithms

A *nonlinear B&B* (B&B-NL) algorithm for a MI-convex problem solves a nonlinear subproblem that includes all of the convex constraints at every node of the search tree. The Bonmin solver package [Bonami et al., 2008] implements a B&B-NL variant by calling the derivative-oracle-based nonlinear programming (NLP) solver Ipopt to solve the subproblems. The relatively new SCIP-SDP [Gally et al., 2018] B&B-NL implementation for mixed-integer semidefinite (MISDP) problems uses a primal-dual conic interior-point solver for the SDP subproblems.

Typically, B&B-NL methods need to solve a large number of very similar nonlinear subproblems to near-global optimality and feasibility in order to obtain accurate objective bounds. Linear optimization (LP) solvers based on the Simplex algorithm are able to rapidly reoptimize after variable bounds are changed or linear cuts are added, thus typically benefiting from warm-starting much more so than state-of-the-art NLP or conic solvers. *B&B LP outer approximation* (B&B-OA) algorithms take advantage of LP warm-starting by solving a polyhedral relaxation of the nonlinear subproblem at every node. Implementations often take advantage of the speed and stability of advanced MILP branch-and-cut solvers.

B&B-OA algorithms differ in how they refine the polyhedral relaxations of the nonlinear constraints and how they obtain feasible solutions. In a *separation-based* algorithm, no nonlinear solver is used. At each node the LP optimal point is first checked for feasibility for the convex constraints; if the violation exceeds a positive tolerance, valid cuts separating the point are added to the LP, otherwise the point may be accepted as a new incumbent if it is integral.¹ Quesada and Grossmann [1992] and Leyffer [1993] describe *subproblem-based* B&B-OA algorithms that solve smooth subproblems at a subset of the nodes.

¹Commercial mixed-integer second-order cone optimization (MISOCP) solvers use separation-based algorithms, but also occasionally solve SOCP subproblems to obtain feasible solutions and fathom nodes. SCIP-SDP offers both B&B-NL and separation-based B&B-OA methods for mixed-integer semidefinite problems.

The subproblems provide points at which cuts based on gradient inequalities can be derived.²

Bonami et al. [2008] found that Bonmin’s B&B-OA method generally outperforms its B&B-NL method. Since both of these methods rely on NLP subproblems, they frequently fail in the presence of nonsmoothness. Continuous conic solvers are more numerically robust than derivative-oracle-based NLP solvers on nonsmooth problems (such as SOCPs and SDPs). For the special case of MISOCP, Drewes and Ulbrich [2012] propose a conic subproblem-based B&B-OA algorithm that derives cuts from subgradients satisfying subproblem KKT optimality conditions, and hence does not require smoothness assumptions.

Another advantage of conic solvers is that they return simple *certificates* proving primal or dual infeasibility or optimality of a primal-dual solution pair [Permenter et al., 2015]. Using the theory of conic duality, it is possible to describe an elegant OA algorithm for generic MI-conic problems that uses conic certificates returned by primal-dual conic solvers, with no need to examine KKT conditions or solve a second modified subproblem in the case of infeasibility (as in the algorithm by Drewes and Ulbrich [2012]). Lubin et al. [2016] propose this idea in an iterative OA algorithm. However a B&B algorithm using a single search tree, instead of solving a sequence of MILP instances each with their own search tree, is more flexible and likely to be significantly faster in practice. We fill this gap with the first conic-certificate-based B&B-OA algorithm.

1.2 Mixed-Integer Conic Form

We use the following general form for a mixed-integer conic (MI-conic) problem:

$$\mathfrak{M} \begin{cases} \inf_{\mathbf{x} \in \mathbb{R}^N} \mathbf{c}^T \mathbf{x} : & (2a) \\ \mathbf{b} - \mathbf{A}\mathbf{x} \in \mathcal{K} \subset \mathbb{R}^M & (2b) \\ x_i \in \mathbb{Z} & \forall i \in \llbracket I \rrbracket, & (2c) \end{cases}$$

where \mathcal{K} is a *closed convex cone*, i.e. a closed subset of \mathbb{R}^M that contains all conic (nonnegative) combinations of its points [Ben-Tal and Nemirovski, 2001a]:

$$\alpha_1 \mathbf{y}_1 + \alpha_2 \mathbf{y}_2 \in \mathcal{K} \quad \forall \alpha_1, \alpha_2 \geq 0 \quad \forall \mathbf{y}_1, \mathbf{y}_2 \in \mathcal{K}. \quad (3)$$

The decision variables in \mathfrak{M} are represented by the column vector $\mathbf{x} \in \mathbb{R}^N$, so the objective (2a) minimizes a linear function of \mathbf{x} subject to (denoted by ‘:’) the constraints (2b) and (2c). The index set of integer decision variables is $\llbracket I \rrbracket = \{1, \dots, I\}$, so the integrality constraints (2c) restrict only the first I variables x_1, \dots, x_I to the set of integers \mathbb{Z} . The conic constraint (2b), which restricts the affine transformation $\mathbf{b} - \mathbf{A}\mathbf{x}$ of \mathbf{x} to \mathcal{K} , is a convex constraint, so

²For a convex function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, the set $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^n : f(\mathbf{x}) \leq 0\}$ is convex. If f is smooth, then given a point $\bar{\mathbf{x}} \in \mathbb{R}^n$, the following *gradient cut* yields a polyhedral relaxation of \mathcal{X} :

$$f(\bar{\mathbf{x}}) + (\nabla f(\bar{\mathbf{x}}))^T (\mathbf{x} - \bar{\mathbf{x}}) \leq 0. \quad (1)$$

relaxing the integrality constraints (2c) results in a convex conic optimization problem.

Any MI-convex problem can be expressed in MI-conic form, by homogenizing the convex constraints, for example through perspective transformations [Boyd and Vandenberghe, 2004, Lubin et al., 2016]. *Disciplined Convex Programming* (DCP) modeling packages such as CVX [Grant and Boyd, 2014], CVXPY [Diamond and Boyd, 2016], and Convex.jl [Udell et al., 2014] perform conic transformations automatically, conveniently enabling modelers to access powerful conic solvers such as ECOS [Domahidi et al., 2013], SCS [O’Donoghue et al., 2016], MOSEK [Mosek ApS, 2016], and CSDP [Borchers, 1999].

Conic solvers recognize the cone \mathcal{K} as a Cartesian product of standard *primitive* cones [Friberg, 2016]. A primitive closed convex cone cannot be written as a Cartesian product of two or more lower-dimensional closed convex cones. Lubin et al. [2017a] claim that the following classes of standard primitive cones are extremely versatile, encoding all of the problems in the Conic Benchmark Library (CBLIB) compiled by Friberg [2016], and all 333 MI-convex problems in MINLPLIB2 [Vigerske, 2018].

Linear cones naturally express affine constraints; any mixed-integer linear optimization (MILP) problem can be written in \mathfrak{M} form using nonnegative, nonpositive, and zero cones.

Second-order cones (and rotated-second-order cones) are widely used to model rational powers, norms, and geometric means, as well as convex quadratic objectives and constraints [Ben-Tal and Nemirovski, 2001a].

Positive semidefinite cones can model robust norms and functions of eigenvalues [Ben-Tal and Nemirovski, 2001a], and sum-of-squares constraints for polynomial optimization problems [Parrilo, 2003].

Exponential cones can model exponentials, logarithms, entropy, and powers, as well as log-sum-exp functions that arise from convex transformations of geometric programs [Serrano, 2015].

We note that conic representations are useful for constructing tight formulations for disjunctions or unions of convex sets [Lubin et al., 2017b,c, Vielma, 2018].

1.3 Overview And Contributions

In section 2, we start by reviewing the relevant foundations of conic duality and certificates. We then introduce the notion of \mathcal{K}^* cuts, and describe how to refine LP OAs of conic constraints using certificates obtained from continuous primal-dual conic solvers. For a MI-conic problem \mathfrak{M} , we propose the first B&B-OA algorithm based on conic certificates. We show that our algorithm detects infeasibility or unboundedness or terminates with an optimal solution in finite time under minimal assumptions.

In section 3, we demonstrate that a \mathcal{K}^* cut from a conic certificate implies useful guarantees about the infeasibility or optimal objective of an LP OA,

suggesting that our algorithm can often fathom a node immediately after solving the LP rather than proceeding to the expensive conic subproblem solve. We consider how these guarantees may be lost in the more realistic setting of an LP solver with a positive feasibility tolerance, and propose a practical methodology for scaling a certificate \mathcal{K}^* cut to recover similar guarantees.

In section 4, we describe how to strengthen the LP OAs by *disaggregating* \mathcal{K}^* cuts, and show that this methodology maintains the guarantees from section 3. We argue for initializing the LP OAs using *initial fixed* \mathcal{K}^* cuts, and offer a procedure for cheaply obtaining *separation* \mathcal{K}^* cuts to cut off an infeasible LP OA solution. These proposed techniques require minimal modifications to our algorithm and are practical to implement. In appendix A, we specialize these techniques for the second-order, positive semidefinite, and exponential cones.

In section 5, we describe the software architecture and algorithmic implementation of Pajarito, our open source MI-convex solver.³ This section may be of particular interest to advanced users and developers of mathematical optimization software. We emphasize that our implementations diverge from the idealized algorithmic description in section 2, because of our decision to leverage powerful external mixed-integer linear (MILP) solvers through limited, solver-independent interfaces. In appendix B, we describe how Pajarito lifts \mathcal{K}^* cuts for the second-order cone using an *extended formulation*, resulting in tighter LP OAs. In appendix C, we show how Pajarito can optionally tighten OAs for PSD cone constraints by strengthening \mathcal{K}^* cuts to rotated second-order cone constraints, which can be added to an MISOCP OA model.

In section 6, we summarize computational experiments demonstrating the speed and robustness of Pajarito. We benchmark Pajarito and several MISOCP solver packages, and conclude that Pajarito is the fastest and most-reliable open source solver for MISOCP. Finally, we compare the performance of several of Pajarito’s algorithmic variants on MI-conic instances involving mixtures of positive semidefinite, second-order, and exponential cones, demonstrating practical advantages of several methodological contributions from sections 3 and 4 and appendix A.

2 A Branch-And-Bound LP Outer Approximation Algorithm

For a MI-conic problem \mathfrak{M} , we propose a branch-and-bound LP outer approximation (B&B-OA) algorithm, the first such method based on conic certificates. In section 2.1, we describe the continuous conic subproblems that a nonlinear branch-and-bound (B&B-NL) algorithm would solve at each node, and review

³The new version of Pajarito that we implemented for this paper is the first conic-certificate-based OA solver. Although Pajarito solver was introduced in Lubin et al. [2016], this early implementation used NLP solvers instead of primal-dual conic solvers for continuous subproblems, and was built to assess the value of *extended formulations* by counting iterations before convergence. To avoid confusing users, we recently moved this old NLP-based functionality out of Pajarito and into *Pavito* solver at github.com/JuliaOpt/Pavito.jl.

the relevant foundations of conic duality. In section 2.2, we introduce the notion of \mathcal{K}^* cuts and describe how to refine polyhedral relaxations of the conic subproblems using information from conic certificates. Finally, we outline our B&B-OA algorithm in section 2.3, and discuss finiteness of convergence.

2.1 Continuous Subproblems And Conic Duality

Recall from \mathfrak{M} that the first I variables in \mathbf{x} are constrained to be integer. Branch-and-bound algorithms recursively partition the valid integer assignments, so for convenience we assume known finite lower bounds $\mathbf{l}^0 \in \mathbb{Z}^I$ and upper bounds $\mathbf{u}^0 \in \mathbb{Z}^I$ on the integer variables x_1, \dots, x_I . At a node of the branch-and-bound search tree with lower bounds $\mathbf{l} \in \mathbb{Z}^I$ and upper bounds $\mathbf{u} \in \mathbb{Z}^I$ on x_1, \dots, x_I , the natural continuous conic subproblem is $\mathfrak{C}(\mathbf{l}, \mathbf{u})$:

$$\mathfrak{C}(\mathbf{l}, \mathbf{u}) \left\{ \begin{array}{ll} \inf_{\mathbf{x}} & \mathbf{c}^T \mathbf{x} : & (4a) \\ & \mathbf{b} - \mathbf{A}\mathbf{x} \in \mathcal{K} & (4b) \\ & l_i - x_i \in \mathbb{R}_- & \forall i \in \llbracket I \rrbracket & (4c) \\ & u_i - x_i \in \mathbb{R}_+ & \forall i \in \llbracket I \rrbracket & (4d) \\ & \mathbf{x} \in \mathbb{R}^N, & (4e) \end{array} \right.$$

where the bound constraints (4c) and (4d) are expressed in conic form using the nonpositive cone \mathbb{R}_- (the nonpositive reals) and the nonnegative cone \mathbb{R}_+ (the nonnegative reals).

There exist primal-dual conic algorithms for $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ that are powerful in both theory and practice. The foundation for these methods and for much of this paper is the elegant theory of conic duality, described by Ben-Tal and Nemirovski [2001a], Boyd and Vandenberghe [2004]. Recall that the cone \mathcal{K} in $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ is a closed convex cone; we let \mathcal{K}^* denote the *dual cone* of \mathcal{K} , i.e. the set of points that have nonnegative inner product with all points in \mathcal{K} :

$$\mathcal{K}^* = \{\mathbf{z} \in \mathbb{R}^M : \mathbf{y}^T \mathbf{z} \geq 0, \forall \mathbf{y} \in \mathcal{K}\}. \quad (5)$$

\mathcal{K}^* is also a closed convex cone [Boyd and Vandenberghe, 2004]. The standard conic dual of $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ can be written as $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$:

$$\mathfrak{C}^*(\mathbf{l}, \mathbf{u}) \left\{ \begin{array}{ll} \sup_{\mathbf{z}, \boldsymbol{\mu}, \boldsymbol{\nu}} & -\mathbf{b}^T \mathbf{z} - \mathbf{l}^T \boldsymbol{\mu} - \mathbf{u}^T \boldsymbol{\nu} : & (6a) \\ & \mathbf{z} \in \mathcal{K}^* & (6b) \\ & \boldsymbol{\mu} \in \mathbb{R}_-^I & (6c) \\ & \boldsymbol{\nu} \in \mathbb{R}_+^I & (6d) \\ & \mathbf{c} + \mathbf{A}^T \mathbf{z} + \boldsymbol{\mu}' + \boldsymbol{\nu}' \in \{0\}^N, & (6e) \end{array} \right.$$

where for ease of exposition we let $\boldsymbol{\mu}' = (\mu_1, \dots, \mu_I, 0, \dots, 0) \in \mathbb{R}^N$ and similarly for $\boldsymbol{\nu}'$. Note that the nonnegative and nonpositive cones are both *self-dual*, i.e.

$\mathbb{R}_-^* = \mathbb{R}_-$ and $\mathbb{R}_+^* = \mathbb{R}_+$. The zero cone $\{0\}$ (containing only the origin) is dual to the free cone \mathbb{R} . The variables \mathbf{z} in the dual constraint (6b) are associated with the primal constraint (4b), and similarly for (6c) and (4c), (6d) and (4d), and (4e) and (6e).

If the conic primal-dual pair $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ – $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$ is *well-posed*, then conic duality can be thought of as a simple generalization of LP duality.⁴ In particular, the inf and sup can be replaced with min and max, and the possible status combinations for $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ and $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$ are: both infeasible, one unbounded and the other infeasible, or both feasible and bounded with equal objective values attained by optimal solutions. The conditions for well-posedness in conic duality are described by Friberg [2016], and are outside the scope of this paper, so we assume that any primal-dual subproblem we encounter has the well-posed property.

Friberg [2016] discusses certificates that provide easily-verifiable proofs of unboundedness or infeasibility of the primal or dual problems or of optimality of a given pair of primal and dual points. In terms of the primal subproblem $\mathfrak{C}(\mathbf{l}, \mathbf{u})$, the three possible mutually-exclusive cases and their interpretations are as follows.

A dual improving ray certifies that $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ is infeasible, via the conic generalization of Farkas’ lemma. The improving ray $(\bar{\mathbf{z}}, \bar{\boldsymbol{\mu}}, \bar{\boldsymbol{\nu}}) \in \mathbb{R}^{M+2I}$ of $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$ is a feasible direction for $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$ along which the objective value of any feasible point of $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$ can be improved indefinitely. It satisfies the following conditions:

$$-\mathbf{b}^T \bar{\mathbf{z}} - \mathbf{l}^T \bar{\boldsymbol{\mu}} - \mathbf{u}^T \bar{\boldsymbol{\nu}} > 0 \quad (7a)$$

$$\bar{\mathbf{z}} \in \mathcal{K}^* \quad (7b)$$

$$\bar{\boldsymbol{\mu}} \leq \mathbf{0} \quad (7c)$$

$$\bar{\boldsymbol{\nu}} \geq \mathbf{0} \quad (7d)$$

$$\mathbf{A}^T \bar{\mathbf{z}} + \bar{\boldsymbol{\mu}}' + \bar{\boldsymbol{\nu}}' = \mathbf{0}. \quad (7e)$$

Clearly, if $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$ itself has a feasible point, then it is unbounded, otherwise it is infeasible.⁵

A primal improving ray and a primal feasible point certifies that $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ is unbounded, because the improving ray is a feasible direction along which the objective value of the feasible point can be improved indefinitely. The improving ray $\bar{\mathbf{x}} \in \mathbb{R}^N$ of $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ also implies infeasibility of $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$ and satisfies the following conditions:

$$\mathbf{c}^T \bar{\mathbf{x}} < 0 \quad (8a)$$

$$-\mathbf{A} \bar{\mathbf{x}} \in \mathcal{K} \quad (8b)$$

$$\bar{x}_i = 0 \quad \forall i \in \llbracket I \rrbracket, \quad (8c)$$

⁴If \mathcal{K} is polyhedral, then \mathcal{K}^* is polyhedral, and hence $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ and $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$ are both LPs. All LPs are well-posed.

