A Simple Effective Heuristic for Embedded Mixed-Integer Quadratic Programming

Reza Takapoui   Nicholas Moehle   Stephen Boyd

Alberto Bemporad

September 29, 2015

Abstract

In this paper we propose a fast optimization algorithm for approximately minimizing convex quadratic functions over the intersection of affine and separable constraints (*i.e.*, the Cartesian product of possibly nonconvex real sets). This problem class contains many NP-hard problems such as mixed-integer quadratic programming. Our heuristic is based on a variation of the alternating direction method of multipliers (ADMM), an algorithm for solving convex optimization problems. We discuss the favorable computational aspects of our algorithm, which allow it to run quickly even on very modest computational platforms such as embedded processors. We give several examples for which an approximate solution should be found very quickly, such as management of a hybrid-electric vehicle drivetrain and control of switched-mode power converters. Our numerical experiments suggest that our method is very effective in finding a feasible point with small objective value; indeed, we find that in many cases, it finds the global solution.

1   Introduction

1.1   The problem

We consider the problem

$$
\begin{array}{ll}
\text{minimize} & (1/2)x^T P x + q^T x + r \\
\text{subject to} & Ax = b \\
& x \in \mathcal{X}
\end{array}
\tag{1}
$$

with decision variable $x \in \mathbf{R}^n$. The problem parameters are the symmetric positive semidefinite matrix $P \in \mathbf{R}^{n \times n}$, the matrix $A \in \mathbf{R}^{m \times n}$, the vectors $b \in \mathbf{R}^m$ and $q \in \mathbf{R}^n$, and the real number $r \in \mathbf{R}$. The constraint set $\mathcal{X}$ is the Cartesian product of (possibly nonconvex) real, closed, nonempty sets, *i.e.*, $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_n$, where $\mathcal{X}_i \subseteq \mathbf{R}$ are closed, nonempty subsets of $\mathbf{R}$ for $i = 1, \dots, n$. If $\mathcal{X}_i$ is a convex set, we refer to variable x_i as a *convex variable*, and if $\mathcal{X}_i$ is a nonconvex set, we call variable x_i a *nonconvex variable*.

Some applications. Many problems can be put into the form of problem (1). For example, if some of the sets $\mathcal{X}_i$ are subsets of integers, our formulation addresses mixed-integer quadratic and mixed-integer linear programs. This includes applications such as admission control [OCP07], economic dispatch [PF07], scheduling [CPM10], hybrid vehicle control [MJSB12], thermal unit commitment problems [CA06], Boolean satisfiability problems (SAT) [JW90], and hybrid model predictive control [BM99]. Another application is embedded signal decoding in communication systems, when the nonconvex sets are signal constellations (*e.g.*, QAM constellations; see [GG10, pg. 416]).

Complexity. If $\mathcal{X}$ is a convex set, problem (1) is a convex optimization problem and can be readily solved using standard convex optimization techniques. Otherwise, the problem (1) can be hard in general. It trivially generalizes mixed-integer quadratic programming, an NP-complete problem, and can therefore be used to encode other NP-complete problems such as the traveling salesman problem (TSP) [PS98], Boolean satisfiability (SAT) [LZD04, Kar72], set cover [Hoc82], and set packing [Pad73]. Hence, any algorithm that guarantees finding the global solution to (1) suffers from non-polynomial worst-case time (unless P = NP).

1.2 Solve techniques

Exact methods. There are a variety of methods for solving (1) exactly. When all of the nonconvex sets $\mathcal{X}_i$ in (1) are finite, the simplest method is brute force; enumerating through all possible combinations of discrete variables and solve a convex optimization problem for each possible combination and finding the point with the smallest objective value. Other methods such as branch-and-bound [LW66] and branch-and-cut [SM99] are guaranteed to find the global solution. Cutting plane methods [G$^+$58, CCH89] rely on solving the relaxation and adding a linear constraint to drive the solution towards being integer. Special purpose methods have been introduced for some specific subclasses of (1). Unfortunately, these methods have non-polynomial worst-case runtime, and are often burdensome to use in practice, especially for embedded optimization, where runtime, memory limits, and code simplicity are prioritized. Also, these methods suffer from a large variance in the algorithm runtime.

Heuristics. On the other hand, many heuristics have been introduced that can deliver a good, but suboptimal (and possibly infeasible) point in a very short amount of time. For example, the *relax-and-round* heuristic consists of replacing each $\mathcal{X}_i$ by its convex hull, solving the resulting relaxation (a convex quadratic program), and projecting the solution onto the nonconvex constraint sets. Another heuristic is to fix the nonconvex variables for several reasonable guesses and solve the convex optimization problem for convex variables. (Each of these problems may not find a feasible point, even if one exists.) The *feasibility pump* is a heuristic to find a feasible solution to a generic mixed integer program and is discussed in [FGL05, BFL07, AB07]. Such heuristics are often quite effective, and can be implemented on very modest computational hardware, making them very attractive for

embedded applications (even without any theoretical guarantees).