⁵Conditions (7b) to (7e) imply that $(\bar{\mathbf{z}}, \bar{\boldsymbol{\mu}}, \bar{\boldsymbol{\nu}})$ is feasible for a modified $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$ problem in which $\mathbf{c} = \mathbf{0}$.

and the feasible point $\hat{\mathbf{x}} \in \mathbb{R}^N$ of $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ simply satisfies the primal feasibility conditions (4b) to (4e). Note that if the objective coefficients of the continuous variables are all zero ($c_{I+1} = \dots = c_N = 0$), then conditions (8a) and (8c) can never be satisfied, so there cannot be a primal improving ray.⁶

A complementary solution pair certifies optimality for $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ of the primal feasible point $\hat{\mathbf{x}} \in \mathbb{R}^N$ in the pair $(\hat{\mathbf{x}}, (\hat{\mathbf{z}}, \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\nu}}))$, via conic weak duality. The dual feasible point $(\hat{\mathbf{z}}, \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\nu}}) \in \mathbb{R}^{M+2I}$ is also optimal for $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$, and the pair have equal primal and dual objective values:

$$\mathbf{c}^T \hat{\mathbf{x}} = -\mathbf{b}^T \hat{\mathbf{z}} - \mathbf{l}^T \hat{\boldsymbol{\mu}} - \mathbf{u}^T \hat{\boldsymbol{\nu}}. \quad (9)$$

2.2 Dynamic Polyhedral Relaxations

Recall from equation (5) that $\mathbf{y} \in \mathcal{K}$ if and only if $\mathbf{z}^T \mathbf{y} \geq 0, \forall \mathbf{z} \in \mathcal{K}^*$. This implies that a nonpolyhedral conic constraint $\mathbf{b} - \mathbf{A}\mathbf{x} \in \mathcal{K}$ has the following equivalent semi-infinite linear representation:

$$\mathbf{z}^T (\mathbf{b} - \mathbf{A}\mathbf{x}) \geq 0 \quad \forall \mathbf{z} \in \mathcal{K}^*. \quad (10)$$

We refer to a point $\mathbf{z} \in \mathcal{K}^*$ as a \mathcal{K}^* *point*, and call the corresponding linear constraint $\mathbf{z}^T (\mathbf{b} - \mathbf{A}\mathbf{x}) \geq 0$ a \mathcal{K}^* *cut*. A \mathcal{K}^* cut cannot exclude any point \mathbf{x} that satisfies $\mathbf{b} - \mathbf{A}\mathbf{x} \in \mathcal{K}$, so any finite set of \mathcal{K}^* cuts defines a valid polyhedral relaxation of the conic constraint (4b).

Given a finite set $\mathcal{Z} \subset \mathcal{K}^*$ of \mathcal{K}^* points, consider modifying the subproblem $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ by relaxing the conic constraint and instead imposing the finite number of \mathcal{K}^* cuts implied by \mathcal{Z} . We refer to the resulting LP as $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$, which we choose to write in inequality form rather than conic form:

$$\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u}) \left\{ \begin{array}{ll} \min_{\mathbf{x}} & \mathbf{c}^T \mathbf{x} : \end{array} \right. \quad (11a)$$

$$x_i \geq l_i \quad \forall i \in \llbracket I \rrbracket \quad (11b)$$

$$x_i \leq u_i \quad \forall i \in \llbracket I \rrbracket \quad (11c)$$

$$\mathbf{z}^T (\mathbf{b} - \mathbf{A}\mathbf{x}) \geq 0 \quad \forall \mathbf{z} \in \mathcal{Z}. \quad (11d)$$

Since $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ and $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ have the same objective function, and the feasible set of $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ is a polyhedral relaxation of the feasible set of $\mathfrak{C}(\mathbf{l}, \mathbf{u})$, solving $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ with an LP solver may give us useful information about $\mathfrak{C}(\mathbf{l}, \mathbf{u})$. If $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ is infeasible, then $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ must be infeasible. If $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ has an optimal objective value of L , then $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ is either infeasible or has an optimal objective no smaller than L . In these cases, we may be able to immediately fathom the node by infeasibility or by bound, or even use a fractional optimal solution for $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ to make a branching decision, without needing to

⁶This matches intuition because if the continuous variables have zero objective coefficients and the integer variables are bounded, \mathfrak{M} cannot be unbounded.

solve $\mathfrak{C}(\mathbf{l}, \mathbf{u})$. However, if $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ is unbounded, it does not provide useful information about the status or optimal value of $\mathfrak{C}(\mathbf{l}, \mathbf{u})$.

An LP solver based on the Simplex algorithm is able to rapidly reoptimize $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ at each new node, after the integer variable bounds are updated and any new \mathcal{K}^* cuts are added. As noted by Skajaa et al. [2013], state-of-the-art conic solvers benefit much less from warm-starting, so it may be computationally faster to sacrifice some information about the conic subproblems in order to avoid some expensive conic subproblem solves.

In analogy to B&B-OA algorithms based on gradient cuts, we add a new cut after every infeasible or bounded conic subproblem solve. Our key innovation, however, is to obtain this cut directly from the conic certificate found by the conic subproblem solver. Suppose that at some node, a primal-dual conic subproblem solver yields a dual improving ray $(\bar{\mathbf{z}}, \bar{\boldsymbol{\mu}}, \bar{\boldsymbol{\nu}})$: from condition (7b), $\bar{\mathbf{z}} \in \mathcal{K}^*$, so $\bar{\mathbf{z}}$ is a \mathcal{K}^* point. Now suppose that the subproblem solver yields a complementary solution $(\hat{\mathbf{x}}, (\hat{\mathbf{z}}, \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\nu}}))$: by the dual feasibility condition (6b), $\hat{\mathbf{z}} \in \mathcal{K}^*$, so $\hat{\mathbf{z}}$ is a \mathcal{K}^* point. In both cases, a subvector of the ray or solution for the dual subproblem $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$ allows us to augment $\mathcal{Z} \subset \mathcal{K}^*$, refining our LP OA model $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$. In section 3, we use conic duality theory to show that these \mathcal{K}^* cuts derived from certificates encode important information about conic subproblems into the subsequent polyhedral relaxations.

2.3 The Conic-Certificate-Based Algorithm

Our conic-certificate-based B&B-OA algorithm for the MI-conic problem \mathfrak{M} is outlined in algorithm 1. Recall that \mathfrak{M} is in minimization form. Algorithm 1 maintains an upper bound U (initially ∞), a corresponding best feasible solution set \mathcal{X} (initially empty), and a set of active nodes \mathcal{N} of the search tree. A node $(\mathbf{l}, \mathbf{u}, L)$ is characterized by the finite variable bound vectors \mathbf{l} and \mathbf{u} and a lower bound value L . The node's lower bound L signifies that all feasible solutions for \mathfrak{M} that satisfy the node's bounds on integer variables have an objective value of at least L . The node set \mathcal{N} is initialized to contain only the root node $(\mathbf{l}^0, \mathbf{u}^0, -\infty)$, where $\mathbf{l}^0, \mathbf{u}^0 \in \mathbb{R}^I$ are the finite initial global bounds on the integer variables.

On line 5, the main loop removes a node $(\mathbf{l}, \mathbf{u}, L)$ from \mathcal{N} . If the node's lower bound L is no smaller than the current global best upper bound U , line 7 fathoms the node by bound as it cannot yield a better incumbent. Otherwise, line 8 solves the node's LP OA model $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$, taking advantage of an LP warm-start from a previous node.

If $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ is infeasible, line 10 immediately fathoms the node by infeasibility. If $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ has an optimal solution $\hat{\mathbf{x}}$, then its optimal objective value is the tightest lower bound known for $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ (in section 3.1.2, we prove $\mathbf{c}^T \hat{\mathbf{x}} \geq L$ is a consequence of the \mathcal{K}^* cuts), so line 12 updates L to $\mathbf{c}^T \hat{\mathbf{x}}$. Line 14 fathoms the node by bound if L is no better than the incumbent value U , otherwise if $\hat{\mathbf{x}}$ is fractional (i.e. it violates an integrality constraint (2c)), line 17 branches on it.⁷ The

⁷We could instead remove lines 15-17 and solve the conic subproblem even if the LP solution

Algorithm 1: Conic-certificate-based branch-and-bound LP outer approximation for \mathfrak{M} .

```

1 initialize incumbent solution set  $\mathcal{X}$  to  $\emptyset$ , upper bound  $U$  to  $\infty$ 
2 initialize  $\mathcal{K}^*$  point set  $\mathcal{Z}$  to  $\emptyset$ 
3 initialize node list  $\mathcal{N}$  with root node  $(\mathbf{l}^0, \mathbf{u}^0, -\infty)$ 
4 while  $\mathcal{N}$  contains nodes do
5     remove a node  $(\mathbf{l}, \mathbf{u}, L)$  from  $\mathcal{N}$ 
6     if lower bound  $L \geq U$  then
7         continue ▷ fathomed by bound
8     call LP solver on  $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ 
9     if get an infeasibility proof then
10         continue ▷ fathomed by infeasibility
11     else if get an optimal solution  $\hat{\mathbf{x}}$  then
12         update  $L$  to  $\mathbf{c}^T \hat{\mathbf{x}}$ 
13         if  $L \geq U$  then
14             continue ▷ fathomed by bound
15         else if  $\hat{\mathbf{x}}$  is fractional then
16             add branch nodes to  $\mathcal{N}$  using  $\hat{\mathbf{x}}$  and  $L$ 
17             continue ▷ branched
18     call primal-dual continuous conic solver on  $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ 
19     if get a dual improving ray  $(\bar{\mathbf{z}}, \bar{\boldsymbol{\mu}}, \bar{\nu})$  then
20         add  $\mathcal{K}^*$  point  $\bar{\mathbf{z}}$  to  $\mathcal{Z}$ 
21         continue ▷ fathomed by infeasibility
22     else if get a primal improving ray  $\bar{\mathbf{x}}$  and feasible point  $\hat{\mathbf{x}}$  then
23         if  $\hat{\mathbf{x}}$  is integral then
24             update  $U$  to  $-\infty$ 
25             break ▷ proven unbounded
26     else if get a complementary solution  $(\hat{\mathbf{x}}, (\hat{\mathbf{z}}, \hat{\boldsymbol{\mu}}, \hat{\nu}))$  then
27         add  $\mathcal{K}^*$  point  $\hat{\mathbf{z}}$  to  $\mathcal{Z}$ 
28         update  $L$  to  $\mathbf{c}^T \hat{\mathbf{x}}$ 
29         if  $L \geq U$  then
30             continue ▷ fathomed by bound
31         else if  $\hat{\mathbf{x}}$  is integral then
32             update  $\mathcal{X}$  to  $\{\hat{\mathbf{x}}\}$  and  $U$  to  $\mathbf{c}^T \hat{\mathbf{x}}$ 
33             continue ▷ fathomed by integrality
34     add branch nodes to  $\mathcal{N}$  using  $\hat{\mathbf{x}}$  (fractional) and  $L$ 
35 return  $\mathcal{X}, U$ 

```

branch procedure strictly partitions the node's integer bounds \mathbf{l} and \mathbf{u} by picking an $i \in \llbracket I \rrbracket : \hat{x}_i \notin \mathbb{Z}$ and adding two child nodes to \mathcal{N} : $(\mathbf{l}, (u_1, \dots, \lfloor \hat{x}_i \rfloor, \dots, u_N), L)$ and $((l_1, \dots, \lceil \hat{x}_i \rceil, \dots, l_N), \mathbf{u}, L)$.

If the node is not fathomed or branched on immediately after the LP solve (before line 18), then $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ is either unbounded or has an optimal solution $\hat{\mathbf{x}}$ that is integral (i.e. $\hat{x}_i \in \mathbb{Z}, \forall i \in \llbracket I \rrbracket$) with optimal value $\mathbf{c}^T \hat{\mathbf{x}} < U$. Then line 18 solves the conic subproblem $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ with the primal-dual continuous conic solver. Recall from section 2.1 our assumption that the primal-dual subproblem pair $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ – $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$ is well-posed, so the conic solver returns one of the three possible certificates, which we handle as follows.

A dual improving ray on line 19 provides a \mathcal{K}^* point, which line 20 adds to \mathcal{Z} (as described in section 2.2). This certificate proves that $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ is infeasible, so line 21 fathoms the node by infeasibility.

A primal improving ray and feasible point on line 22 certifies that $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ is unbounded. Since the primal improving ray conditions (8a) to (8c) are the same for any conic subproblem, every subproblem is infeasible or unbounded, so \mathfrak{M} is either infeasible or unbounded. The incumbent solution set must be empty and $U = \infty$. Line 23 checks whether the feasible point $\hat{\mathbf{x}}$ is integral. If so, it is a feasible solution for \mathfrak{M} , so \mathfrak{M} is unbounded and line 25 terminates the main loop.

A complementary solution on line 26 provides a \mathcal{K}^* point that line 27 adds to \mathcal{Z} (as described in section 2.2) and an optimal solution $\hat{\mathbf{x}}$ for $\mathfrak{C}(\mathbf{l}, \mathbf{u})$. The optimal objective value gives the tightest lower bound known for the node, so line 28 updates L to $\mathbf{c}^T \hat{\mathbf{x}}$, and line 30 fathoms by bound if this value is no better than U . Line 31 checks if $\hat{\mathbf{x}}$ is integral, in which case it becomes the new incumbent solution for \mathfrak{M} on line 32, and the node is fathomed by integrality on line 33.

If the node is not fathomed immediately after the conic solve (before line 34), then $\hat{\mathbf{x}}$ is a feasible solution for $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ that is fractional. L is either ∞ (in the primal improving ray case) or finite (in the complementary solution case), and is the best known lower bound for the node. Line 34 branches on $\hat{\mathbf{x}}$ using the same branch procedure we describe above for line 17.

Since the initial bounds on the integer variables are finite, and the main loop of algorithm 1 either fathoms each node or strictly partitions its integer bounds or terminates the algorithm, it follows that algorithm 1 terminates finitely. From the fact that $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ is a valid polyhedral relaxation of $\mathfrak{C}(\mathbf{l}, \mathbf{u})$, and from the correctness of our inferences from the subproblem certificates, it is clear algorithm 1 terminates correctly, under the assumption of well-posed conic subproblems. On line 35, if $U = \infty$, then \mathfrak{M} is proven infeasible, otherwise if U is finite, then \mathcal{X} contains an optimal solution for \mathfrak{M} , otherwise $U = -\infty$ and \mathfrak{M} is proven unbounded.

is fractional, rather than branching. This variation may perform better if the conic subproblem solves are quite fast in practice.

We note that without using the LP $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ (i.e. removing lines 8 to 17 and not creating and augmenting \mathcal{Z} on lines 2, 20 and 27), we get a simple conic-certificate-based B&B-NL algorithm for $\mathfrak{C}(\mathbf{l}, \mathbf{u})$, for which finite termination guarantees and correctness follow from the same assumptions and arguments. We have omitted any discussion of node selection or fractional variable selection for branching. MILP solvers can use LP certificates to make intelligent selections, and we expect that some of these LP-based criteria are generalizeable to the conic case, as conic duality theory is a simple extension of LP duality under the well-posed assumption.

3 Polyhedral Relaxation Guarantees From Conic Certificates

Recall from section 2.2 that \mathcal{K}^* cuts yield *valid* polyhedral relaxations of the conic constraint $\mathbf{b} - \mathbf{A}\mathbf{x} \in \mathcal{K}$, and a *certificate* \mathcal{K}^* cut can be obtained directly from the conic certificate for an infeasible or bounded and feasible subproblem $\mathfrak{C}(\mathbf{l}, \mathbf{u})$. We demonstrate in section 3.1 that a certificate \mathcal{K}^* cut implies useful *guarantees* about the infeasibility or optimal objective of the LP OAs, suggesting that algorithm 1 can often fathom a node immediately after solving the LP OA rather than proceeding to the expensive conic subproblem solve.⁸ In section 3.2, we consider how these guarantees may be lost in the more realistic setting of an LP solver with a positive feasibility tolerance, and propose a practical methodology for *scaling* a certificate \mathcal{K}^* cut to recover similar guarantees.

3.1 Under An Exact LP Solver

We continue to assume well-posedness of every conic subproblem at every node. We consider what a certificate \mathcal{K}^* from the conic subproblem $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ at a node with integer variable bounds \mathbf{l}, \mathbf{u} implies about the LP OA $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ at a different node with bounds \mathbf{l}, \mathbf{u} .

3.1.1 Certificate Cuts From Dual Improving Rays

Suppose $(\bar{\mathbf{z}}, \bar{\boldsymbol{\mu}}, \bar{\boldsymbol{\nu}})$ is an improving ray of the dual subproblem $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$, certifying infeasibility of $\mathfrak{C}(\mathbf{l}, \mathbf{u})$. Using properties (7a) to (7e) of this certificate, any point $\mathbf{x} \in \mathbb{R}^N$ satisfying the integer variable bounds $\bar{l}_i \leq x_i \leq \bar{u}_i, \forall i \in \llbracket I \rrbracket$ at the new

⁸By similar arguments, we expect that the certificate \mathcal{K}^* cut may be useful at nearby nodes for duality based preprocessing such as reduced cost fixing [Gally et al., 2018, sec. 7] or conflict analysis [Witzig et al., 2017].

node also satisfies:

$$\bar{\mathbf{z}}^T(\mathbf{b} - \mathbf{A}\mathbf{x}) = \mathbf{b}^T \bar{\mathbf{z}} - \mathbf{x}^T \mathbf{A}^T \bar{\mathbf{z}} \quad (12a)$$

$$= \mathbf{b}^T \bar{\mathbf{z}} + \mathbf{x}^T \bar{\boldsymbol{\mu}}' + \mathbf{x}^T \bar{\boldsymbol{\nu}}' \quad (12b)$$

$$\leq \mathbf{b}^T \bar{\mathbf{z}} + \mathbf{x}^T \bar{\boldsymbol{\mu}}' + \mathbf{x}^T \bar{\boldsymbol{\nu}}' + \sum_{i \in \llbracket I \rrbracket} ((\underline{l}_i - x_i) \bar{\mu}_i + (\underline{u}_i - x_i) \bar{\nu}_i) \quad (12c)$$

$$= \mathbf{b}^T \bar{\mathbf{z}} + \underline{\mathbf{l}}^T \bar{\boldsymbol{\mu}} + \underline{\mathbf{u}}^T \bar{\boldsymbol{\nu}} \quad (12d)$$

$$= (\mathbf{b}^T \bar{\mathbf{z}} + \mathbf{l}^T \bar{\boldsymbol{\mu}} + \mathbf{u}^T \bar{\boldsymbol{\nu}}) - (\mathbf{l} - \underline{\mathbf{l}})^T \bar{\boldsymbol{\mu}} - (\mathbf{u} - \underline{\mathbf{u}})^T \bar{\boldsymbol{\nu}}. \quad (12e)$$

From property (7a) of the certificate, $\mathbf{b}^T \bar{\mathbf{z}} + \mathbf{l}^T \bar{\boldsymbol{\mu}} + \mathbf{u}^T \bar{\boldsymbol{\nu}} < 0$. If $l_i \leq \underline{l}_i \leq u_i \leq \bar{u}_i, \forall i \in \llbracket I \rrbracket$, then $(\mathbf{l} - \underline{\mathbf{l}})^T \bar{\boldsymbol{\mu}} \geq 0$ and $(\mathbf{u} - \underline{\mathbf{u}})^T \bar{\boldsymbol{\nu}} \geq 0$. In this case, the value (12e) is negative, so from (12a) to (12e), the certificate \mathcal{K}^* cut $\bar{\mathbf{z}}^T(\mathbf{b} - \mathbf{A}\mathbf{x}) \geq 0$ is violated. Therefore, the certificate \mathcal{K}^* cut from the infeasible subproblem $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ guarantees infeasibility of any LP OA $\mathfrak{P}(\mathcal{Z}, \underline{\mathbf{l}}, \underline{\mathbf{u}})$ in the subtree of the node with integer variable bounds $\underline{\mathbf{l}}, \underline{\mathbf{u}}$.

More importantly for algorithm 1, the certificate \mathcal{K}^* cut is likely to remain violated at ‘nearby’ nodes outside of this subtree, as the conditions (12a) to (12e) have a natural interpretation from global sensitivity analysis. Perturbing the integer variable bounds from $\underline{\mathbf{l}}, \underline{\mathbf{u}}$ to \mathbf{l}, \mathbf{u} changes the upper bound on $\bar{\mathbf{z}}^T(\mathbf{b} - \mathbf{A}\mathbf{x})$ through a linear dependence on the values $\boldsymbol{\mu} \leq 0$ and $\boldsymbol{\nu} \geq 0$ of the dual variables in the improving ray of $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$.

3.1.2 Certificate Cuts From Dual Optimal Solutions

Suppose $(\hat{\mathbf{x}}, (\hat{\mathbf{z}}, \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\nu}}))$ is a complementary solution pair for the subproblem $\mathfrak{C}(\mathbf{l}, \mathbf{u})$, certifying optimality of the solution pair. Using the strong duality conditions (property (9) and feasibility for $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ and $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$), any point $\mathbf{x} \in \mathbb{R}^N$ satisfying the integer variable bounds $\underline{l}_i \leq x_i \leq \underline{u}_i, \forall i \in \llbracket I \rrbracket$ at the new node and the certificate \mathcal{K}^* cut $\hat{\mathbf{z}}^T(\mathbf{b} - \mathbf{A}\mathbf{x}) \geq 0$ has objective value:

$$\mathbf{c}^T \mathbf{x} = -(\mathbf{A}^T \hat{\mathbf{z}} + \hat{\boldsymbol{\mu}}' + \hat{\boldsymbol{\nu}}')^T \mathbf{x} \quad (13a)$$

$$= -\hat{\mathbf{z}}^T \mathbf{A}\mathbf{x} - \mathbf{x}^T (\hat{\boldsymbol{\mu}}' + \hat{\boldsymbol{\nu}}') \quad (13b)$$

$$= -\mathbf{b}^T \hat{\mathbf{z}} + \hat{\mathbf{z}}^T (\mathbf{b} - \mathbf{A}\mathbf{x})^T - \mathbf{x}^T (\hat{\boldsymbol{\mu}}' + \hat{\boldsymbol{\nu}}') \quad (13c)$$

$$\geq -\mathbf{b}^T \hat{\mathbf{z}} - \mathbf{x}^T (\hat{\boldsymbol{\mu}}' + \hat{\boldsymbol{\nu}}') \quad (13d)$$

$$\geq -\mathbf{b}^T \hat{\mathbf{z}} - \mathbf{x}^T (\hat{\boldsymbol{\mu}}' + \hat{\boldsymbol{\nu}}') - \sum_{i \in \llbracket I \rrbracket} ((\underline{l}_i - x_i) \hat{\mu}_i + (\underline{u}_i - x_i) \hat{\nu}_i) \quad (13e)$$

$$= -\mathbf{b}^T \hat{\mathbf{z}} - \underline{\mathbf{l}}^T \hat{\boldsymbol{\mu}} - \underline{\mathbf{u}}^T \hat{\boldsymbol{\nu}} \quad (13f)$$

$$= (-\mathbf{b}^T \hat{\mathbf{z}} - \mathbf{l}^T \hat{\boldsymbol{\mu}} - \mathbf{u}^T \hat{\boldsymbol{\nu}}) + (\mathbf{l} - \underline{\mathbf{l}})^T \hat{\boldsymbol{\mu}} + (\mathbf{u} - \underline{\mathbf{u}})^T \hat{\boldsymbol{\nu}} \quad (13g)$$

$$= \mathbf{c}^T \hat{\mathbf{x}} + (\mathbf{l} - \underline{\mathbf{l}})^T \hat{\boldsymbol{\mu}} + (\mathbf{u} - \underline{\mathbf{u}})^T \hat{\boldsymbol{\nu}}. \quad (13h)$$

If $l_i \leq \underline{l}_i \leq \underline{u}_i \leq u_i, \forall i \in \llbracket I \rrbracket$, then $(\mathbf{l} - \underline{\mathbf{l}})^T \hat{\boldsymbol{\mu}} \geq 0$ and $(\mathbf{u} - \underline{\mathbf{u}})^T \hat{\boldsymbol{\nu}} \geq 0$. In this case, the value (13h) is no smaller than $\mathbf{c}^T \hat{\mathbf{x}}$, the lower bound from the subproblem $\mathfrak{C}(\mathbf{l}, \mathbf{u})$. Therefore, the certificate \mathcal{K}^* cut from the feasible

subproblem $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ guarantees that the optimal value of any LP OA $\mathfrak{P}(\mathcal{Z}, \mathbf{l}, \mathbf{u})$ in the subtree of the node with integer variable bounds \mathbf{l}, \mathbf{u} does not decrease, but may actually improve.⁹

More importantly for algorithm 1, at ‘nearby’ nodes outside of this subtree, the objective bounds implied by the certificate \mathcal{K}^* cut in the LP OA model are likely to remain fairly tight. Perturbing the integer variable bounds from \mathbf{l}, \mathbf{u} to $\bar{\mathbf{l}}, \bar{\mathbf{u}}$ changes the lower bound on $\mathbf{c}^T \mathbf{x}$ through a linear dependence on the values $\boldsymbol{\mu} \leq 0$ and $\boldsymbol{\nu} \geq 0$ of the dual variables in the complementary solution pair for $\mathfrak{C}(\mathbf{l}, \mathbf{u})$.

3.2 Under An LP Solver With A Feasibility Tolerance

So far, we have been assuming that the LP solver computes a solution that satisfies all the \mathcal{K}^* cuts in the LP OAs exactly. In practice, LP solvers based on the Simplex method (except those that use rational arithmetic) enforce constraints up to an absolute constraint-wise violation tolerance $\delta > 0$ (typically set by the user). Therefore, a more realistic assumption is that any solution returned by the LP solver does not violate any \mathcal{K}^* cut by more than δ , i.e. a \mathcal{K}^* point \mathbf{z} effectively yields a ‘relaxed \mathcal{K}^* cut’ $\mathbf{z}^T(\mathbf{b} - \mathbf{A}\mathbf{x}) \geq -\delta$. Under this relaxed condition, we may lose the ‘within-subtree’ guarantees described in section 3.1. However, noting that any positive scaling of a \mathcal{K}^* point is still a \mathcal{K}^* point, we demonstrate how to recover the infeasibility guarantee from section 3.1.1 exactly, and the objective bound guarantee from section 3.1.2 to within a given relative objective gap tolerance. Such an analysis appears to be novel in the MI-convex literature.

3.2.1 Certificate Cuts From Dual Improving Rays

Suppose $(\bar{\mathbf{z}}, \bar{\boldsymbol{\mu}}, \bar{\boldsymbol{\nu}})$ is an improving ray of the dual subproblem $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$. From the property (7a) of the certificate and the conditions (12a) to (12e), any point $\mathbf{x} \in \mathbb{R}^N$ satisfying the integer variable bounds $l_i \leq x_i \leq u_i, \forall i \in \llbracket I \rrbracket$ and the relaxed certificate \mathcal{K}^* cut condition $\bar{\mathbf{z}}^T(\mathbf{b} - \mathbf{A}\mathbf{x}) \geq -\delta$ must satisfy:

$$0 > \mathbf{b}^T \bar{\mathbf{z}} + \mathbf{l}^T \bar{\boldsymbol{\mu}} + \mathbf{u}^T \bar{\boldsymbol{\nu}} \geq \bar{\mathbf{z}}^T(\mathbf{b} - \mathbf{A}\mathbf{x}) \geq -\delta. \quad (14)$$

Therefore, if $\delta > 0$ is sufficiently large, the relaxed certificate \mathcal{K}^* cut condition fails to enforce the infeasibility guarantee from section 3.1.1.

However, for a positive multiplier $\bar{\gamma} > 0$ satisfying:

$$\bar{\gamma} > \frac{\delta}{-\mathbf{b}^T \bar{\mathbf{z}} - \mathbf{l}^T \bar{\boldsymbol{\mu}} - \mathbf{u}^T \bar{\boldsymbol{\nu}}} > 0, \quad (15)$$

we have $\bar{\gamma}(\mathbf{b}^T \bar{\mathbf{z}} + \mathbf{l}^T \bar{\boldsymbol{\mu}} + \mathbf{u}^T \bar{\boldsymbol{\nu}}) < -\delta$. Therefore, the relaxed *scaled* certificate \mathcal{K}^* cut condition $\bar{\gamma} \bar{\mathbf{z}}^T(\mathbf{b} - \mathbf{A}\mathbf{x}) \geq -\delta$ recovers the infeasibility guarantee within

⁹If algorithm 1 branches on line 34 after solving a bounded and feasible conic subproblem to get the tightest lower bound, then when examining a child node, this objective guarantee ensures the node’s lower bound L does not decrease when we update it to the optimal value of the LP OA on line 12.

the subtree of the node from which the certificate is obtained. Note that the scaling factor (15) depends only on δ , problem data, and the certificate for the infeasible subproblem $\mathfrak{C}(\mathbf{l}, \mathbf{u})$. We can modify algorithm 1 on line 20 to add the scaled \mathcal{K}^* point $\bar{\gamma}\hat{\mathbf{z}}$ to \mathcal{Z} .

3.2.2 Certificate Cuts From Dual Optimal Solutions

Suppose $(\hat{\mathbf{x}}, (\hat{\mathbf{z}}, \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\nu}}))$ is a complementary solution pair for the subproblem $\mathfrak{C}(\mathbf{l}, \mathbf{u})$. From the conditions (13a) to (13h), any point $\mathbf{x} \in \mathbb{R}^N$ satisfying the integer variable bounds $l_i \leq x_i \leq u_i, \forall i \in \llbracket I \rrbracket$ and the relaxed certificate \mathcal{K}^* cut condition $\hat{\mathbf{z}}^T(\mathbf{b} - \mathbf{A}\mathbf{x}) \geq -\delta$ has objective value:

$$\mathbf{c}^T \mathbf{x} \geq -\mathbf{b}^T \hat{\mathbf{z}} + \hat{\mathbf{z}}^T(\mathbf{b} - \mathbf{A}\mathbf{x}) - \mathbf{l}^T \hat{\boldsymbol{\mu}} - \mathbf{u}^T \hat{\boldsymbol{\nu}} \geq L - \delta. \quad (16)$$

Recall $L = \mathbf{c}^T \hat{\mathbf{x}} = -\mathbf{b}^T \hat{\mathbf{z}} - \mathbf{l}^T \hat{\boldsymbol{\mu}} - \mathbf{u}^T \hat{\boldsymbol{\nu}}$ is the optimal objective value of $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ and $\mathfrak{C}^*(\mathbf{l}, \mathbf{u})$. Therefore, the relaxed certificate \mathcal{K}^* cut condition only enforces the objective guarantee from section 3.1.2 to an absolute tolerance of δ . In general, it makes little sense for an objective guarantee to depend on the LP solver's feasibility tolerance.

Instead, for a relative optimality gap tolerance $\epsilon > 0$, we can easily motivate a relative objective gap condition such as:

$$\frac{L - \mathbf{c}^T \mathbf{x}}{|L| + \theta} \leq \epsilon. \quad (17)$$

Consider a positive multiplier $\hat{\gamma} > 0$ satisfying:

$$\hat{\gamma} \geq \frac{\delta}{\epsilon(|L| + \theta)} > 0. \quad (18)$$

Modifying the conditions (16) for the relaxed *scaled* certificate \mathcal{K}^* cut condition $\hat{\gamma}\hat{\mathbf{z}}^T(\mathbf{b} - \mathbf{A}\mathbf{x}) \geq -\delta$, we get $\mathbf{c}^T \mathbf{x} \geq L - \delta/\hat{\gamma}$. Rearranging, this implies:

$$\frac{L - \mathbf{c}^T \mathbf{x}}{|L| + \theta} \leq \frac{\delta}{\hat{\gamma}(|L| + \theta)} \leq \epsilon, \quad (19)$$

so by scaling the certificate \mathcal{K}^* cut by $\hat{\gamma}$, we achieve the relative objective gap guarantee (17) within the subtree of the node from which the certificate is obtained. Note that the scaling factor (18) depends only on ϵ , δ , problem data, and the certificate for the bounded and feasible subproblem $\mathfrak{C}(\mathbf{l}, \mathbf{u})$. We can modify algorithm 1 on line 27 to add the scaled \mathcal{K}^* point $\hat{\gamma}\hat{\mathbf{z}}$ to \mathcal{Z} .

4 Tightening Polyhedral Relaxations

In section 4.1, we outline a two-stage procedure for *disaggregating* \mathcal{K}^* cuts to get stronger polyhedral relaxations, and show how to maintain the certificate \mathcal{K}^* cut guarantees from section 3. In section 4.2, we argue for initializing the

polyhedral relaxations using *initial fixed* \mathcal{K}^* cuts, and in section 4.3, we describe a procedure for cheaply obtaining *separation* \mathcal{K}^* cuts to cut off an infeasible LP OA solution. All of our proposed techniques for tightening the LP OAs require minimal modifications to algorithm 1 and are practical to implement.

4.1 Extreme Ray Disaggregation

Consider a set of \mathcal{K}^* points $\mathcal{Z} = \{\mathbf{z}^1, \dots, \mathbf{z}^J\} \subset \mathcal{K}^*$. By *aggregating* the corresponding \mathcal{K}^* cuts, we see they imply infinitely many \mathcal{K}^* cuts:

$$\mathbf{z}^T(\mathbf{b} - \mathbf{A}\mathbf{x}) \geq 0 \quad \forall \mathbf{z} \in \text{cone}(\mathcal{Z}), \quad (20)$$

where $\text{cone}(\mathcal{Z})$ is the *conic hull* of \mathcal{Z} , i.e. the set of conic (nonnegative) combinations of $\mathbf{z}^1, \dots, \mathbf{z}^J$:

$$\text{cone}(\mathcal{Z}) = \{\alpha^1 \mathbf{z}^1 + \dots + \alpha^J \mathbf{z}^J : \alpha^1, \dots, \alpha^J \geq 0\} \subset \mathcal{K}^*. \quad (21)$$

Thus for a redundant \mathcal{K}^* point $\mathbf{z}^{J+1} \in \text{cone}(\mathcal{Z})$, the polyhedral relaxation of the conic constraint $\mathbf{b} - \mathbf{A}\mathbf{x} \in \mathcal{K}$ implied by $\mathcal{Z} \cup \{\mathbf{z}^{J+1}\}$ is no stronger than that implied by \mathcal{Z} alone. An *extreme ray* of \mathcal{K}^* is a point $\mathbf{z} \in \mathcal{K}^*$ that cannot be written as a nontrivial conic combination of other points in \mathcal{K}^* that are not positive rescalings of \mathbf{z} . To maximize the efficiency of our polyhedral relaxations, we propose adding only extreme rays of \mathcal{K}^* to the \mathcal{K}^* point set \mathcal{Z} maintained by algorithm 1.

Recall from section 1.2 that our closed convex cone \mathcal{K} is encoded as a Cartesian product $\mathcal{K} = \mathcal{K}_1 \times \dots \times \mathcal{K}_K$ of standard primitive cones $\mathcal{K}_1, \dots, \mathcal{K}_K$ (e.g. nonnegative, second-order, exponential, and positive semidefinite cones). A primitive closed convex cone cannot be written as a Cartesian product of two or more lower-dimensional closed convex cones [Friberg, 2016]. If \mathcal{K} is separable, then its dual cone \mathcal{K}^* is also separable:

$$\mathcal{K}^* = (\mathcal{K}_1 \times \dots \times \mathcal{K}_K)^* = \mathcal{K}_1^* \times \dots \times \mathcal{K}_K^*. \quad (22)$$

We exploit this separability and our understanding of the structure of the standard primitive cones to disaggregate a \mathcal{K}^* point \mathbf{z} into extreme rays of \mathcal{K}^* .

First, we note that $\mathbf{z} = (\tilde{\mathbf{z}}^1, \dots, \tilde{\mathbf{z}}^K) \in \mathcal{K}^*$, where $\tilde{\mathbf{z}}^k \in \mathcal{K}_k^*, \forall k \in \llbracket K \rrbracket$. Second, for each $k \in \llbracket K \rrbracket$, we disaggregate $\tilde{\mathbf{z}}^k$ into extreme rays of the primitive standard dual cone \mathcal{K}_k^* . This step is trivial for linear cones. For second-order, positive semidefinite, and exponential cones, we describe practical computational procedures for dual disaggregation in appendix A.¹⁰ We have $\tilde{\mathbf{z}}^k = \sum_{j \in \llbracket J_k \rrbracket} \tilde{\mathbf{z}}^{k,j}$, where $\tilde{\mathbf{z}}^{k,j} \neq \mathbf{0}$ is an extreme ray of \mathcal{K}_k^* , for all $j \in \llbracket J_k \rrbracket$. We choose these extreme rays so that none is a positive scaling of another, and J_k does not exceed $\dim(\mathcal{K}_k^*)$. Note that $J_k = 0$ if $\tilde{\mathbf{z}}^k = \mathbf{0}$.