1.3   Embedded applications

We focus on embedded applications where finding a feasible point with relatively small objective will often result in performance that is practically indistinguishable from implementing the global solution. In embedded applications, the computational resources are limited and a solution must be found in a small time. Hence, methods to find the global solution are not favorable, because their large variance in runtime cannot be tolerated.

In an embedded application, it is often required to solve several instances of (1), with different values of the parameters. Here we distinguish two separate use cases, depending on whether one or both of P or A change. This distinction will play an important role in solution methods. In the first use case, we solve many instances of (1) in which any of the parameters may change between instances. In the second use case, we solve instances of (1) in which q, b, and $\mathcal{X}$ change between instances, but P and A are constant. Although this is more restrictive than the first use case, many applications can be well modeled using this approach, including linear, time-invariant model predictive control and moving horizon estimation. Indeed, all of the three examples we present in §3 are of this type.

1.4   Contributions

Our proposed algorithm is a simple and computationally efficient heuristic to find approximate solutions to problem (1) quickly. It is based on the alternating direction method of multipliers (ADMM), an algorithm for solving convex optimization problems. Because the problem class we address includes nonconvex optimization problems, our method is not guaranteed to find the global solution, or even converge.

Numerical experiments suggest that this heuristic is an effective tool to find the global solution in a variety of problem instances. Even if our method does not find the global solution, it usually finds a feasible point with reasonable objective value. This makes it effective for many embedded optimization applications, where finding a feasible point with relatively small objective value often results in performance that is practically indistinguishable from implementing the global solution.

Comparison of the runtime with commercial solvers such as MOSEK [ApS15] and CPLEX [CPL09] show that our method can be substantially faster than solving a global optimization method, while having a competitive practical performance.

1.5   Related work

Fast embedded optimization.   In recent years, much research has been devoted to solving moderately-sized convex optimization problems quickly (*i.e.*, in milliseconds or microseconds), possibly on embedded platforms. Examples include the SOCP solvers ECOS [DCB13], and FiordOs [Ull11], and the QP solver CVXGEN [MB12]. Other algorithms have been developed exclusively for convex optimal control problems; see [WB10, OSB13, JGR$^+$14]. In

3

addition, recent advances in automatic code generation for convex optimization [MWB11, CPDB13] can significantly reduce the cost and complexity of using an embedded solver. Some recent effort has been devoted to (globally) solving mixed-integer convex programs very quickly; see [Bem15], [FDM15] and references therein.

Nonconvex ADMM. Even though ADMM was originally introduced as a tool for convex optimization problems, it turns out to be a powerful heuristic method even for NP-hard nonconvex problems [BPC$^+$11, §5, 9]. Recently, this tool has been used as a heuristic to find approximate solutions to nonconvex problems [CW13, Cha12]. In [DBEY13], the authors study the *Divide and Concur* algorithm as a special case of a message-passing version of the ADMM, and introduce a three weight version of this algorithm which greatly improves the performance for some nonconvex problems such as circle packing and the Sudoku puzzle.

2   Our heuristic

2.1   Algorithm

Our proposed algorithm is an extension of the alternating direction method of multipliers (ADMM) for constrained optimization to the nonconvex setting [BPC$^+$11, §5,9]. ADMM was originally introduced for solving convex problems, but practical evidence suggests that it can be an effective method to approximately solve some nonconvex problems as well. In order to use ADMM, we rewrite problem (1) as

$$
\begin{aligned}
\text{minimize} \quad & (1/2)x^T P x + q^T x + I_{\mathcal{X}}(z) \\
\text{subject to} \quad & \begin{bmatrix} A \\ I \end{bmatrix} x - \begin{bmatrix} 0 \\ I \end{bmatrix} z = \begin{bmatrix} b \\ 0 \end{bmatrix}.
\end{aligned}
\tag{2}
$$

Here $I_{\mathcal{X}}$ denotes the indicator function of $\mathcal{X}$, so that $I_{\mathcal{X}}(x) = 0$ for $x \in \mathcal{X}$ and $I_{\mathcal{X}}(x) = \infty$ for $x \notin \mathcal{X}$. Each iteration in the algorithm consists of the following three steps:

$$
\begin{aligned}
x^{k+1/2} \quad &:= \quad \underset{x}{\operatorname{argmin}} \left((1/2)x^T P x + q^T x + (\rho/2) \left\| \begin{bmatrix} A \\ I \end{bmatrix} x - \begin{bmatrix} 0 \\ I \end{bmatrix} x^k - \begin{bmatrix} b \\ 0 \end{bmatrix} + u^k \right\|_2^2 \right) \\
x^{k+1} \quad &:= \quad \Pi \left(x^{k+1/2} + \begin{bmatrix} 0 & I \end{bmatrix} u^k \right) \\
u^{k+1} \quad &:= \quad u^k + \begin{bmatrix} A \\ I \end{bmatrix} x^{k+1/2} - \begin{bmatrix} 0 \\ I \end{bmatrix} x^k - \begin{bmatrix} b \\ 0 \end{bmatrix}.
\end{aligned}
$$