For some $k \in \llbracket K \rrbracket$ and $j \in \llbracket J_k \rrbracket$, consider the point $\mathbf{z}^{k,j} = (0, \dots, 0, \tilde{\mathbf{z}}^{k,j}, 0, \dots, 0)$, which is nonzero only on the elements corresponding to the k th primitive dual

¹⁰For example, if \mathcal{K}_k is a positive semidefinite cone, we disaggregate $\tilde{\mathbf{z}}^k \in \mathcal{K}_k^*$ by performing an eigendecomposition on it; see appendix A.3.2.

cone. Since any cone contains the origin $\mathbf{0}$, and $\tilde{\mathbf{z}}^{k,j} \in \mathcal{K}_k^*$, $\mathbf{z}^{k,j} \in \mathcal{K}^*$ by equation (22). Furthermore, since $\tilde{\mathbf{z}}^{k,j}$ is an extreme ray of \mathcal{K}_k^* , it cannot be written as a nontrivial sum of extreme rays of \mathcal{K}_k^* , and so $\mathbf{z}^{k,j}$ cannot be written as a nontrivial sum of extreme rays of \mathcal{K}^* . Thus $\mathbf{z}^{k,j}$ is an extreme ray of \mathcal{K}^* .

Our two-stage disaggregation procedure for $\mathbf{z} \in \mathcal{K}^*$ yields $\sum_{k \in \llbracket K \rrbracket} J_k \leq \dim(\mathcal{K}) = M$ extreme rays of \mathcal{K}^* :

$$\mathbf{z} = \sum_{k \in \llbracket K \rrbracket} \sum_{j \in \llbracket J_k \rrbracket} \mathbf{z}^{k,j}. \quad (23)$$

Besides adding potentially multiple \mathcal{K}^* points to \mathcal{Z} , no modifications are needed to the description of algorithm 1. Since \mathbf{z} is clearly contained in the conic hull of these \mathcal{K}^* points, there is no loss of strength in the polyhedral relaxations, so the certificate \mathcal{K}^* guarantees from section 3.1 are maintained. The polyhedral relaxations are potentially much tighter, improving the power of the LP OA for fathoming a node by infeasibility or objective bound without proceeding to an expensive conic subproblem solve.¹¹

We can also recover the guarantees from section 3.2 for an LP solver with a feasibility tolerance $\delta > 0$. We assume \mathbf{z} is a certificate \mathcal{K}^* point that has already been scaled according to section 3.2. After disaggregating \mathbf{z} , we scale each extreme ray up by $J = \sum_{k \in \llbracket K \rrbracket} J_k$ before adding it to \mathcal{Z} . The J relaxed scaled disaggregated \mathcal{K}^* cut conditions are:

$$(J\mathbf{z}^{k,j})^T(\mathbf{b} - \mathbf{A}\mathbf{x}) \geq -\delta \quad \forall k \in \llbracket K \rrbracket, j \in \llbracket J_k \rrbracket. \quad (24)$$

Summing and using equation (23), and dividing by J , we see that these conditions imply the relaxed scaled original \mathcal{K}^* cut condition $\mathbf{z}^T(\mathbf{b} - \mathbf{A}\mathbf{x}) \geq -\delta$.

4.2 Initial Fixed Polyhedral Relaxations

We can modify algorithm 1 on line 2 to initialize a nonempty set \mathcal{Z} of *initial fixed* \mathcal{K}^* extreme rays that are not derived from subproblem certificates, but depend only on the geometry of \mathcal{K}^* . If \mathcal{K} is a separable product of standard primitive cones, we can obtain initial fixed \mathcal{K}^* extreme rays by treating each primitive cone constraint separately. In particular, a linear cone constraint need not be relaxed at all, since it is equivalent to one \mathcal{K}^* cut (for a nonnegative or nonpositive cone) or two \mathcal{K}^* cuts (for the zero cone). In appendix A, we describe simple sets of initial fixed \mathcal{K}^* extreme rays for second-order, positive semidefinite, or exponential primitive cones.¹² We show in appendix A how knowledge of the initial fixed \mathcal{K}^* extreme rays allows us to tailor our extreme ray disaggregation procedures from section 4.1 for certificate \mathcal{K}^* points to further

¹¹The LP solver may need to deal with more cuts at nodes visited early in the search tree, but is ultimately likely to need to examine fewer nodes overall and solve fewer expensive conic subproblems, so the tradeoff can be worthwhile.

¹²For example, for a positive semidefinite cone, we use the extreme rays of the polyhedral cone of diagonally dominant symmetric matrices as initial fixed \mathcal{K}^* extreme rays; see appendix A.3.1.

increase the strength of the polyhedral relaxations and reduce redundancy in \mathcal{Z} .¹³

4.3 Separation Of Infeasible Points

Inspired by separation-based OA algorithms, we can modify algorithm 1 on line 16 to add *separation \mathcal{K}^* points* to \mathcal{Z} that cut off a fractional optimal LP solution $\hat{\mathbf{x}}$ that violates the conic constraint, right before branching on $\hat{\mathbf{x}}$. We show that a separation \mathcal{K}^* point exists when $\mathbf{b} - \mathbf{A}\hat{\mathbf{x}} \notin \mathcal{K}$. Since \mathcal{K} is closed and convex, there exists a hyperplane (\mathbf{z}, θ) that separates $\hat{\mathbf{y}} = \mathbf{b} - \mathbf{A}\hat{\mathbf{x}}$ from \mathcal{K} , i.e. $\mathbf{z}^T \hat{\mathbf{y}} < \theta$ and $\mathbf{z}^T \mathbf{y} \geq \theta, \forall \mathbf{y} \in \mathcal{K}$. Since the problem $\inf_{\mathbf{y} \in \mathcal{K}} \mathbf{z}^T \mathbf{y}$ is homogeneous (as \mathcal{K} is a cone) and the optimal value is bounded below by finite θ , the optimal value must equal zero. So $\theta \leq 0$, implying $\mathbf{z}^T \hat{\mathbf{y}} < 0$ and $\mathbf{z}^T \mathbf{y} \geq 0, \forall \mathbf{y} \in \mathcal{K}$. Thus $\mathbf{z} \in \mathcal{K}^*$ (by definition (5) of \mathcal{K}^*), and it implies a \mathcal{K}^* cut that separates $\hat{\mathbf{x}}$ from the feasible set of the conic constraint.

A separation \mathcal{K}^* point may fail to improve the objective lower bound from the LP OA, and does not in general possess the sort of guarantees from section 3 that a certificate \mathcal{K}^* point implies. However, deriving a separation \mathcal{K}^* point can be much cheaper than solving a continuous conic subproblem. If \mathcal{K} is a separable product of standard primitive cones, we can obtain separation \mathcal{K}^* extreme rays easily by treating each primitive cone constraint separately. In appendix A, we describe practical computational methods for obtaining separation \mathcal{K}^* extreme rays for primitive conic constraints involving second-order, positive semidefinite, or exponential cones.¹⁴

5 Pajarito Solver And Related Software

We describe the software architecture and algorithmic implementation of Pajarito, our open source MI-convex solver.¹⁵ This section may be of particular interest to advanced users and developers of mathematical optimization software. We emphasize that our implementations diverge from the idealized description of algorithm 1 in section 2.3, because of our decision to leverage powerful external mixed-integer linear (MILP/MIP) solvers through limited, solver-independent interfaces. Developers of MI-conic software with low-level control of the MIP search tree are able to implement features of algorithm 1 that we are not capable of in Pajarito.

¹³However, to be able to recover the guarantees from section 3.2 under an LP solver with a feasibility tolerance, we would need the ability to dynamically scale up the initial fixed \mathcal{K}^* points.

¹⁴For example, we obtain separation \mathcal{K}^* extreme rays for a point that violates a positive semidefinite cone constraint by performing an eigendecomposition on it; see appendix A.3.3.

¹⁵In Pajarito's readme file (github.com/JuliaOpt/Pajarito.jl) we provide more guidance on the recommended ways of using the solver, as well as default options and tolerances.

5.1 Julia And MathProgBase

Pajarito is the first MI-convex solver written in the relatively young Julia language [Bezanson et al., 2017]. MI-convex solvers such as α -ECP, Artelys Knitro, Bonmin, DICOPT, FilMINT, MINLP_BB, and SBB, which are reviewed by Bonami et al. [2012], are to our knowledge written in C, C++, or Fortran. Julia is a high-level programming language that can match the performance of these lower-level languages for writing solvers with much less boilerplate code [Lubin and Dunning, 2015]. Pajarito’s compact codebase is thoroughly commented, and conveniently reusable and extensible by other researchers. We implement an extensive testing infrastructure with hundreds of unit tests. Since Pajarito’s first release, several other MINLP solvers have been written in Julia and are available through MathProgBase, such as POD [Nagarajan et al., 2017], Juniper [Kröger et al., 2018], and Katana.¹⁶

Pajarito is integrated with the powerful MathProgBase abstraction layer. MathProgBase is a standardized API in Julia for interacting with optimization solvers, designed in part to allow the user to write solver-independent code.¹⁷ The breadth of problem classes covered by MathProgBase is described at juliaopt.org and distinguishes it from similar abstraction layers such as OSI [Saltzman et al., 2004], a COIN-OR library in C++. It includes specifications for continuous and mixed-integer solvers that use linear/quadratic, conic, or oracle-based NLP (nonlinear programming) forms.

In section 5.2, we describe accessing Pajarito through MathProgBase’s conic interface (see top of figure 1). The user specifies external MIP and continuous primal-dual conic solvers (including solver options) from the available solvers (i.e., those accessible through MathProgBase) and passes each solver object as an option into a function that creates a Pajarito solver object. In section 5.3, we summarize Pajarito’s main algorithmic implementations. Pajarito uses the modeling package JuMP to conveniently build and manage the external MIP solver’s OA model. JuMP itself interacts with the MIP solver via MathProgBase’s linear/quadratic interface (see bottom right of figure 1). To solve a continuous conic subproblem for a conic certificate, Pajarito calls the external primal-dual conic solver through the conic interface (see bottom left of figure 1).¹⁸

5.2 Accessing Pajarito

Pajarito’s use of conic form is a significant architectural difference from most existing MI-convex solvers, which interact with a MI-convex instance almost exclusively through oracles to query values and derivatives of the constraint and objective functions. MathProgBase conic form can be described compactly from a constraint matrix in sparse or dense format, right-hand side and objective

¹⁶See github.com/lanl-ansi/POD.jl, github.com/lanl-ansi/Juniper.jl, and github.com/lanl-ansi/Katana.jl.

¹⁷MathProgBase is being replaced by a redesigned API, MathOptInterface. The process of building Pajarito has motivated many of the planned improvements in MathOptInterface.

¹⁸MathProgBase documents the conic and linear/quadratic interfaces at mathprog-basejl.readthedocs.io/en/latest. JuMP is documented at juliaopt.org/JuMP.jl/0.18/.

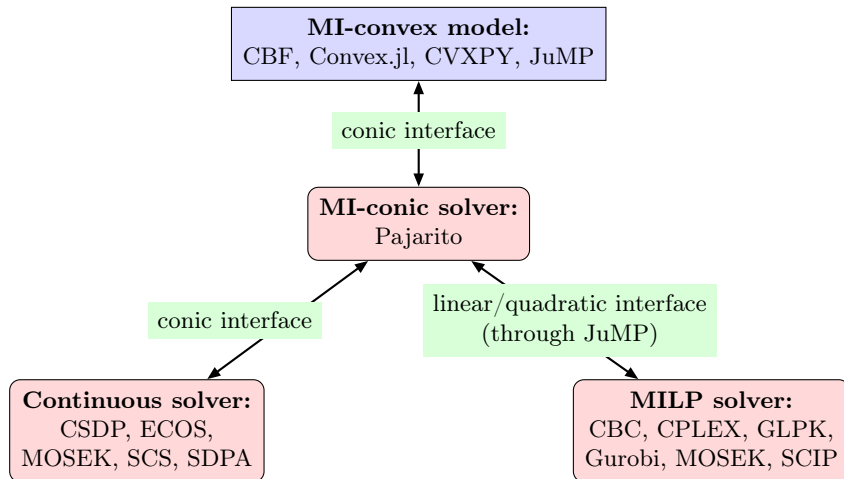


Figure 1: Pajarito’s integration with MathProgBase.

coefficient vectors, variable and constraint cones expressed as lists of standard primitive cones (1-dimensional vector sets) with corresponding ordered row indices, and a vector of variable types (each continuous, binary, or general integer). In addition to the basic linear cones (nonnegative, nonpositive, zero, and free cones), Pajarito recognizes three standard primitive nonpolyhedral cones introduced in section 1.2: exponential cones (see appendix A.1), second-order cones (see appendix A.2), and positive semidefinite cones (see appendix A.3).¹⁹

Friberg [2016] designed the Conic Benchmark Format (CBF) as a file format originally to support mixed-integer second-order cone (SOCP) and positive semidefinite cone (SDP) instances. In collaboration with Henrik Friberg, we extended the format to support exponential cones in Version 2, and developed a Julia interface `ConicBenchmarkUtilities.jl` to provide utilities for translating between CBF and MathProgBase conic format.²⁰ One may use Pajarito to solve any instance in the Conic Benchmark Library (CBLIB), which contains thousands of benchmark problems from a wide variety of sources.

Lubin et al. [2016] demonstrate that all 333 known MI-convex instances in MINLPLib2 [Vigerske, 2018] are representable with linear, second-order, exponential, and *power* cones. Since a power cone constraint is representable with linear and exponential cone constraints, Pajarito can be used to solve any of the MI-convex instances in MINLPLib2. We translated 115 instances from the MINLPLIB2 library to CBF and contributed them to CBLIB.²¹ Many of the

¹⁹As we note in appendix A.2, Pajarito also recognizes rotated second-order cones, but for simplicity converts them to second-order cones during preprocessing.

²⁰Pajarito’s extensive unit tests rely on small example instances loaded from CBF files.

²¹Lubin et al. [2016] first translated these instances from the MINLPLIB2 library into `Convex.jl` models. We used `ConicBenchmarkUtilities.jl` to translate these to CBF. The instances, available at github.com/mlubin/MICPExperiments, are 48 ‘rsyn’ instances, 48 ‘syn’ instances, 6 ‘tls’ instances, 12 ‘clay’ instances, and the challenging ‘gams01’ instance.

MINLPLIB2 instances have tiny values artificially-introduced in order to work around potential numerical issues with smooth derivative-based NLP solvers [Günlük and Linderoth, 2012], which we manually removed before converting to conic form. For example, the instance ‘rsyn0805h’ has a constraint:

$$\left(\frac{x_{289}}{10^{-6} + b_{306}} - \frac{6}{5} \log \left(1 + \frac{x_{285}}{10^{-6} + b_{306}} \right) \right) (10^{-6} + b_{306}) \leq 0, \quad (25)$$

where $b_{306}, x_{285}, x_{289}$ are scalar variables. Without the artificial 10^{-6} values, a conic encoding of the NLP constraint (25) in terms of the exponential cone \mathcal{E} is:

$$(b_{306} + x_{285}, b_{306}, {}^{5/6}x_{289}) \in \mathcal{E}. \quad (26)$$

Within Julia, the modeling packages JuMP [Dunning et al., 2017] and Convex.jl [Udell et al., 2014] each provide a convenient way for users to specify MI-convex problems, call Pajarito solver, and interpret solutions. JuMP is particularly useful for a large, sparse problem involving complex indexing schemes for variables, expressions, or constraints. It efficiently builds a MathProgBase conic form representation of a problem involving second-order or positive semidefinite cones, but currently does not recognize exponential cones.

Convex.jl, unlike JuMP, is a Disciplined Convex Programming (DCP) modeling package. It defines a list of *atoms* for the user to model a MI-convex problem with and performs automatic verification of convexity of the continuous relaxation by applying simple composition rules described by Grant et al. [2006]. It converts the problem into a MI-conic instance in MathProgBase conic form through epigraph and perspective transformations that introduce additional variables and constraints in conic form using only the standard primitive cones recognized by Pajarito. CVXPY [Diamond and Boyd, 2016] is a Python-based DCP modeling package analogous to Convex.jl. In collaboration with Steven Diamond and Baris Ungun, we developed cmpb.jl (github.com/mlubin/cmpb), a prototype C API to MathProgBase that enables Pajarito to be called on a problem modeled with CVXPY.

We illustrate Convex.jl and JuMP modeling using a simple MI-convex example described by Boyd and Vandenberghe [2004, ch. 7.5]: ‘E-optimal experimental design’.²² While we can solve E-optimal experimental design exactly using Pajarito, Boyd and Vandenberghe [2004, ch. 7.5] choose to relax the integrality constraints in order to use a continuous convex solver before rounding the fractional solution heuristically. To begin, we set up the Pajarito solver object using a GLPK MILP solver object and a SCS conic solver object, each with internal options set. Pajarito performs a sanity check on the combination of options and solvers specified.²³ Next, we model and solve the problem using Convex.jl as

²²More Pajarito examples can be found at github.com/JuliaOpt/Pajarito.jl/blob/master/examples.

²³MathProgBase does not attempt to provide an abstraction for solver parameters like convergence tolerances. In cases where we need certain tolerances on the continuous conic and MIP solvers in order for Pajarito to converge to a requested tolerance, it is the user’s responsibility to set these tolerances. For example, we ask users to manually adjust the MIP solver’s linear feasibility tolerance and integer feasibility tolerance for improved convergence behavior. These cases are documented in Pajarito’s readme file.

follows, where $p, m, n \in \mathbb{R}$ and $V \in \mathbb{R}^{n \times p}$ are problem data. Alternatively, we model and solve the problem using JuMP as follows. Pajarito manipulates the conic data and performs sanity checks. We refer to the resulting preprocessed representation of the instance as \mathfrak{M} . After Pajarito executes one of the OA algorithms described in section 5.3 on \mathfrak{M} , the user can use Convex.jl or JuMP to conveniently query information such as Pajarito’s solve status, objective bound, objective value, and solution.

5.3 Basic Algorithmic Implementations

We discuss the main conic-certificate-based methods Pajarito uses to solve the preprocessed MI-conic model \mathfrak{M} . We omit many options, enhancements, and numerical details that can be understood from the Pajarito readme file and from browsing the high-level Julia code and comments.²⁴ In section 5.3.1, we summarize the initialization procedure for the OA model, an MILP relaxation of \mathfrak{M} that Pajarito constructs and later refines (with extreme ray \mathcal{K}^* cuts) using JuMP. In section 5.3.2, we describe the ‘iterative’ method, an extension of the simple sequential OA algorithm by Lubin et al. [2017a]. In section 5.3.3, we describe the ‘MIP-solver-driven’ (MSD) method, so-called because it relies on the power of the branch-and-cut MIP solver to manage convergence in a single tree. Since MathProgBase’s solver-independent abstraction for MIP solver callbacks is designed primarily around shared behavior between CPLEX and Gurobi, Pajarito is limited to interacting with the MIP solver through a *lazy cut callback* function and a *heuristic callback* function. Although the MSD method is generally much faster than the iterative method, the latter may be used with MILP solvers for which callback functionality is unavailable or unreliable.

5.3.1 Initializing The MIP OA Model

We first solve the continuous relaxation of \mathfrak{M} (in which only the integrality constraints are relaxed), using the primal-dual conic solver (see top of figure 2).²⁵ If the conic solver indicates this relaxation is infeasible, then \mathfrak{M} must be infeasible, so we terminate with an ‘infeasible’ status. If the conic solver returns a complementary solution pair, the optimal value gives an objective lower bound $L > -\infty$ for \mathfrak{M} . Otherwise, we set $L = -\infty$. We initialize the objective upper bound U for \mathfrak{M} to ∞ .

Using JuMP, we build the initial OA model, adding the variables and integrality constraints and setting the objective (see bottom of figure 2). We then add initial fixed cuts for each primitive cone, as we describe in section 4.2. Primitive linear cone constraints are imposed entirely (as equivalent LP equality

²⁴For explaining computational experiments, section 6.3 briefly introduces several other algorithmic variants that we do not discuss here, such as separation-based methods that do not utilize conic certificates.

²⁵The conic relaxation is analogous to the first node subproblem in algorithm 1, but without finite integer variable bounds. We preprocess this conic model slightly to tighten any non-integral bounds on the integer variables.

or inequality constraints), and for each primitive nonpolyhedral cone, we add a small number of initial fixed cuts (defined in appendix A).

A complementary solution pair from the conic relaxation solve yields a \mathcal{K}^* point, so we perform an extreme ray disaggregation from section 4.1 and add certificate cuts for each primitive nonpolyhedral cone (using the procedures in appendix A). These continuous relaxation certificate cuts technically guarantee that the root node of the OA model has an optimal value no smaller than L .²⁶ This is important because we cannot handle unboundedness of the OA model.

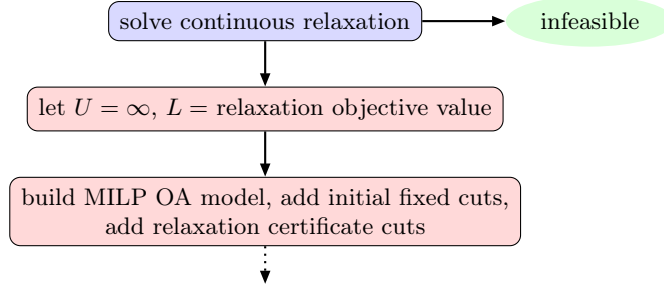


Figure 2: Pajarito’s OA model initialization.

5.3.2 Iterative Method

The iterative method, following initialization in figure 2, is outlined in figure 3. At each iteration of the main loop, Pajarito solves the current OA model using the MIP solver.²⁷ If the OA model is infeasible, \mathfrak{M} must be infeasible, so we terminate with an ‘infeasible’ status. If it is unbounded, Pajarito fails with an ‘OA fail’ error status, as we are unable to handle unbounded rays. If the MIP solver returns an optimal solution to the OA model, this OA solution satisfies the integrality constraints and initial fixed cuts, but in general not all of nonpolyhedral primitive cone constraints. The MIP solver’s objective bound provides a lower bound for \mathfrak{M} , so we update L . Pajarito terminates with an ‘optimal’ status if the relative optimality gap condition (17) on L, U is satisfied.²⁸

If after solving the OA model we have an optimal OA solution and the objective bounds haven’t converged, we check whether the OA sub-solution on the integer variables has been encountered before. If so, we check the conic feasibility of the OA solution. We calculate the absolute violation on each primitive nonpolyhedral cone constraint as the violation of the appropriate separation cut (defined in appendix A). If the worst absolute violation does not exceed Pajarito’s feasibility tolerance (set by the user), then the OA solution

²⁶This can be seen from a simple modification of the complementary solution case polyhedral relaxation guarantee we prove in section 3.1.2, with trivial integer variable bounds.

²⁷We suggest the user set the MIP solver’s relative optimality gap tolerance to its smallest possible value.

²⁸Pajarito uses $\theta = 10^{-5}$ (to avoid division by 0). The gap tolerance $\epsilon > 0$ is specified by the user, but defaults to 10^{-5} .

is considered feasible. In this case, since the solution is optimal for the OA model, we can consider it optimal for \mathfrak{M} , so we update the incumbent and upper bound and terminate the solve immediately. If the OA solution is not considered feasible, we add all of the separation cuts that are (significantly) violated to the OA model.

If the integer sub-solution has not already been encountered at a previous iteration, then we solve a continuous conic subproblem in which the integer variables are fixed to their values in the integer sub-solution. This subproblem is analogous to $\mathfrak{C}(\mathbf{l}, \mathbf{u})$ from section 2.1, with $\mathbf{l} = \mathbf{u}$.²⁹ Since it is more constrained than the OA model, the conic subproblem is bounded or infeasible. If the conic subproblem solver fails to return a certificate, we backtrack and perform the separation procedure (as if the integer sub-solution repeated). Otherwise, we scale the certificate’s dual solution or dual ray according to section 3.2 (using the tolerance values set as Pajarito options), then disaggregate the scaled \mathcal{K}^* point and add extreme ray certificate cuts to the OA model (as we described for the continuous relaxation certificate in section 5.3.1). In the case of a complementary solution pair, the primal solution yields a feasible point for \mathfrak{M} , since it satisfies both the integrality and conic constraints. If it has an objective value better than U , we update U and the incumbent solution and check the relative optimality gap condition again.³⁰

After adding separation or certificate cuts, we warm-start the MIP solver with our incumbent and re-execute the main loop. The procedure in figure 3 is iterated until L and U converge or the MIP solver detects infeasibility.³¹ Note that since we only add cuts to the OA model on every loop, if the first OA model is bounded, then all subsequent (refined) OA models are bounded or infeasible, and the sequence of lower bounds L is nondecreasing.³²

5.3.3 MIP-Solver-Driven Method

The MSD method, following initialization in figure 2, is outlined in figure 4. As in the iterative method, Pajarito returns an ‘OA fail’ status if the MIP solver detects unboundedness (as we are unable to handle unbounded rays), or an ‘infeasible’ status in the case of infeasibility.³³ The MIP-solver-independent call-

²⁹In preprocessing, we remove any subproblem equality constraints that effectively have no variables when an integer sub-solution is fixed. For efficient loading of the subproblem data at each iteration, we only change the constant vector \mathbf{b} of the preprocessed conic subproblem, as this is the only data that changes.

³⁰Conic solvers typically do not use an absolute primitive constraint-wise feasibility tolerance, as Pajarito does for checking feasibility of OA solutions for the conic constraint. Our incumbent may not satisfy this notion of feasibility, since we do not perform a feasibility check on the conic solver’s primal subproblem solutions.

³¹If the user sets a time limit, Pajarito may terminate with the status ‘user limit’. Pajarito sets the time limit on each MIP or conic solve to the remaining time. Note that the vast majority of Pajarito execution time is spent in MIP or conic solves.

³²In case of failures of strong duality at some conic subproblems, Pajarito may fail to converge, as there exists no finite set of cuts that can tighten the lower bound sufficiently to meet the upper bound. See Lubin et al. [2016] for a discussion of strong duality in OA.

³³If the user sets a time limit, Pajarito sets a time limit on the MIP solver, and terminates with a ‘user limit’ status if this limit is reached.

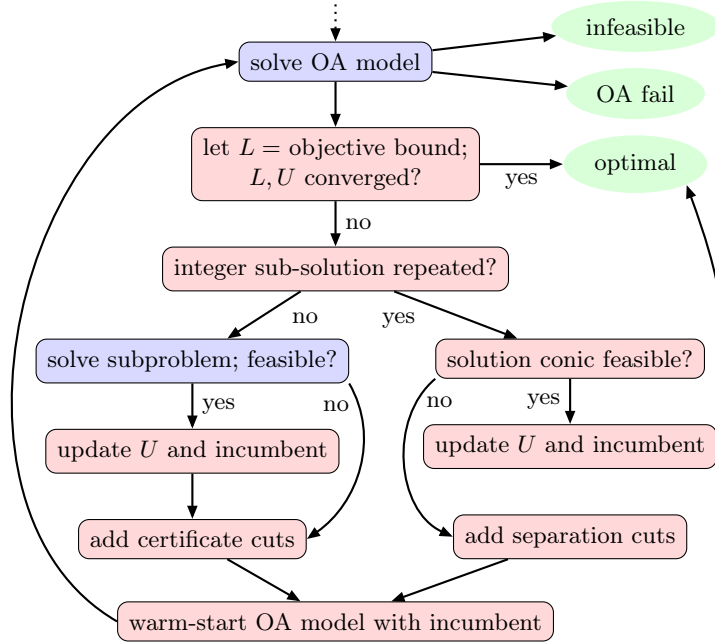


Figure 3: Pajarito's iterative method, following initialization.

back interface allows us to pass in lazy cuts in during a lazy callback and feasible solutions during a heuristic callback, however we cannot exert any control over branching decisions, node selection, fathoming, or node lower bound updating.

The MIP solver calls the lazy callback function whenever it finds an integer-feasible OA solution at a node. During a lazy callback, we first check whether the integral OA solution from the MIP solver is repeated. If so, we derive separation cuts to add as lazy constraints; if none can be added, the MIP solver considers the solution feasible and may update its incumbent. If the integer sub-solution is repeated, we solve a new (bounded or infeasible) conic subproblem. Since we lack the ability to query the node's integer variable bounds, we only solve subproblems with fixed integer sub-solutions, as in the iterative method. If the conic solver returns a certificate, we scale and disaggregate the \mathcal{K}^* point (as we described for the iterative method in section 5.3.2), and add extreme ray cuts as lazy constraints.³⁴ In the case of a complementary solution pair, the primal solution yields a feasible point for \mathfrak{M} , which we store. During a heuristic callback, if there is a stored feasible solution to \mathfrak{M} that has never been added as a heuristic solution, we add it.

Since there are no guarantees on when or how frequently the MIP solver

³⁴The MIP solver is not guaranteed to respect the cuts that we add, and we may need to re-add the same cuts during multiple lazy callbacks (unlike in the iterative method, where cuts previously added are respected). We actually store a dictionary from the integer sub-solution to the cuts. For each repeated integer sub-solution, we re-add these saved certificate cuts, in addition to the new separation cuts.

calls the heuristic callback function, we may not be able to indirectly update the MIP solver’s incumbent and upper bound when we are able to. Partly for this reason, Pajarito maintains its own upper bound and incumbent (not illustrated in figure 4, which we update during lazy callbacks. During each lazy callback, we ask the MIP solver for its lower bound and check our relative optimality gap condition (as we described for the iterative method). If the condition is met, we force the MIP solver to terminate early. In this case, or if the MIP solver terminates with an optimal solution and we verify that the relative optimality gap condition is met, we return our incumbent solution with an ‘optimal’ status.³⁵

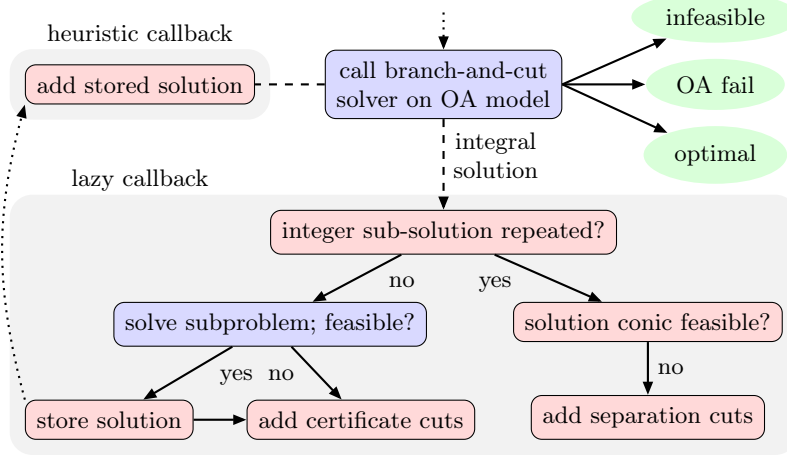


Figure 4: Pajarito’s MIP-solver-driven (MSD) method, following initialization.

5.4 Some Advanced Algorithmic Enhancements

We conclude with several key optional OA enhancements we implemented in Pajarito. First, Pajarito by default uses an *extended formulation* for each second-order cone constraint. Vielma et al. [2017] demonstrate on a testset of mixed-integer second-order cone (MISOCP) problems that OA algorithms tend to converge much faster when using this extended representation for each second-order cone constraint.³⁶ In appendix B, we describe how to lift a \mathcal{K}^* cut for the second-order cone into \mathcal{K}^* cuts for the extended formulation. This technique also allows us to describe a much more economical set of initial fixed \mathcal{K}^* cuts for the second-order cone.

Second, Pajarito can optionally use a MISOCP OA model instead of a MILP OA model. There exist several powerful MISOCP solvers that can be used, or

³⁵Note the user is responsible for setting the desired relative optimality gap tolerance on both the MIP solver and on Pajarito directly.

³⁶DCP modeling software implementations such as Convex.jl do not perform this transformation because they are simply designed to access conic solvers. Pajarito keeps the original second-order cone formulation in the conic subproblems because conic solvers are likely to perform better with this representation than with the higher-dimensional extended formulation.

Pajarito itself may be used. Since a second-order cone constraint can imply an infinite number of \mathcal{K}^* cuts, Pajarito can achieve tighter relaxations of the conic constraint in the OA model. Of course, this can only make practical sense for certain types of problems that aren't pure MISOCP. One potential use case is where the MI-conic problem has second-order cones as well as exponential and/or positive semidefinite cones, but the only stable and efficient continuous conic solvers we have access to are for SOCP problems. In this case, we can use Pajarito with the SOCP solver as an MISOCP solver, and pass this into a second Pajarito solver that uses a conic solver for mixed-cone problems. This arrangement helps minimize the number of calls to the less-effective conic solver. Another use case is for problems involving positive semidefinite cone constraints. We demonstrate in appendix C how to strengthen \mathcal{K}^* cuts for PSD constraints to rotated-second-order cone constraints.³⁷

6 Computational Experiments

Our computational experiments demonstrate the speed and robustness of our open source MI-conic solver Pajarito. As we emphasize in section 5.3, our algorithmic implementations differ from the description of algorithm 1, because of our practical decision to use branch-and-cut MILP solvers through a limited, solver-independent interface. In section 6.1, we summarize our metrics for comparing the practical performance of different MI-conic solvers and describe our presentation of tables and performance profile plots. In section 6.2, we benchmark Pajarito (version 0.5.1) and several open source and commercial mixed-integer second-order cone (MISOCP) solver packages accessible through MathProgBase on a MISOCP library, and conclude that Pajarito is the fastest and most-reliable open source solver for MISOCP. In section 6.3, we compare the performance of several of Pajarito's algorithmic variants on MI-conic instances involving mixtures of positive semidefinite, second-order, and exponential cones, demonstrating practical advantages of the methodological extensions we describe in sections 3 and 4 and appendix A. The scripts and data we use to run our experiments are available in the supplement github.com/chriscoey/PajaritoSupplement.

6.1 Presentation Of Results

We define a 'solver' as a MathProgBase solver object given a particular complete set of algorithmic options.³⁸ We define an 'instance' as a particular MI-conic problem (stored in CBF format; see section 5.2) that is known to be feasible and bounded (but an optimal solution or the optimal objective value is not necessarily known). For a particular instance, a solver may return a 'solution',

³⁷For the MSD method, since most MISOCP solvers don't currently allow adding lazy quadratic constraints, only the initial fixed cuts can be strengthened in this way.

³⁸Each solver we test is deterministic, i.e. it performs consistently across different runs on a particular dedicated system.

which is a vector of real floating point numbers representing an assignment of the variables of the instance (not necessarily feasible for the constraints).

First, we compare the performances of a group of MI-conic solvers on a particular testset of instances by counting the number of instances for which each solver returns and apparently proves ‘approximate optimality’ of a solution. To be more precise, we use the following four categories to characterize a solver’s apparent success or failure on an instance.

ex (exclude) means either the solver incorrectly claims the instance is infeasible or unbounded, or the solver returns a solution it claims is approximately optimal but we detect one of the following inconsistencies.

- The solution significantly violates at least one primitive cone constraint or integrality constraint.³⁹
- The relative objective gap condition (equation (17)) for optimality is significantly violated.⁴⁰
- The objective value or objective bound significantly differs from that of a preponderance of other solvers.⁴¹

co (converge) means the solver returns a solution that it claims is (approximately) optimal (and it is not excluded for the reasons above).

li (reach limit) means the solver does not terminate before the time limit, or (rarely) the solver reaches a memory limit and is forced to terminate.

er (error) means the solver crashes or terminates with an error message.

Second, we compare aggregate quantitative measures of solver performance. We define the shifted geometric mean \tilde{g} of L positive values p_1, \dots, p_L as:

$$\tilde{g}(\mathbf{p}, q) = \prod_{l \in [L]} (p_l + q)^{\frac{1}{L}} - q, \quad (27)$$

where $q > 0$ is the shift [Achterberg, 2009]. Unlike the standard geometric mean $\tilde{g}(\mathbf{p}, 0)$, the shifted geometric mean decreases the relative influence of smaller values in \mathbf{p} , thus giving less weight to very ‘easy’ instances (small values are preferable for all of our metrics). We shift by $q = 10$ seconds for execution times, $q = 1$ iterations for iteration counts, and $q = 10$ nodes for MIP-solver-reported node counts. For comparing a particular group of solvers S_1, \dots, S_n on a particular performance metric (such as execution time), we calculate for

³⁹ Absolute violation of a primitive cone constraint is calculated as worst violation of the inequalities defining the standard cone (see appendix A), and our tolerances are 10^{-6} for linear cones, 10^{-5} for second-order and exponential cones, and 10^{-4} for positive semidefinite cones. Variable-wise integrality violation is calculated as distance to the nearest integer, and our tolerance is 10^{-6} .