Here, Π denotes the projection onto $\mathcal{X}$. Note that if $\mathcal{X}$ is not convex, the projection onto $\mathcal{X}$ may not be unique; for our purposes, we only need that $\Pi(z) \in \operatorname{argmin}_{x \in \mathcal{X}} \|x - z\|_2$ for all $z \in \mathbf{R}^n$. Since $\mathcal{X}$ is the Cartesian product of subsets of the real line, *i.e.*, $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_n$, we can take $\Pi(z) = \Pi_1(z_1) \times \cdots \times \Pi_n(z_n)$, where Π_i is a projection function onto $\mathcal{X}_i$. Usually evaluating $\Pi_i(z)$ is inexpensive; for example, if $\mathcal{X}_i = [a, b]$ is an interval, $\Pi_i(z) = \min\{\max\{z, a\}, b\}$. If $\mathcal{X}_i$ is the set of integers, Π_i rounds its argument to the nearest integer. For any finite set $\mathcal{X}_i$ with k elements, $\Pi_i(z)$ is a closest point to z that belongs to $\mathcal{X}_i$, which can be found by $\lceil \log_2 k \rceil$ comparisons.

2.2 Convergence

If the set $\mathcal{X}$ is convex and problem (1) is feasible, the algorithm is guaranteed to converge to an optimal point [BPC$^+$11, §3]. However, for $\mathcal{X}$ nonconvex, there is no such guarantee. Indeed, because problem (1) can be NP-hard, any algorithm that finds the global solution suffers from nonpolynomial worst-case runtime. Our approach is to give up the accuracy and use methods that find an approximate solution in a small time.

Our numerical results verify that even for simple examples, the algorithm may fail to converge, converge to a suboptimal point, or fail to find a feasible point, even if one exists. Since the objective value need not decrease monotonically (or at all), it is critical to keep track of the best point found runtime. That is, for a selected primal feasibility tolerance ϵ^{tol}, we shall reject all points x such that $\|Ax - b\| > \epsilon^{\text{tol}}$, and among those primal feasible points x that $\|Ax - b\| \leq \epsilon^{\text{tol}}$, we choose the point with the smallest objective value. Here, ϵ^{tol} is a tolerance for accepted feasibility. We should remind the reader again, that this point need not be the global minimum.

2.3 Initialization

To initialize x^0, one can randomly choose a point in $\mathbf{Co}\,\mathcal{X}$, where $\mathbf{Co}\,\mathcal{X}$ denotes the convex hull of $\mathcal{X}$. More specifically, this means that we need to have access to a subroutine that generates random points in $\mathbf{Co}\,\mathcal{X}$. Our numerical results show that running the algorithm multiple times with different random initializations increases the chance of finding a feasible point with smaller objective value. Hence, we suggest running the algorithm multiple times initialized with random starting points and report the best point as the approximate solution. We always initialize $u^0 = 0$.

2.4 Computational cost

In this subsection, we make a few comments about the computational cost of each iteration. The first step involves minimizing a strongly convex quadratic function and is actually a linear operator. The point $x^{k+1/2}$ can be found by solving the following system of equations:

$$\begin{bmatrix} P + \rho I & A^T \\ A & -(1/\rho)I \end{bmatrix} \begin{bmatrix} x^{k+1/2} \\ v \end{bmatrix} = \begin{bmatrix} -q + \rho\left(x^k + A^T b - \begin{bmatrix} A^T & I \end{bmatrix} u^k\right) \\ 0 \end{bmatrix}.$$

Since the matrix on the lefthand side remains constant for all iterations, we can precompute the LDL^T factorization of this matrix once and cache the factorization for use in subsequent iterations. When P and A are dense, the factorization cost is $O(n^3)$ yet each subsequent iteration costs only $O(n^2)$. (Both factorization and solve costs can significantly smaller if P or A is sparse.) Amortizing the factorization step over all iterations means that the first step is quite efficient. Also notice that the matrix on the lefthand side is quasi-definite and hence favorable for LDL^T factorization.

In many applications, P and A do not change across problem instances. In this case, for different problem instances, we solve (1) for the same P and A and varying b and q. This lets us use the same LDL^T factorization, which results in a significant saving in computation.

The second step involves projection onto $\mathcal{X} = \mathcal{X}_1 \times \cdots \times \mathcal{X}_n$ and can typically be done much more quickly than the first step. It can be done in parallel since the projection onto $\mathcal{X}$ can be found by projections onto $\mathcal{X}_i$ for $i = 1, \ldots, n$. The third step is simply a dual update and is computationally inexpensive.

2.5 Preconditioning

Both theoretical analysis and practical evidence suggest that the precision and convergence rate of first-order methods can be significantly improved by preconditioning the problem. Here, we use diagonal scaling as preconditioning as discussed in [Bec14] and [WN99]. Diagonal scaling can be viewed as applying an appropriate linear transformation before running the algorithm. When the set $\mathcal{X}$ is convex, the preconditioning can substantially affect the speed of convergence, but does not affect the quality of the point returned, (which must be a solution to the convex problem). In other words, for convex problems, preconditioning is simply a tool to help the algorithm converge faster. Optimal choice of preconditioners, even in the convex case, is still an active research area [GB14a, GB14c, Gis14, GB14b, GTSJ15, SLY+14, HL12, Bol13, DY12]. In the nonconvex case, however, preconditioning can have a critical role in the *quality* of approximate solution, as well as the speed at which this solution is found.