⁴⁰ Our optimality condition (17) matches that used by most MIP solvers. We set the constant $\theta = 10^{-5}$ (to avoid division by zero) and use the tolerance $\epsilon = 10^{-5}$. We ensure we do not exclude in the case that the gap we calculate is sensitive to a solver’s different value of θ .

⁴¹ This is assessed semi-manually from output of our scripts.

each solver S_i the following three shifted geomeans, each over a different subset of the testset.

aco (all solvers converge) is calculated over the instances for which S_1, \dots, S_n all have a ‘co’ status.

tco (this solver converges) is calculated over the instances for which S_i has a ‘co’ status.

all (all instances) is calculated over all instances. Missing execution times are set to the time limit, and missing iteration/node counts are ignored.

Finally, we employ ‘performance profiles’, described by Dolan and Moré [2002], Gould and Scott [2016], to visually compare the relative execution times and iteration or node counts of pairs of solvers. Again, we decrease the relative influence of very easy instances by shifting the metrics by the same shift values q we use for shifted geomeans. A performance profile is a plot that should be interpreted as follows: for a fixed factor F on the horizontal axis (a linear scale from 1 to the value at the bottom right of the plot), the level of solver S_i on the vertical axis (a linear scale from 0 to 1) represents the proportion P_i of instances (out of the instances for which at least one of the pair of solvers has a ‘co’ status) for which S_i has a ‘co’ status and a (shifted) performance metric that is within a factor of F of per-instance best achieved by either solver. So, at $F = 1$ (i.e. on the left vertical axis), P_i is the fraction of solved instances on which solver S_i has the best performance. As F increases, we can infer that solver S_i has reliably better performance than solver S_j if P_i remains above P_j .

6.2 MISOCP Solver Performance Comparisons

Our open source Pajarito solvers, ‘Iter-GLPK’ and ‘Iter-CBC’, use the iterative method (see section 5.3.2) with ECOS [Domahidi et al., 2013] for continuous conic subproblems and CBC or GLPK for MILPs.⁴² Our two restricted-license Pajarito solvers, ‘Iter-CPLEX’ (using the iterative method) and ‘MSD-CPLEX’ (using the MIP-solver-driven method; see section 5.3.3), call MOSEK’s continuous conic solver and CPLEX’s MILP solver.

The open source Bonmin solver package is described in detail by Bonami et al. [2008] and uses CBC to manage branching and Ipopt to solve continuous NLP (derivative-based nonlinear programming) subproblems.⁴³ Our ‘Bonmin-BB’ solver uses the nonlinear B&B method (no polyhedral approximation), ‘Bonmin-OA’ uses the B&B OA method, and ‘Bonmin-OA-D’ is equivalent to the ‘Bonmin-OA’ solver but applied to transformed instances that use the second-order cone extended formulation we describe in section 5.4. Our two restricted-license MISOCP solvers are ‘SCIP’ and ‘CPLEX’. Unlike Bonmin,

⁴²We do not test the Pajarito’s MIP-solver-driven method with CBC or GLPK MIP solvers because their support for MathProgBase callbacks is limited.

⁴³We are unaware of any mainstream open source solvers designed for MISOCP. The functional representation of the second-order cone has points of nondifferentiability that may cause Bonmin to crash or suffer numerical issues.

these MISOCP solvers use the second-order cone extended formulation internally.⁴⁴

These nine MISOCP solvers are each given a relative optimality gap tolerance of 10^{-5} . The ‘SCIP’ and ‘CPLEX’ solvers are given an absolute linear-constraint-wise feasibility tolerance of 10^{-8} , and ‘CPLEX’ is given an integrality tolerance of 10^{-9} . The MILP solvers used by Pajarito are given an absolute linear-constraint-wise feasibility tolerance of 10^{-8} , an integrality tolerance of 10^{-9} , and a relative optimality gap tolerance of 0 for ‘Iter-GLPK’, ‘Iter-CBC’, and ‘Iter-CPLEX’ and 10^{-5} for ‘MSD-CPLEX’. Due to limited resources, we set a one hour time limit for each run of a solver on an instance, and run all solvers (including the MILP and conic solvers called by Pajarito) in single-threaded mode.

We use a testset of 120 MISOCP instances drawn from the larger CBLIB library, recently compiled by Friberg [2016]. The testset contains randomly selected subsets of most of the major families of models in CBLIB. We exclude instances that are not bounded and feasible, or are solved in under 5 seconds by all solvers, or are unable to be solved by all solvers in under an hour. Our computations are performed on the Amazon EC2 cloud computing platform with ‘m4.xlarge’ computing nodes having 16GB of RAM.⁴⁵ As the computing nodes are virtual machines, timing results on EC2 are subject to random variability, but repeated runs suggest the variation is sufficiently small to avoid impacting our conclusions. The nodes run Ubuntu 16.04 with Julia version 0.6.0. Version information for the Julia packages can be obtained from the supplement.

Table 1 summarizes the status counts and shifted geomeans of performance metrics on instance subsets (explained in section 6.1) for the nine MISOCP solvers on the 120 MISOCP instances. The Bonmin solvers fail on most instances, and overall solve significantly fewer instances than the open source Pajarito solvers. Pajarito tends perform faster using CBC rather than GLPK.⁴⁶ Figure 5a is a performance profile (explained in section 6.1) comparing the execution times of the open source Pajarito (with CBC) solver and the instance-wise best of the three Bonmin solvers. From these results, we claim that Pajarito with ECOS and CBC is the fastest and most reliable open source MISOCP solver.

Using CPLEX, Pajarito’s MSD method is significantly faster and more reliable than its iterative method.⁴⁷ The performance profile figure 5b compares the execution times of Pajarito’s MSD method using CPLEX’s MILP solver against CPLEX’s specialized MISOCP solver. The execution time comparisons between ‘CPLEX’ and ‘MSD-CPLEX’ are ambiguous, however we argue that, at least by our metrics, Pajarito is a more reliable MISOCP solver.

⁴⁴CPLEX is available under an academic or commercial licence, and SCIP is an academic solver that is not released under an OSI-approved open source license. We use CPLEX version 12.7.0 and SCIP version 4.0.0.


⁴⁵See aws.amazon.com/ec2/instance-types.

⁴⁶However, for most of the 9 excluded instances from ‘Iter-CBC’, we verify that CBC is responsible for the significant integrality violations that result in exclusion.

⁴⁷For ‘MSD-CPLEX’, the two errors occur where Pajarito claims a solution is suboptimal and has an objective gap no worse than 1.04×10^{-5} , and the one exclusion occurs where Pajarito’s solution violates a linear constraint by 9.78×10^{-6} .

	solver	statuses				time (s)		
		co	li	er	ex	aco	tco	all
open source	Bonmin-BB	34	44	11	31	38.0	83.8	463
	Bonmin-OA	25	53	29	13	64.2	64.5	726
	Bonmin-OA-D	30	48	29	13	15.1	61.6	610
	Iter-GLPK	56	60	3	1	2.0	29.7	377
	Iter-CBC	78	30	3	9	1.6	50.3	163
restricted	SCIP	74	35	8	3	3.2	41.5	160
	CPLEX	90	16	5	9	0.9	16.1	50
	Iter-CPLEX	86	26	0	8	0.4	37.0	106
	MSD-CPLEX	97	20	2	1	0.4	18.2	56

Table 1: MISOCP solver performance summary.



(a) Open source Bonmin (instance-wise best of 3) and Pajarito iterative solvers. (b) CPLEX MISOCP and Pajarito MSD solvers.

Figure 5: MISOCP solver execution time performance profiles.

6.3 Comparative Testing Of Algorithmic Variants

To compare the performance of several of Pajarito’s algorithmic variants, we use a testset of 95 MI-conic instances involving mixtures of positive semidefinite (PSD), second-order, and exponential cones. These instances are all bounded and feasible and come from the following four sources.⁴⁸

Discrete experimental design (14 instances). Recall from section 5.2 that Boyd and Vandenberghe [2004, Ch. 7.5] describes MI-convex experimental design problems. We generate ‘A-optimal’ and ‘E-optimal’ instances that include PSD cones, and ‘D-optimal’ instances that include PSD and exponential cones.

Portfolios with mixed risk constraints (16 instances). We formulate a portfolio problem that maximizes expected returns subject to some combinatorial constraints on stocks and three types of convex risk constraints on

⁴⁸Formulations for instances we generated can be found at github.com/JuliaOpt/Pajarito.jl/tree/master/examples.

subsets of stocks with known covariances. Each instance includes multiple exponential cones from entropy risk constraints, second-order cones from norm risk constraints, and PSD cones from robust norm risk constraints.

Retrofit-synthesis of process networks (32 instances). We select a representative subset of the two CBLIB families ‘syn’ and ‘rsyn’. Each instance includes exponential cones.

A subset of the MISOCP testset (33 instances). We select a representative subset of the CBLIB families ‘estein’, ‘ccknapsack’, ‘sssd’, ‘uflquad’, and ‘portfoliocard’.

We use Pajarito with Gurobi (version 7.5.2) as the MILP solver and MOSEK (version 9.0.0.29-alpha) as the continuous conic solver.⁴⁹ Pajarito is given a relative optimality gap tolerance of 10^{-5} . Gurobi is given an absolute linear-constraint-wise feasibility tolerance of 10^{-8} , an integrality tolerance of 10^{-9} , and a relative optimality gap tolerance of 0 when the iterative method is used and 10^{-5} when the MSD method is used. We set a one hour time limit for each run of a solver on an instance, and limit Gurobi and MOSEK to 8 threads. We run the computations on dedicated hardware with 16 Intel Xeon E5-2650 CPUs (2GHz) and 64GB of RAM. Repeated runs suggest the variation is sufficiently small to avoid impacting our conclusions. The machine runs Ubuntu 17.10 and Julia 0.6.2. Version information for the Julia packages can be obtained from the supplement.

6.3.1 Initial Fixed Cuts, Certificate Cuts, And Separation Cuts

Recall from section 5.3 that Pajarito by default uses three different types of \mathcal{K}^* cuts: initial fixed cuts, certificate cuts, and separation cuts. For both the iterative and MSD methods, we compare the following four important algorithmic variants of OA that use different combinations of these three cut types.

c means initial fixed cuts on linear primitive cones only, and certificate cuts on nonpolyhedral primitive cones.

cs means initial fixed cuts on linear primitive cones only, and certificate cuts on nonpolyhedral primitive cones, and separation cuts when apparently needed for convergence. The separation cuts allow us to cut off significantly infeasible OA solutions, so Pajarito can also obtain (approximately) feasible solutions from OA solutions found by the MILP solver.

ics means initial fixed cuts on all primitive cones, certificate cuts on nonpolyhedral primitive cones, and separation cuts when apparently needed for convergence. This is Pajarito’s default approach, as described in section 5.3.

⁴⁹MOSEK 9 is the first version to recognize exponential cones.

is means initial fixed cuts on all primitive cones, and separation cuts only. No conic solver is used, hence all (approximately) feasible solutions found are OA solutions.

Table 2 summarizes the status counts and shifted geomeans of performance metrics on instance subsets. Although the MSD method is significantly faster than the iterative method, we see similar relative performances for the four types of cuts under iterative versus MSD. When using certificate cuts only (‘c’), Pajarito often failed to converge to the desired optimality gap (though it typically came very close), likely due to the inexactness of the certificates from the numerical continuous conic solver. By also using separation cuts on repeated integer sub-solutions and accepting (approximately) conic feasible OA solutions as incumbents, Pajarito is able to converge on many more instances. Starting with initial fixed cuts (‘ics’) further increases Pajarito’s robustness, particularly for the MSD method. Comparing the ‘ics’ variant with the separation-based variant with initial fixed cuts (‘is’), we see significantly faster overall performance and fewer iterations or nodes when using the continuous conic solver and adding certificate cuts. The performance profiles figures 6a to 6d compare the execution times or iteration/node counts for the ‘ics’ and ‘is’ solvers, unambiguously demonstrating superiority of Pajarito’s default ‘ics’ method.

	cuts	statuses				time (s)			subproblems			iters or nodes		
		co	li	er	ex	aco	tco	all	aco	tco	all	aco	tco	all
Iter	c	72	1	21	1	5.59	6.47	7.17	5.31	5.36	4.23	5.48	5.53	4.46
	cs	88	1	3	3	5.56	12.72	14.38	5.31	6.97	6.36	5.48	7.22	6.84
	ics	89	2	0	4	4.73	11.57	14.77	4.17	5.93	6.03	4.32	6.15	6.32
	is	84	1	0	10	8.35	14.53	22.08	-	-	-	13.41	16.52	18.07
MSD	c	76	0	18	1	2.37	3.40	3.50	12.63	15.80	12.70	223	438	348
	cs	88	0	5	2	3.33	6.47	7.76	18.96	26.77	24.87	295	843	815
	ics	92	0	1	2	2.20	6.31	6.52	15.62	24.58	24.95	273	796	857
	is	84	1	0	10	3.12	5.29	7.49	-	-	-	522	932	1345

Table 2: \mathcal{K}^* cut types performance summary.

6.3.2 Extreme Ray Disaggregation

To test the efficacy of the \mathcal{K}^* extreme ray disaggregation technique we describe in section 4.1 and appendix A, we run Pajarito using only certificate cuts (the ‘c’ variant described in section 6.3.1), with and without disaggregation. Note that disabling disaggregation disables use of the second-order cone extended formulation, which has no benefit without disaggregation.

Table 3 summarizes the status counts and shifted geomeans of performance metrics on instance subsets, and the performance profiles figures 7a to 7d compare the execution times or iteration/node counts. For both the iterative and

MSD methods, disaggregation improves performance on nearly every solved instance. For the iterative method, it enables the pure-certificate-based variant to converge on more than double the number of instances, and it more than halves the execution time and iteration count. Without disaggregation, the MSD method manages to converge on many more instances than the iterative method. Disaggregation greatly improves the performance of the MSD method, though the comparison is not quite as striking as for the iterative method.

	disag	statuses				time (s)			subproblems			iters or nodes		
		co	li	er	ex	aco	tco	all	aco	tco	all	aco	tco	all
Iter	off	33	10	52	0	10.95	11.11	17.20	11.63	12.41	11.82	11.95	12.74	12.51
	on	72	1	21	1	4.47	6.47	7.17	4.04	5.36	4.23	4.25	5.53	4.46
MSD	off	51	3	41	0	1.74	6.18	6.71	15.51	50.51	27.98	70	613	261
	on	76	0	18	1	1.06	3.40	3.50	7.57	15.80	12.70	36	438	348

Table 3: \mathcal{K}^* cut disaggregation performance summary.


6.3.3 Certificate-Based Scaling

To test the efficacy of the \mathcal{K}^* certificate cut scaling technique for an LP solver with a feasibility tolerance we describe in section 3.2, we run Pajarito using only certificate cuts (the ‘c’ variant described in section 6.3.1), with and without scaling. We set a larger feasibility tolerance on these four Pajarito solvers ($\delta = 10^{-6}$ instead of 10^{-8} , which we used for all other tests), to reduce the chance that any observed effects are caused by numerical issues near machine epsilon.


Table 4 summarizes the status counts and shifted geomeans of performance metrics on instance subsets, and the performance profiles figures 8a to 8d compare the execution times or iteration/node counts. For both the iterative and MSD methods, using scaling improves the robustness of the pure-certificate-based variant, allowing us to converge on 6 or 7 additional instances. On the subset of instances solved by all four solvers (the ‘aco’ columns), scaling slightly reduces conic subproblem counts and iteration or node counts, but has small and ambiguous effects on the execution times.

	scale	statuses				time (s)			subproblems			iters or nodes		
		co	li	er	ex	aco	tco	all	aco	tco	all	aco	tco	all
Iter	off	63	1	28	3	4.54	4.41	6.59	5.15	5.03	4.23	5.18	5.06	4.40
	on	69	1	22	3	4.35	5.20	6.73	4.90	4.92	3.88	4.99	4.99	4.04
MSD	off	60	0	30	5	2.68	2.77	3.15	12.44	14.48	12.78	193	240	366
	on	67	0	26	2	2.92	4.02	3.86	11.88	15.77	12.07	188	392	393


Table 4: \mathcal{K}^* certificate cut scaling performance summary (larger δ).




(a) Iter - execution time.



(b) MSD - execution time.



(c) Iter - iteration count.



(d) MSD - node count.

Figure 6: \mathcal{K}^* cut types performance profiles.

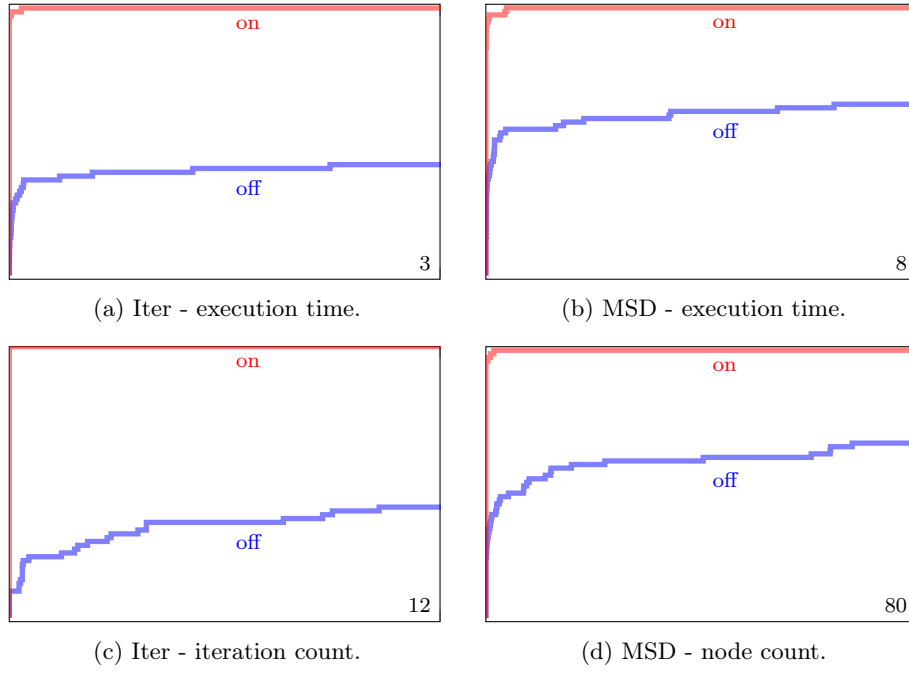


Figure 7: \mathcal{K}^* cut disaggregation performance profiles.

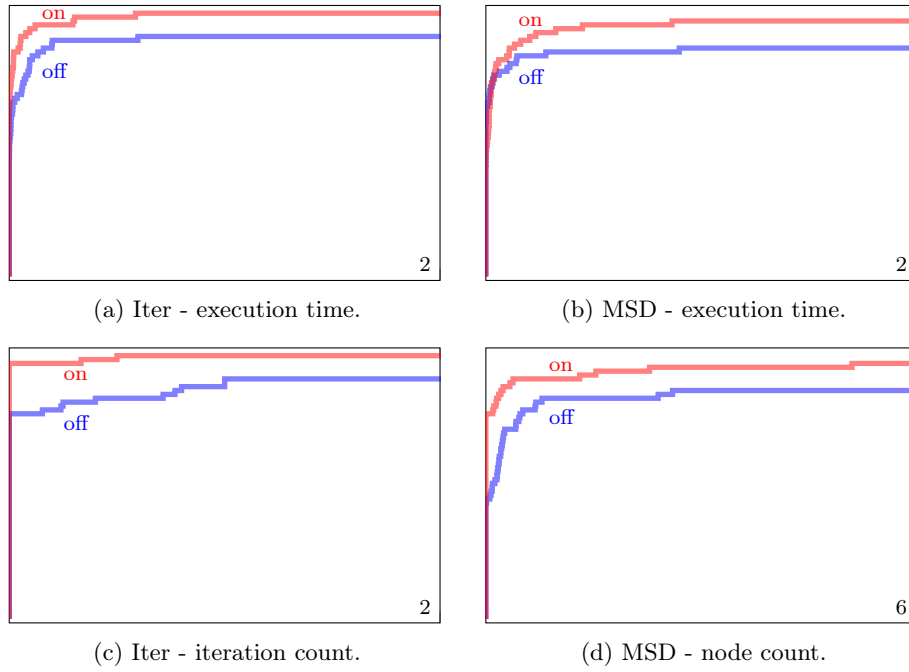


Figure 8: \mathcal{K}^* certificate cut scaling performance profiles.

A The Standard Primitive Nonpolyhedral Cones

As we discussed in section 5.2, Pajarito recognizes three standard primitive nonpolyhedral cones defined by MathProgBase: exponential, second-order, and positive semidefinite cones. Here, we tailor the general techniques for tightening OAs from section 4 to a primitive cone constraint involving one of these three cones. In particular, we describe the initial fixed OAs (see section 4.2), extreme ray disaggregations (see section 4.1), and separation procedures (see section 4.3) implemented in Pajarito.

These ideas could be adapted to other primitive nonpolyhedral cones if the user desires. For example, consider a convex constraint (in NLP form) $f(\mathbf{t}) \leq r$, where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a non-homogeneous convex function. If we define a closed convex cone using the closure of the epigraph of the perspective of f :

$$\mathcal{K}_f = \text{cl}\{(r, s, \mathbf{t}) \in \mathbb{R}^{2+n} : s > 0, r \geq sf(\mathbf{t}/s)\}, \quad (28)$$

then the equivalent conic constraint is $(r, 1, \mathbf{t}) \in \mathcal{K}_f$. From Zhang [2014], the dual cone of \mathcal{K}_f is the closure of the perspective of the epigraph of the convex conjugate of f , $f^*(\mathbf{w}) = \sup\{-\mathbf{w}^T \mathbf{t} - f(\mathbf{t}) : \mathbf{t} \in \text{dom } f\}$:

$$\mathcal{K}_f^* = \text{cl}\{(u, v, \mathbf{w}) \in \mathbb{R}^{2+n} : u > 0, v \geq uf^*(\mathbf{w}/u)\}. \quad (29)$$

Note $\mathcal{K}_f^* \neq \mathcal{K}_{f^*}$ because of the permuting of the first two indices.

A.1 Exponential Cone

The exponential cone \mathcal{E} is defined from the convex univariate exponential function $f(t) = \exp(t)$ in equation (28):

$$\begin{aligned} \mathcal{E} &= \text{cl}\{(r, s, t) \in \mathbb{R}^3 : s > 0, r \geq s \exp(t/s)\} \\ &= \{(r, 0, t) : r \geq 0, t \leq 0\} \cup \{(r, s, t) : s > 0, r \geq s \exp(t/s)\}. \end{aligned} \quad (30)$$

The convex conjugate is $f^*(w) = w - w \log(-w)$, so by equation (29), the dual cone of the exponential cone is:

$$\begin{aligned} \mathcal{E}^* &= \text{cl}\{(u, v, w) \in \mathbb{R}^3 : u > 0, w < 0, v \geq w - w \log(-w/u)\} \\ &= \{(u, v, 0) : u, v \geq 0\} \cup \{(u, v, w) : u > 0, w < 0, v \geq w - w \log(-w/u)\}. \end{aligned} \quad (31)$$

A.1.1 Initial Fixed Polyhedral Relaxation

Suppose we have a primitive cone constraint $(r, s, t) \in \mathcal{E}$. We use the two \mathcal{E}^* extreme rays $(1, 0, 0)$, $(0, 1, 0)$ to impose the simple bound constraints $r, s \geq 0$. We use more \mathcal{E}^* extreme rays of the form $(1, w - w \log(-w), w)$ by picking several different values $w < 0$. Note the corresponding cuts separate any point $(0, 0, t)$ satisfying $t > 0$.

A.1.2 Extreme Ray Disaggregation

Suppose we have the \mathcal{E}^* point (u, v, w) . If $w = 0$, the point is already a non-negative combination of the initial fixed \mathcal{E}^* points from appendix A.1.1, so we discard it. If $w < 0$, we use the \mathcal{E}^* extreme ray $(u, w - w \log(-w/u), w)$, which when added to some nonnegative multiple of $(0, 1, 0)$, gives (u, v, w) .⁵⁰

A.1.3 Separation Of An Infeasible Point

Suppose we want to separate a point $(r, s, t) \notin \mathcal{E}$ that satisfies the initial fixed cuts. Then $r, s \geq 0$ and if $r = s = 0$ then $t \leq 0$. If $s = 0$, then $t > 0$ and $r > 0$, and we use the \mathcal{E}^* extreme ray $(t/r, -2 + 2 \log(2r/t), -2)$. If $s > 0$, then $r < s \exp(t/s)$, and we use the \mathcal{E}^* extreme ray $(1, (t/s - 1) \exp(t/s), -\exp(t/s))$.

A.2 Second-Order Cone

For $n \geq 2$, the second-order cone is the epigraph of the ℓ_2 -norm, which is convex and homogeneous:

$$\mathcal{L}^{1+n} = \{(r, \mathbf{t}) \in \mathbb{R}^{1+n} : r \geq \|\mathbf{t}\|_2\}. \quad (32)$$

We sometimes drop the dimension $1 + n$ when implied by context. This cone is self-dual ($\mathcal{L}^* = \mathcal{L}$). We also define the (self-dual) rotated second-order cone:

$$\mathcal{V}^{2+n} = \{(r, s, \mathbf{t}) \in \mathbb{R}^{2+n} : r, s \geq 0, 2rs \geq \|\mathbf{t}\|_2^2\}. \quad (33)$$

Note that \mathcal{V} is an invertible linear transformation of \mathcal{L} , since $(r, s, \mathbf{t}) \in \mathcal{V}^{2+n}$ if and only if $(r + s, r - s, \sqrt{2}t_1, \dots, \sqrt{2}t_n) \in \mathcal{L}^{2+n}$, so for simplicity we restrict attention to \mathcal{L} .⁵¹

A.2.1 Initial Fixed Polyhedral Relaxation

Suppose we have a primitive cone constraint $(r, \mathbf{t}) \in \mathcal{L}^{1+n}$. First, we note that the ℓ_∞ -norm lower-bounds the ℓ_2 -norm, since for any $\mathbf{t} \in \mathbb{R}^n$ we have:

$$\|\mathbf{t}\|_\infty = \max_{i \in \llbracket n \rrbracket} |t_i| \leq \|\mathbf{t}\|_2. \quad (34)$$

Let $\mathbf{e}(i) \in \mathbb{R}^n$ be the i th unit vector in n dimensions. We use the $2n$ \mathcal{L}^* extreme rays $(1, \pm \mathbf{e}(i)), \forall i \in \llbracket n \rrbracket$, which imply the conditions $r \geq |t_i|, \forall i \in \llbracket n \rrbracket$, equivalent to the homogenized box relaxation $r \geq \|\mathbf{t}\|_\infty$. Second, we note that the ℓ_1 -norm also provides a lower bound for the ℓ_2 -norm, since for any $\mathbf{t} \in \mathbb{R}^n$ we have:

$$\|\mathbf{t}\|_1 = \sum_{i \in \llbracket n \rrbracket} |t_i| \leq \sqrt{n} \|\mathbf{t}\|_2. \quad (35)$$

⁵⁰This also projects $(u, v, w) \notin \mathcal{E}^*$ with $u > 0, w < 0$ onto \mathcal{E}^* .

⁵¹As noted in section 5.2, Pajarito transforms any \mathcal{V} constraints to equivalent \mathcal{L} constraints during preprocessing.

We use the 2^n \mathcal{L}^* extreme rays $(1, \sigma/\sqrt{n}), \forall \sigma \in \{-1, 1\}^n$, which imply the homogenized diamond relaxation $r \geq \|t\|_1/\sqrt{n}$. Although the number of initial fixed cuts is exponential in the dimension n , in appendix B we describe how to use an extended formulation introduced by Vielma et al. [2017] with n auxiliary variables to imply an initial fixed OA that is no weaker but uses only a polynomial number of cuts. Note that the \mathcal{L}^* point $(1, \mathbf{0})$, which corresponds to the simple variable bound $r \geq 0$, is a nontrivial conic combination of these initial fixed \mathcal{L}^* extreme rays.

A.2.2 Extreme Ray Disaggregation

Suppose we have the \mathcal{L}^* point (u, \mathbf{w}) . If $\mathbf{w} = \mathbf{0}$, the point is already a nonnegative multiple of the \mathcal{L}^* point $(1, \mathbf{0})$, so we discard it. Otherwise, we use the \mathcal{L}^* extreme ray $(\|\mathbf{w}\|_2, \mathbf{w})$, which when added to some nonnegative multiple of $(1, \mathbf{0})$, gives the original point (u, \mathbf{w}) .⁵²

A.2.3 Separation Of An Infeasible Point

Suppose we want to separate a point $(r, \mathbf{t}) \notin \mathcal{L}$ that satisfies the initial fixed cuts. Then $r \geq 0$ and so $\mathbf{t} \neq \mathbf{0}$, and we use the \mathcal{L}^* extreme ray $(1, -\mathbf{t}/\|\mathbf{t}\|_2)$.

A.3 Positive Semidefinite Cone

For $n \geq 2$, we define the $n \times n$ -dimensional positive semidefinite (PSD) matrix cone \mathbb{S}_+^n as a subset of the symmetric matrices $\mathbb{S}^n = \{\mathbf{T} \in \mathbb{R}^{n \times n} : \mathbf{T} = \mathbf{T}^T\}$, avoiding the need to enforce symmetry constraints. From the minimum eigenvalue function $\lambda_{\min} : \mathbb{S}^n \rightarrow \mathbb{R}$, we have:

$$\mathbb{S}_+^n = \{\mathbf{T} \in \mathbb{S}^n : \lambda_{\min}(\mathbf{T}) \geq 0\}. \quad (36)$$

For $\mathbf{W}, \mathbf{T} \in \mathbb{S}^n$, we use the trace inner product $\langle \mathbf{W}, \mathbf{T} \rangle = \sum_{i,j \in [n]} W_{i,j} T_{i,j}$. \mathbb{S}_+^n is self-dual and its extreme rays are the rank-1 PSD matrices [Ben-Tal and Nemirovski, 2001a], i.e. any $\omega \omega^T$ for $\omega \in \mathbb{R}^n$. An extreme ray $(\mathbb{S}_+^n)^*$ cut has the form:

$$\langle \omega \omega^T, \mathbf{T} \rangle = \omega^T \mathbf{T} \omega \geq 0. \quad (37)$$

In appendix C, we describe how to strengthen extreme ray \mathbb{S}_+^* cuts to rotated second-order cone constraints (for MISOCP OA; see section 5.4).

Recall that our MI-conic form \mathfrak{M} uses vector cone definitions, as does MathProgBase. Mosek ApS [2016] refers to the matrix cone $\mathbb{S}_+^n \subset \mathbb{S}^n$ as the *smat PSD cone*, and to its equivalent vectorized definition $\mathcal{S}^{n(n+1)/2} \subset \mathbb{R}^{n(n+1)/2}$ as the *svec PSD cone*. In svec space, we use the usual vector inner product, and \mathcal{S} is also self-dual. The invertible linear transformations for an smat-space point

⁵²This also projects $(u, \mathbf{w}) \notin \mathcal{L}^*$ with $u < \|\mathbf{w}\|_2$ onto \mathcal{L}^* .

$\mathbf{T} \in \mathbb{S}^n$ and an svec-space point $\mathbf{t} \in \mathbb{R}^{n(n+1)/2}$ are:

$$\text{svec}(\mathbf{T}) = (T_{1,1}, \sqrt{2}T_{2,1}, \dots, \sqrt{2}T_{n,1}, T_{2,2}, \sqrt{2}T_{3,2}, \dots, T_{n,n}), \quad (38a)$$

$$\text{smat}(\mathbf{t}) = \begin{bmatrix} t_1 & t_1/\sqrt{2} & \dots & t_n/\sqrt{2} \\ t_2/\sqrt{2} & t_{n+1} & \dots & t_{2n-1}/\sqrt{2} \\ \vdots & \vdots & \ddots & \vdots \\ t_n/\sqrt{2} & t_{n-1}/\sqrt{2} & \dots & t_{n(n+1)/2} \end{bmatrix}. \quad (38b)$$

A.3.1 Initial Fixed Polyhedral Relaxation

Suppose we have a primitive cone constraint $\mathbf{T} \in \mathbb{S}_+^n$. Let $\mathbf{e}(i) \in \mathbb{R}^n$ be the i th unit vector in n dimensions. For each $i \in \llbracket n \rrbracket$, we let $\boldsymbol{\omega} = \mathbf{e}(i)$ in the extreme ray \mathbb{S}_+^* cut (37), which imposes the diagonal nonnegativity condition $T_{i,i} \geq 0$ necessary for PSDness. For each $i, j \in \llbracket n \rrbracket : i > j$, we let $\boldsymbol{\omega} = \mathbf{e}(i) \pm \mathbf{e}(j)$ in (37), which enforces the condition $T_{i,i} + T_{j,j} \geq 2|T_{i,j}|$ necessary for PSDness. Ahmadi and Hall [2015] discuss an LP inner approximation of the PSD cone called the cone of diagonally dominant (DD) matrices. Our initial fixed \mathbb{S}_+^* points $\boldsymbol{\omega}\boldsymbol{\omega}^T$ are exactly the extreme rays of the DD cone, so our initial fixed OA is the dual cone of the DD cone.

A.3.2 Extreme Ray Disaggregation

Suppose we have the \mathbb{S}_+^* point \mathbf{W} , not necessarily an extreme ray of \mathbb{S}_+^* . We perform an eigendecomposition $\mathbf{W} = \sum_{i \in \llbracket n \rrbracket} \lambda_i \tilde{\boldsymbol{\omega}}_i \tilde{\boldsymbol{\omega}}_i^T$, where for all $i \in \llbracket n \rrbracket$, λ_i is the i th eigenvalue and $\tilde{\boldsymbol{\omega}}_i$ is its corresponding eigenvector.⁵³ Since \mathbf{W} is PSD, every eigenvalue is nonnegative, and there are $\text{rank}(\mathbf{W}) \leq n$ positive eigenvalues. For each $i \in \llbracket n \rrbracket : \lambda_i > 0$, we let $\boldsymbol{\omega} = \sqrt{\lambda_i} \tilde{\boldsymbol{\omega}}_i$ in (37).⁵⁴ These extreme ray \mathbb{S}_+^* cuts aggregate to imply the original \mathbb{S}_+^* cut $\langle \mathbf{W}, \mathbf{T} \rangle \geq 0$.

A.3.3 Separation Of An Infeasible Point

Suppose we want to separate a point $\mathbf{T} \in \mathbb{S}^n \setminus \mathbb{S}_+^n$. We perform an eigendecomposition $\mathbf{T} = \sum_{i \in \llbracket n \rrbracket} \lambda_i \boldsymbol{\tau}_i \boldsymbol{\tau}_i^T$, for which at least one eigenvalue is negative. For each $i \in \llbracket n \rrbracket : \lambda_i < 0$, we let $\boldsymbol{\omega} = \boldsymbol{\tau}_i$ in (37) (note $\langle \boldsymbol{\tau}_i \boldsymbol{\tau}_i^T, \mathbf{T} \rangle = \boldsymbol{\tau}_i^T \mathbf{T} \boldsymbol{\tau}_i = \lambda_i < 0$).

B The Second-Order Cone Extended Formulation

Recall the definitions of the second-order cone \mathcal{L} and the rotated-second-order cone \mathcal{V} in appendix A.2. Both \mathcal{L} and \mathcal{V} are self-dual. As discussed in section 5.4, Pajarito can optionally use an *extended formulation* (EF) for second-order cone constraints, leading to tighter polyhedral relaxations. Vielma et al. [2017] show

⁵³Note that for real symmetric matrices, all eigenvalues are real. We select the eigenvectors to be orthonormal.