Specifically, let $F \in \mathbf{R}^{n \times n}$, $E \in \mathbf{R}^{m \times m}$ be diagonal matrices with positive diagonal entries. The goal is to choose F and E such that running ADMM on the following problem has better convergence properties

$$
\begin{aligned}
\text{minimize} \quad & (1/2)x^T P x + q^T x + I_{\mathcal{X}}(z) \\
\text{subject to} \quad & \begin{bmatrix} EA \\ F \end{bmatrix} x - \begin{bmatrix} 0 \\ F \end{bmatrix} z = \begin{bmatrix} Eb \\ 0 \end{bmatrix}.
\end{aligned}
\tag{3}
$$

We use the choice of E and F recommended in [GB14a] to minimize the effective condition number (the ratio of the largest singular value to the smallest non-zero singular value) of the following matrix

$$
\begin{bmatrix} E & 0 \\ 0 & F \end{bmatrix} \begin{bmatrix} A \\ I \end{bmatrix} P^\dagger \begin{bmatrix} A^T & I \end{bmatrix} \begin{bmatrix} E & 0 \\ 0 & F \end{bmatrix},
$$

where $P^\dagger$ denotes the pseudo-inverse of P. Given matrix $M \in \mathbf{R}^{n \times n}$, minimizing the condition number of DMD for diagonal $D \in \mathbf{R}^{n \times n}$ can be cast as a semidefinite program. However, a heuristic called *matrix equilibration* can be used to avoid the computational cost of solving a semidefinite program. (See [Slu69, Bra10] and references therein.) Since for embedded applications computational resources are limited, we avoid finding $P^\dagger$ or equilibrating completely. We instead find E to normalize the rows of A (usually in ℓ_1 or ℓ_2 norm) and set F to be the identity.

After finding E and F, preconditioned ADMM has the following form:

$$x^{k+1/2} := \begin{bmatrix} I & 0 \end{bmatrix} \begin{bmatrix} P + \rho F^2 & A^T E \\ EA & -(1/\rho)I \end{bmatrix}^{-1} \begin{bmatrix} -q + \rho \left(F^2 x^k + A^T E^2 b - \begin{bmatrix} A^T E & F \end{bmatrix} u^k \right) \\ 0 \end{bmatrix}$$

$$x^{k+1} := \Pi \left(x^{k+1/2} + \begin{bmatrix} 0 & F^{-1} \end{bmatrix} u^k \right)$$

$$u^{k+1} := u^k + \begin{bmatrix} EA \\ F \end{bmatrix} x^{k+1/2} - \begin{bmatrix} 0 \\ F \end{bmatrix} x^k - \begin{bmatrix} Eb \\ 0 \end{bmatrix}.$$

$$(4)$$

2.6 The overall algorithm

We use the update rules (4) for $k = 1, \ldots, N$, where N denotes the (fixed) maximum number of iterations. Also, as described above, the algorithm is repeated for M number of random initializations. The computational cost of the algorithm consists of a factorization and MN matrix products and projections. Here is a description of the overall algorithm with $f(x) = (1/2)x^T P x + q^T x + r$.

Algorithm 1 Approximately solving nonconvex constraint QP (1)

---

if A or P changed **then**

    find E and F by equilibrating $\begin{bmatrix} A \\ I \end{bmatrix} P^\dagger \begin{bmatrix} A^T & I \end{bmatrix}$

    find and store LDL factorization of $\begin{bmatrix} P + \rho F^2 & A^T E \\ EA & -(1/\rho)I \end{bmatrix}$

end if

$x_{\text{best}} := \emptyset$, $f(x_{\text{best}}) := \infty$

for random initialization $1, 2, \ldots, N$ **do**

    for iteration $1, 2, \ldots, M$ **do**

        update x from (4)

        if $\|Ax - b\| \le \epsilon^{\text{tol}}$ and $f(x) < f(x_{\text{best}})$ **then**

            $x_{\text{best}} = x$

        end if

    end for

end for

return x_{best}.

We mention a solution refinement technique here that can be used to find a solution with possibly better objective value after the algorithm stops. This technique, sometimes known as *polishing* consists of fixing the nonconvex variable and solving the resulting convex optimization problem. Using this technique, one may use larger ϵ^{tol} during the N iterations and only reduce ϵ^{tol} at the refinement step. Depending on the application, it might be computationally sensible to solve the resulting convex optimization problem. Another effective technique is to introduce a notion of *no-good cut* during iterations for problems with binary

variables. A no-good cut prohibits the integer part to be equal to the previous one, by imposing one additional inequality constraint $\sum_{i \in T} x_{b_i}^{k+1/2} - \sum_{i \in F} x_{b_i}^{k+1/2} \leq B - 1$, where $x_{b_1}, x_{b_2}, \ldots$ are binary variables and $T = \{i | x_{b_i}^k = 1\}$, $F = \{i | x_{b_i}^k = 0\}$, and B is the number of elements of T. We do not use either of these techniques in the following examples.