⁵⁴By dropping any $i \in \llbracket n \rrbracket : \lambda_i < 0$, this projects $\mathbf{W} \in \mathbb{S} \setminus \mathbb{S}_+^*$ onto \mathbb{S}_+^* .

that the constraint $(r, \mathbf{t}) \in \mathcal{L}^{1+n}$ is equivalent to the following $1 + n$ constraints on r , \mathbf{t} , and the auxiliary variables $\boldsymbol{\pi} \in \mathbb{R}^n$:

$$\sum_{i \in [n]} 2\pi_i \leq r \quad (39a)$$

$$(r, \pi_i, t_i) \in \mathcal{V}^3 \quad \forall i \in [n]. \quad (39b)$$

By projecting out the $\boldsymbol{\pi}$ variables, the equivalence is obvious. Constraints (39b) imply $r \geq 0$ and $\pi_i \geq 0$ and $2r\pi_i \geq t_i^2$ for all $i \in [n]$. Aggregating the latter conditions and using the linear inequality (39a), we see $r^2 \geq \sum_{i \in [n]} 2r\pi_i \geq \sum_{i \in [n]} t_i^2$, which is equivalent to the original constraint (for $r \geq 0$). We only use \mathcal{L}^{1+n} and \mathcal{V}^3 here, so for convenience we drop the dimensions.⁵⁵

Suppose (u, \mathbf{w}) is a \mathcal{L}^* extreme ray, so from appendix A.2.2, we have $\mathbf{w} \neq \mathbf{0}$ and $u = \|\mathbf{w}\|_2 > 0$. Then $ur + \mathbf{w}^T \mathbf{t} \geq 0$ is a \mathcal{L}^* cut. Note that the linear constraint (39a) in the EF implies:

$$ur + \mathbf{w}^T \mathbf{t} \geq \frac{ur}{2} + u \sum_{i \in [n]} \pi_i + \mathbf{w}^T \mathbf{t} = \sum_{i \in [n]} \left(\frac{w_i^2 r}{2u} + u\pi_i + w_i t_i \right). \quad (40)$$

For each $i \in [n]$, consider the \mathcal{V}^* extreme ray $(w_i^2/2u, u, w_i)$, which implies a \mathcal{V}^* cut for the i th constraint (39b) in the EF. The RHS of (40) is an aggregation of these n \mathcal{V}^* cuts, which means the \mathcal{V}^* cuts imply the \mathcal{L}^* cut condition $ur + \mathbf{w}^T \mathbf{t} \geq 0$. Therefore, there is no loss of strength in the polyhedral relaxations, and we maintain the certificate \mathcal{K}^* cut guarantees from section 3.1.⁵⁶

We now apply this lifting procedure to the initial fixed \mathcal{L}^* points described in appendix A.2.1. The \mathcal{L}^* points for the ℓ_∞ -norm relaxation are $(1, \pm \mathbf{e}(i))$, $\forall i \in [n]$; for each $i \in [n]$, we get three unique \mathcal{V}^* extreme rays $(0, 1, 0)$ (for $w_i = 0$) and $(1/2, 1, \pm 1)$ (for $w_i = \pm 1$). The \mathcal{L}^* points for the ℓ_1 -norm relaxation are $(1, \boldsymbol{\sigma}/\sqrt{n})$, $\forall \boldsymbol{\sigma} \in \{-1, 1\}^n$; for each $i \in [n]$, we get two unique \mathcal{V}^* extreme rays $(1/2n, 1, \pm 1/\sqrt{n})$ (for $w_i = \pm 1/\sqrt{n}$). The polyhedral relaxation implied by these $5n$ \mathcal{V}^* points in the EF (39a) and (39b) is at least as strong as that implied by the $2n + 2^n$ \mathcal{L}^* points from appendix A.2.1, so our initial fixed OA can be imposed much more economically with the EF.

C SOCP Outer Approximation For PSD Cones

Recall the definitions of the self-dual smat-space PSD cone \mathbb{S}_+ in appendix A.3 and the self-dual rotated-second-order cone \mathcal{V} in appendix A.2.⁵⁷ For a primitive cone constraint $\mathbf{T} \in \mathbb{S}_+^n$, we demonstrate how to strengthen an $(\mathbb{S}_+^n)^*$ extreme

⁵⁵Ben-Tal and Nemirovski [2001b] introduced an alternative extended formulation for the \mathcal{L}^{1+n} . See Vielma et al. [2017] for a discussion and computational comparison of various \mathcal{L}^{1+n} extended formulations in the context of a separation-based B&B-OA algorithm for MISOCP.

⁵⁶Without the ability to rescale the linear constraint (39a), we cannot recover the guarantees under an LP solver with a feasibility tolerance from section 3.2. However, Pajarito heuristically scales up each \mathcal{V}^* point by a factor n .

⁵⁷ \mathcal{V}^3 is in fact a simple linear transformation of \mathbb{S}_+^2 .

ray cut $\langle \omega \omega^T, \mathbf{T} \rangle \geq 0$ to up to n different \mathcal{V}^3 constraints. As discussed in section 5.4, Pajarito can optionally solve an MISOCP OA model including these \mathcal{V}^3 constraints, leading to tighter relaxations of a challenging \mathbb{S}_+^n constraint.

Fix the index $i \in \llbracket n \rrbracket$. Let $\omega = \omega_i$ be the i th element of ω , and $\underline{\omega} = (\omega_j)_{j \in \llbracket n \rrbracket \setminus \{i\}} \in \mathbb{R}^{n-1}$ be the (column) subvector of ω with the i th element removed. Similarly, let $\underline{t} = T_{i,i}$ and $\underline{\mathbf{t}} = (T_{i,j})_{j \in \llbracket n \rrbracket \setminus \{i\}} \in \mathbb{R}^{n-1}$, and let $\underline{\mathbf{T}} = (T_{k,j})_{k,j \in \llbracket n \rrbracket \setminus \{i\}} \in \mathbb{S}^{n-1}$ be the submatrix of \mathbf{T} with the i th column and row removed. Kim et al. [2003] prove a variant of the standard Schur-complement result that $\mathbf{T} \in \mathbb{S}_+^n$ if and only if $\underline{t} \geq 0$ and:

$$\underline{\mathbf{T}} \in \mathbb{S}_+^{n-1} \quad (41a)$$

$$\underline{t}\underline{\mathbf{T}} - \underline{\mathbf{t}}\underline{\mathbf{t}}^T \in \mathbb{S}_+^{n-1}. \quad (41b)$$

Consider the 3-dimensional rotated second-order cone constraint:

$$(\underline{t}, \underline{\omega}^T \underline{\mathbf{T}} \underline{\omega}, \sqrt{2} \underline{\omega}^T \underline{\mathbf{t}}) \in \mathcal{V}^3. \quad (42)$$

By the definition of \mathcal{V} , constraint (42) is equivalent to the conditions $\underline{t} \geq 0$ and:

$$\underline{\omega}^T \underline{\mathbf{T}} \underline{\omega} \geq 0 \quad (43a)$$

$$\underline{\omega}^T (\underline{t}\underline{\mathbf{T}}) \underline{\omega} \geq (\underline{\omega}^T \underline{\mathbf{t}})^2. \quad (43b)$$

Condition (41a) implies condition (43a), by the dual cone definition (5). Since $(\underline{\omega}^T \underline{\mathbf{t}})^2 = \underline{\omega}^T \underline{\mathbf{t}} \underline{\mathbf{t}}^T \underline{\omega}$, condition (43b) is equivalent to $\underline{\omega}^T (\underline{t}\underline{\mathbf{T}} - \underline{\mathbf{t}}\underline{\mathbf{t}}^T) \underline{\omega} \geq 0$. Thus, by the dual cone definition again, condition (41b) implies condition (43b). Therefore, constraint (42) is a valid relaxation of the PSD constraint $\mathbf{T} \in \mathbb{S}_+^n$. Furthermore, from Theorem 3.3 of Kim et al. [2003], constraint (42) holds if and only if:

$$\langle \mathbf{W}, \mathbf{T} \rangle \geq 0 \quad \forall \mathbf{W} \in \mathbb{S}_+^n : (W_{k,j} = \omega_k \omega_j, \forall k, j \in \llbracket n \rrbracket \setminus \{i\}). \quad (44)$$

Thus constraint (42) potentially implies an infinite family of $(\mathbb{S}_+^n)^*$ cuts, including the original $(\mathbb{S}_+^n)^*$ cut $\langle \omega \omega^T, \mathbf{T} \rangle \geq 0$.⁵⁸ Note that the choice of $i \in \llbracket n \rrbracket$ is arbitrary, so we can derive n different \mathcal{V} constraints of the form (42).⁵⁹

We now apply this strengthening procedure to the initial fixed $(\mathbb{S}_+^n)^*$ extreme rays described in appendix A.3.1. Letting $\omega = e(i) \pm e(j)$ for each $i, j \in \llbracket n \rrbracket : j > i$ in constraint (42), we get the $n(n-1)/2$ initial fixed \mathcal{V} constraints:

$$(T_{i,i}, T_{j,j}, \sqrt{2} T_{i,j}) \in \mathcal{V}^3 \quad \forall i, j \in \llbracket n \rrbracket : j > i. \quad (45)$$

These constraints enforce that every 2×2 principal matrix of \mathbf{T} is PSD, a necessary but insufficient condition for $\mathbf{T} \in \mathbb{S}_+^n$. Ahmadi and Hall [2015] discuss an SOCP inner approximation of the PSD cone called the cone of scaled diagonally dominant (SDD) matrices. Our initial fixed SOCP OA is the dual SDD matrix cone, a strict subset of the polyhedral dual DD matrix cone that corresponds to our initial fixed LP OA from appendix A.3.1.

⁵⁸We do not explore how to scale these \mathcal{V} constraints to recover the guarantees from section 3.2 for an SOCP solver with an absolute feasibility tolerance.

⁵⁹For strengthening separation or certificate \mathcal{K}^* cuts, Pajarito heuristically picks one of the n possible \mathcal{V} constraints by choosing i as the coordinate of the largest absolute value in ω .

Bibliography

- T. Achterberg. SCIP: Solving constraint integer programs. *Mathematical Programming Computation*, 1(1):1–41, 2009.
- A. A. Ahmadi and G. Hall. Sum of squares basis pursuit with linear and second order cone programming. *arXiv preprint arXiv:1510.01597*, 2015.
- P. Belotti, C. Kirches, S. Leyffer, J. Linderoth, J. Luedtke, and A. Mahajan. Mixed-integer nonlinear optimization. *Acta Numerica*, 22:1–131, 2013. ISSN 1474-0508.
- A. Ben-Tal and A. Nemirovski. *Lectures on Modern Convex Optimization*. Society for Industrial and Applied Mathematics, 2001a.
- A. Ben-Tal and A. Nemirovski. On polyhedral approximations of the second-order cone. *Mathematics of Operations Research*, 26(2):193–205, 2001b.
- J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah. Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1):65–98, 2017.
- P. Bonami, L. T. Biegler, A. R. Conn, G. Cornuéjols, I. E. Grossmann, C. D. Laird, J. Lee, A. Lodi, F. Margot, N. Sawaya, and A. Wächter. An algorithmic framework for convex mixed integer nonlinear programs. *Discrete Optimization*, 5(2):186–204, 2008. ISSN 1572-5286.
- P. Bonami, M. Kılınç, and J. Linderoth. Algorithms and software for convex mixed integer nonlinear programs. In J. Lee and S. Leyffer, editors, *Mixed Integer Nonlinear Programming*, volume 154 of *The IMA Volumes in Mathematics and its Applications*, pages 1–39. Springer New York, 2012. ISBN 978-1-4614-1926-6.
- B. Borchers. Csdp, ac library for semidefinite programming. *Optimization methods and Software*, 11(1-4):613–623, 1999.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. ISBN 0521833787.
- S. Diamond and S. Boyd. CVXPY: A python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
- E. D. Dolan and J. J. Moré. Benchmarking optimization software with performance profiles. *Mathematical Programming*, 91(2):201–213, 2002. ISSN 0025-5610.
- A. Domahidi, E. Chu, and S. Boyd. ECOS: An SOCP solver for embedded systems. In *European Control Conference (ECC)*, pages 3071–3076, 2013.

- S. Drewes and S. Ulbrich. Subgradient based outer approximation for mixed integer second order cone programming. In J. Lee and S. Leyffer, editors, *Mixed Integer Nonlinear Programming*, volume 154 of *The IMA Volumes in Mathematics and its Applications*, pages 41–59. Springer New York, 2012. ISBN 978-1-4614-1926-6. doi: 10.1007/978-1-4614-1927-3_2.
- I. Dunning, J. Huchette, and M. Lubin. JuMP: A modeling language for mathematical optimization. *SIAM Review*, 59(2):295–320, 2017. doi: 10.1137/15M1020575.
- H. A. Friberg. CBLIB 2014: a benchmark library for conic mixed-integer and continuous optimization. *Mathematical Programming Computation*, 8(2): 191–214, 2016. ISSN 1867-2957.
- T. Gally, M. E. Pfetsch, and S. Ulbrich. A framework for solving mixed-integer semidefinite programs. *Optimization Methods and Software*, 33(3):594–632, 2018. doi: 10.1080/10556788.2017.1322081.
- N. Gould and J. Scott. A note on performance profiles for benchmarking software. *ACM Transactions on Mathematical Software*, 43(2), 2016. ISSN 0098-3500.
- M. Grant and S. Boyd. Cvx: Matlab software for disciplined convex programming, version 2.1. cvxr.com/cvx, 2014.
- M. Grant, S. Boyd, and Y. Ye. Disciplined convex programming. In L. Liberti and N. Maculan, editors, *Global Optimization*, volume 84 of *Nonconvex Optimization and Its Applications*, pages 155–210. Springer US, 2006. ISBN 978-0-387-28260-2.
- O. Günlük and J. Linderoth. Perspective reformulation and applications. In J. Lee and S. Leyffer, editors, *Mixed Integer Nonlinear Programming*, volume 154 of *The IMA Volumes in Mathematics and its Applications*, pages 61–89. Springer New York, 2012. ISBN 978-1-4614-1926-6.
- S. Kim, M. Kojima, and M. Yamashita. Second order cone programming relaxation of a positive semidefinite constraint. *Optimization Methods and Software*, 18(5):535–541, 2003.
- O. Kröger, C. Coffrin, H. Hijazi, and H. Nagarajan. Juniper: An open-source nonlinear branch-and-bound solver in julia, 2018.
- S. Leyffer. *Deterministic Methods for Mixed Integer Nonlinear Programming*. PhD thesis, University of Dundee, 12 1993.
- M. Lubin and I. Dunning. Computing in operations research using Julia. *INFORMS Journal on Computing*, 27(2):238–248, 2015.

- M. Lubin, E. Yamangil, R. Bent, and J. P. Vielma. Extended formulations in mixed-integer convex programming. In Q. Louveaux and M. Skutella, editors, *Integer Programming and Combinatorial Optimization: 18th International Conference, IPCO 2016, Liège, Belgium, June 1-3, 2016, Proceedings*, pages 102–113. Springer International Publishing, 2016. ISBN 978-3-319-33461-5. doi: 10.1007/978-3-319-33461-5_9.
- M. Lubin, E. Yamangil, R. Bent, and J. P. Vielma. Polyhedral approximation in mixed-integer convex optimization. *Mathematical Programming*, Sep 2017a. ISSN 1436-4646. doi: 10.1007/s10107-017-1191-y. URL <https://doi.org/10.1007/s10107-017-1191-y>.
- M. Lubin, I. Zadik, and J. P. Vielma. Mixed-integer convex representability. In F. Eisenbrand and J. Könnemann, editors, *Integer Programming and Combinatorial Optimization - 19th International Conference, IPCO 2017, Waterloo, ON, Canada, June 26-28, 2017, Proceedings*, volume 10328 of *Lecture Notes in Computer Science*, pages 392–404. Springer, 2017b. ISBN 978-3-319-59249-7. doi: 10.1007/978-3-319-59250-3_32. URL https://doi.org/10.1007/978-3-319-59250-3_32.
- M. Lubin, I. Zadik, and J. P. Vielma. Regularity in mixed-integer convex representability. *arXiv preprint arXiv:1706.05135*, 2017c.
- Mosek ApS. Modeling Cookbook revision 2.0.1. docs.mosek.com/MOSEKModeling_Cookbook-letter.pdf accessed 2017-04-12, 2016.
- H. Nagarajan, M. Lu, S. Wang, R. Bent, and K. Sundar. An adaptive, multivariate partitioning algorithm for global optimization of nonconvex programs. *arXiv preprint arXiv:1707.02514*, 2017.
- B. O’Donoghue, E. Chu, N. Parikh, and S. Boyd. Conic optimization via operator splitting and homogeneous self-dual embedding. *Journal of Optimization Theory and Applications*, 169(3):1042–1068, 2016. ISSN 1573-2878. doi: 10.1007/s10957-016-0892-3. URL <http://dx.doi.org/10.1007/s10957-016-0892-3>.
- P. A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical programming*, 96(2):293–320, 2003.
- F. Permenter, H. A. Friberg, and E. D. Andersen. Solving conic optimization problems via self-dual embedding and facial reduction: a unified approach. *Optimization Online*, September, 2015.
- I. Quesada and I. Grossmann. An LP/NLP based branch and bound algorithm for convex MINLP optimization problems. *Computers & Chemical Engineering*, 16(10):937–947, 1992. ISSN 0098-1354. doi: 10.1016/0098-1354(92)80028-8. URL www.sciencedirect.com/science/article/pii/0098135492800288.

- M. Saltzman, L. Ladányi, and T. Ralphs. The COIN-OR Open Solver Interface: Technology overview. Presented at CORS/INFORMS Banff, 5 2004. URL <https://www.coin-or.org/Presentations/CORS2004-OSI.pdf>.
- S. A. Serrano. *Algorithms for unsymmetric cone optimization and an implementation for problems with the exponential cone*. PhD thesis, Stanford University, 2015.
- A. Skajaa, E. D. Andersen, and Y. Ye. Warmstarting the homogeneous and self-dual interior point method for linear and conic quadratic problems. *Mathematical Programming Computation*, pages 1–25, 2013.
- M. Udell, K. Mohan, D. Zeng, J. Hong, S. Diamond, and S. Boyd. Convex optimization in Julia. In *Proceedings of HPTCDL '14*, pages 18–28, Piscataway, NJ, USA, 2014. IEEE Press. ISBN 978-1-4799-7020-9.
- J. P. Vielma. Small and strong formulations for unions of convex sets from the cayley embedding. *Mathematical Programming*, Mar 2018. ISSN 1436-4646. doi: 10.1007/s10107-018-1258-4. URL <https://doi.org/10.1007/s10107-018-1258-4>.
- J. P. Vielma, I. Dunning, J. Huchette, and M. Lubin. Extended formulations in mixed integer conic quadratic programming. *Mathematical Programming Computation*, 9(3):369–418, Sep 2017. ISSN 1867-2957. doi: 10.1007/s12532-016-0113-y. URL <https://doi.org/10.1007/s12532-016-0113-y>.
- S. Vigerske. MINLPLIB2 library. www.gamsworld.org/minlp/minlplib2/html/ accessed 2016-05-13, 2018.
- J. Witzig, T. Berthold, and S. Heinz. Experiments with conflict analysis in mixed integer programming. In *International Conference on AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, pages 211–220. Springer, 2017.
- S. Zhang. Lecture 4: The dual cone and dual problem. University Lecture, 2014. URL <http://www.isye.umn.edu/courses/ie8534/pdf/Lecture-4.pdf>.