3 Numerical examples

In this section, we explore the performance of our proposed algorithm on some example problems. For each example, ρ was chosen between 0.1 and 10 to yield good performance; all other algorithm parameters were kept constant. As a benchmark, we compare our results to the commercial solver MOSEK, which can globally solve MIQPs. All experiments were carried out on a system with two 3.06 GHz cores with 4 GB of RAM.

The results suggest that this heuristic is effective in finding approximate solutions for mixed integer quadratic programs.

3.1 Randomly generated QP

First we demonstrate the performance of our algorithm qualitatively for a random mixed-Boolean quadratic program. The matrix P in (1) was chosen as $P = QQ^T$, where the entries of $Q \in \mathbf{R}^{n \times n}$, as well as those of q and A, were drawn from a standard normal distribution. The constant r was chosen such that the optimal value of the unconstrained quadratic minimization is 0. The vector b was chosen as $b = Ax_0$, where $x_0 \in \mathcal{X}$ was chosen uniformly randomly, thus ensuring that the problem is feasible. We used $n = 200$ and $m = 50$ with $\mathcal{X}_i = \{0, 1\}$ for $i = 1, \ldots, 100$, $\mathcal{X}_i = \mathbf{R}_+$ for $i = 101, \ldots, 150$, and $\mathcal{X}_i = \mathbf{R}$ for the other indices i.

We used MOSEK to find the optimal value for the problem. After 60832 seconds (more than 16 hours), MOSEK certifies that the optimal value is equal to 2040. We ran algorithm 1 for 10 different initializations and 200 iterations for each initialization, with step size $\rho = 0.5$. For a naive implementation in MATLAB, it took 120 milliseconds to complete all precomputations (preconditioning and factorization), and 800 milliseconds to do all 2000 iterations. The best objective value found for the problem was 2067 (1.3% suboptimal).

One interesting observation is that the parameter ρ tends to trade off feasibility and optimality: with small values of ρ, the algorithm often fails to find a feasible point, but feasible points found tend to have low objective value. On the other hand, with large values of ρ, feasible points are found more quickly, but tend to have higher objective value.

3.2 Hybrid vehicle control

We consider a simple hybrid electric vehicle drivetrain (similar to that of [BV04, Exercise 4.65]), which consists of a battery, an electric motor/generator, and a heat engine, in a parallel configuration. We assume that the demanded power P_t^{des} at the times $t = 0, \ldots, T-1$

is known in advance. Our task is to plan out the battery and engine power outputs P_t^{batt} and P_t^{eng}, for $t = 0, \ldots, T-1$, so that

$$P_t^{\text{batt}} + P_t^{\text{eng}} \geq P_t^{\text{des}}.$$

(Strict inequality above corresponds to braking.)

Battery. The battery has stored energy E_t at time t, which evolves according to

$$E_{t+1} = E_t - \tau P_t^{\text{batt}}, \qquad t = 0, \ldots, T-1,$$

where τ is the length of each discretized time interval. The battery capacity is limited, so that $0 \leq E_t \leq E^{\text{max}}$ for all t, and the initial energy E_0 is known. We penalize the terminal energy state of the battery according to $g(E_T)$, where

$$g(E) = \eta(E^{\text{max}} - E)^2,$$

for $\eta \geq 0$.

Engine. At time t, the engine may be on or off, which is modeled with binary variable z_t. If the engine is on ($z_t = 1$), then we have $0 \leq P_t^{\text{eng}} \leq P^{\text{max}}$, and $\alpha(P_t^{\text{eng}})^2 + \beta P_t^{\text{eng}} + \gamma$ units of fuel are consumed, for nonnegative constants α, β, and γ. If the engine is off ($z_t = 0$), it consumes no fuel, and $P_t^{\text{eng}} = 0$. Because $z_t \in \{0, 1\}$, the power constraint can be written as $0 \leq P^{\text{eng}} \leq P^{\text{max}} z_t$, and the fuel cost as $f(P_t^{\text{eng}}, z_t)$, where

$$f(P, z) = \alpha P^2 + \beta P + \gamma z.$$

Additionally, we assume that turning the engine on after it has been off incurs a cost $\delta \geq 0$, i.e., at each time t, we pay $\delta(z_t - z_{t-1})_+$, where $(\cdot)_+$ denotes the positive part.

Optimal power split problem. The hybrid vehicle control problem can be formulated as

$$
\begin{array}{ll}
\text{minimize} & \eta(E_T - E^{\text{max}})^2 + \sum_{t=0}^{T-1} f(P_t^{\text{eng}}, z_t) + \delta(z_t - z_{t-1})_+ \\
\text{subject to} & E_{t+1} = E_t - \tau P_t^{\text{batt}} \\
& P_t^{\text{batt}} + P_t^{\text{eng}} \geq P_t^{\text{des}} \\
& z_t \in \{0, 1\},
\end{array}
\tag{5}
$$

where all constraints must hold for $t = 0, \ldots, T-1$. The variables are P_t^{batt}, P_t^{eng}, and z_t for $t = 0, \ldots, T-1$, and E_t, for $t = 1, \ldots, T$. In addition to the parameters given above, we take z_{-1} to be a parameter denoting the initial engine state.

We used the parameter values $\alpha = 1$, $\beta = 10$, $\gamma = 1.5$, $\delta = 10$, $\eta = 0.1$, $\tau = 5$, $P^{\text{max}} = 1$, $E^{\text{max}} = 200$, $E_0 = 200$, and $z_{-1} = 0$. The demanded power trajectory P_t^{des} is not shown, but can be obtained by summing the engine power and battery power in Figure 1. We ran the algorithm with $\rho = 0.4$ for 1000 iterations from 5 different initializations, with primal optimality threshold $\epsilon^{\text{tol}} = 10^{-4}$. The global solution found by MOSEK generates an objective value of 339.2 and the best objective value with our algorithm was 375.7. In Figure 1, we see that qualitatively, the optimal trajectory and the trajectory generated by ADMM are very similar.
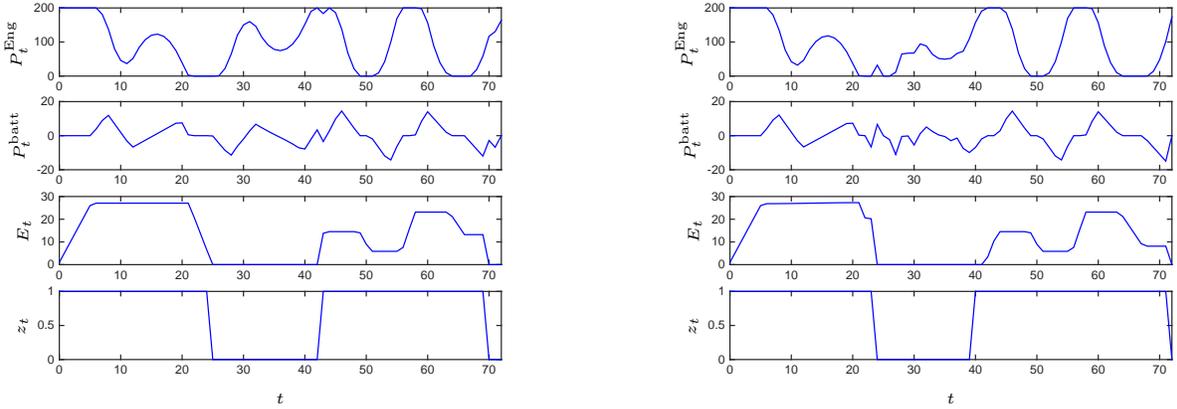
Figure 1: Engine power, battery power, battery energy, and engine on/off signals versus time. Left: the global solution. Right: the solution found using ADMM.
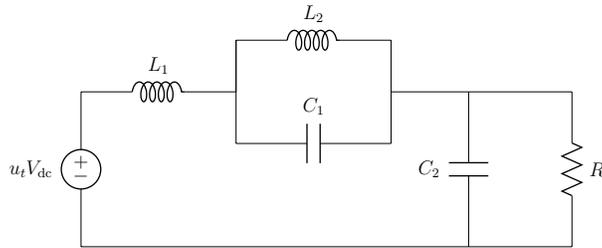
Figure 2: Converter circuit model.

3.3 Power converter control

We consider control of the switched-mode power converter shown in Figure 2. The circuit dynamics are

$$\xi_{t+1} = G\xi_t + Hu_t, \qquad t = 0, 1, \ldots, T-1,$$

where $\xi_t = (i_{1,t}, v_{1,t}, i_{2,t}, v_{2,t})$ is the system state at epoch t, consisting of all inductor currents and capacitor voltages, and $u_t \in \{-1, 0, 1\}$ is the control input. The dynamics matrices $G \in \mathbf{R}^{4\times4}$ and $H \in \mathbf{R}^{4\times1}$ are obtained by discretizing the dynamics of the circuit in Figure 2.

We would like to control the switch configurations so that v_2 tracks a desired sinusoidal waveform. This can be done by solving

$$
\begin{array}{ll}
\text{minimize} & \sum_{t=0}^{T}(v_{2,t} - v_{\text{des}})^2 + \lambda|u_t - u_{t-1}| \\
\text{subject to} & \xi_{t+1} = G\xi_t + Hu_t \\
& \xi_0 = \xi_T \\
& u_0 = u_T \\
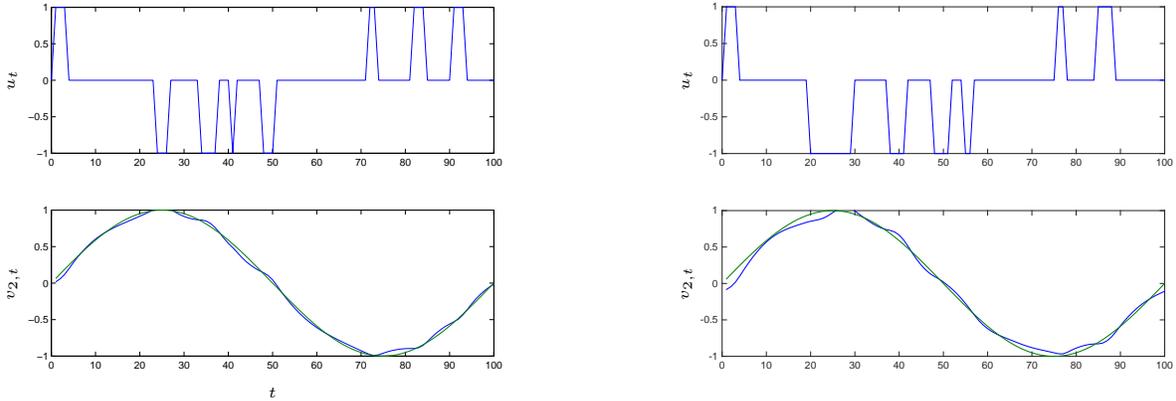& u_t \in \{-1, 0, 1\},
\end{array}
\tag{6}
$$

Figure 3: The switch configuration and the output voltage. Left: the global solution. Right: the solution using ADMM.

where $\lambda \geq 0$ is a tradeoff parameter between output voltage regulation and switching frequency. The variables are ξ_t for $t = 0, \ldots, T$ and u_t for $t = 0, \ldots, T-1$.

Note that if we take $\lambda = 0$, and take the input voltage u_t to be unconstrained (*i.e.*, allow u_t to take any values in $\mathbf{R}$), (6) can be solved as a convex quadratic minimization problem, with solution ξ_t^{ls}. Returning to our original problem, we can penalize deviation from this ideal waveform by including a regularization term $\mu \| \xi - \xi_t^{\text{ls}} \|^2$ to (6), where $\mu > 0$ is a positive weighting parameter. We solved this regularized version of (6), with $L_1 = 10 \, \mu\text{H}$, $C_1 = 1 \, \mu\text{F}$, $L_2 = 10 \, \mu\text{H}$, $C_2 = 10 \, \mu\text{F}$, $R = 1 \, \Omega$, $V_{\text{dc}} = 10 \, \text{V}$, $T = 100$ (with a discretization interval of $0.5 \, \mu\text{s}$), $\lambda = 1.5 \, \text{V}^2$, and $\mu = 0.1$. We run algorithm 1 with $\rho = 2.7$ and 500 iterations for three different initializations. An approximate solution is found via our heuristic in less than 2 seconds, whereas it takes MOSEK more than 4 hours to find the global solution. Figure 3 compares the approximate solution derived by the heuristic with the global solution.

3.4   Signal decoding

We consider maximum-likelihood decoding of a message passed through a linear multiple-input and multiple-output (MIMO) channel. In particular, we have

$$y = Hx + v,$$

where $y \in \mathbf{R}^p$ is the message received, $H \in \mathbf{R}^{p \times n}$ is the channel matrix, $x \in \mathbf{R}^n$ is the message sent, and the elements of the noise vector $v \in \mathbf{R}^p$ are independent, identically distributed Gaussian random variables. We further assume that the elements of x belong to the *signal constellation* $\{-3, -1, 1, 3\}$. The maximum likelihood estimate of x is given by the solution to the problem

$$\begin{aligned}
&\text{minimize} && \| H\hat{x} - y \|^2 \\
&\text{subject to} && \hat{x}_i \in \{-3, -1, 1, 3\}, \quad i = 1, \ldots, n,
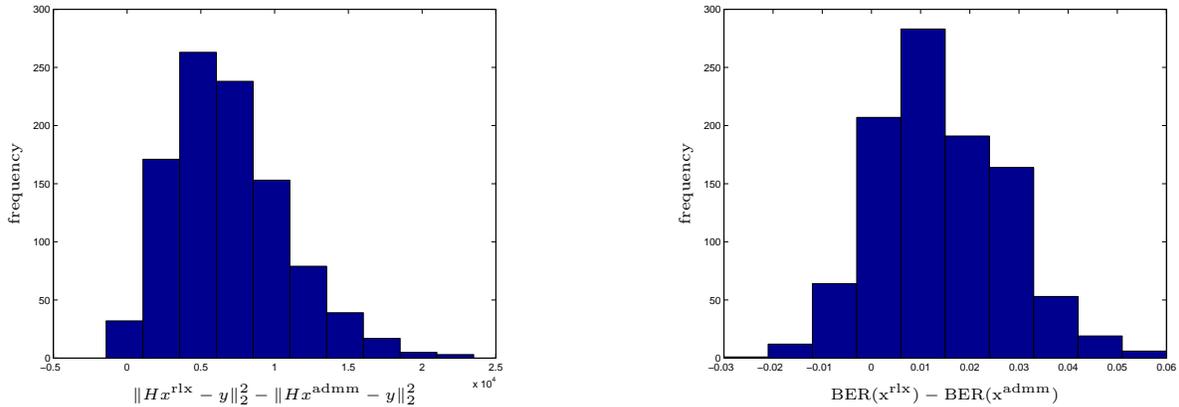\end{aligned} \tag{7}$$

11

Figure 4: Comparison of ADMM heuristic and relax-and-round. Left: The difference in objective values. Right: The difference in bit error rates (BER).

where $\hat{x} \in \mathbf{R}^n$ is the variable.

We generate 1000 random problem instances with $H \in \mathbf{R}^{2000 \times 400}$ chosen from a standard normal distribution. The uncorrupted signal x is chosen uniformly randomly and the additive noise is Gaussian such that the signal to noise ratio (SNR) is 8 dB. For such a problem in embedded application, branch-and-bound methods are not desirable due to their worst-case time complexity. We run the heuristic with only one initialization, with 10 iterations to find x^{admm}. The average runtime for each problem (including preprocessing) is 80 milliseconds, which is substantially faster than branch-and-bound based methods. We compare the performance of the points x^{admm} with the points found by relax-and-round technique x^{rlx}. In Figure 4 we have plotted the histogram of the difference between the objective values evaluated at x^{admm} and x^{rlx}. Depicted in Figure 4, we see that in 95% of the cases, the bit error rate (BER) using our heuristic was at least as good as the bit error rate using relax and round.

4 Conclusions

In this paper, we introduced an effective heuristic for finding approximate solutions to convex quadratic minimization problems over the intersection of affine and nonconvex sets. Our heuristic is significantly faster than branch-and-bound algorithms and has shown effective in a variety of embedded problems including hybrid vehicle control, power converter control, and signal decoding.

References

[AB07]   T. Achterberg and T. Berthold. Improving the Feasibility Pump. *Discrete Optimization*, 4(1):77–86, 2007.

[Ach09]   T. Achterberg. SCIP: Solving Constraint Integer Programs. *Mathematical Programming Computation*, 1(1):1–41, 2009.

[AH04]   D. Axehill and A. Hansson. A preprocessing algorithm for MIQP solvers with applications to MPC. In *Proceedings of the 43rd IEEE Conference on Decision and Control*, volume 3, pages 2497–2502, 2004.

[AL74]   R. N. Adams and M. A. Laughton. Optimal planning of power networks using mixed-integer programming. part 1: static and time-phased network synthesis. *Proceedings of the Institution of Electrical Engineers*, 121(2):139–147, 1974.

[ApS15]   MOSEK ApS. *TheMOSEKoptimization toolbox for MATLAB manual. Version 7.1 (Revision 28)*, 2015.

[Bea98]   J. E. Beasley. Heuristic algorithms for the unconstrained binary quadratic programming problem. *Management School, Imperial College, London, UK*, 1998.

[Bec14]   A. Beck. *Introduction to Nonlinear Optimization: Theory, Algorithms, and Applications with MATLAB*, volume 19. SIAM, 2014.

[Bem15]   A. Bemporad. Solving mixed-integer quadratic programs via nonnegative least squares. *5th IFAC Conference on Nonlinear Model Predictive Control*, page 73?79, 2015.

[BFG+00]   E. R. Bixby, M. Fenelon, Z. Gu, E. Rothberg, and R. Wunderling. MIP: theory and practice, closing the gap. In *System modeling and optimization*, pages 19–49. Springer, 2000.

[BFG+04]   R. E. Bixby, M. Fenelon, Z. Gu, E. Rothberg, and R. Wunderling. Mixed-integer programming: a progress report. *The Sharpest Cut: The Impact of Manfred Padberg and his work, MPS-SIAM Series on Optimization*, 4:309–326, 2004.

[BFL07]   L. Bertacco, M. Fischetti, and A. Lodi. A feasibility pump heuristic for general mixed-integer problems. *Discrete Optimization*, 4(1):63–76, 2007.

[Bie96]   D. Bienstock. Computational study of a family of mixed-integer quadratic programming problems. *Mathematical Programming*, 74(2):121–140, 1996.

[BM99]   A. Bemporad and M. Morari. Control of systems integrating logic, dynamics, and constraints. *Automatica*, 35(3):407–427, 1999.

[Bol13]   D. Boley. Local linear convergence of the alternating direction method of multipliers on quadratic or linear programs. *SIAM Journal on Optimization*, 23(4):2183–2207, 2013.

[BP12]   A. Bemporad and P. Patrinos. Simple and certifiable quadratic programming algorithms for embedded linear model predictive control. In *Nonlinear Model Predictive Control*, volume 4, pages 14–20, 2012.

[BPC⁺11]   S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.

[BR07]   R. Bixby and E. Rothberg. Progress in computational mixed-integer programming, a look back from the other side of the tipping point. *Annals of Operations Research*, 149(1):37–41, 2007.

[Bra10]   A. M. Bradley. *Algorithms for the Equilibration of Matrices and their Application to Limited-Memory Quasi-Newton Methods*. PhD thesis, Stanford University, 2010.

[BRL01]   A. Bemporad, J. Roll, and L. Ljung. Identification of hybrid systems via mixed-integer programming. In *IEEE Conference on Decision and Control*, volume 1, pages 786–792, 2001.

[BTT91]   M. Bierlaire, P. L. Toint, and D. Tuyttens. On iterative algorithms for linear least squares problems with bound constraints. *Linear Algebra and its Applications*, 143:111–143, 1991.

[BV04]   S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[BW05]   D. Bertsimas and R. Weismantel. *Optimization Over Integers*, volume 13. Dynamic Ideas, Belmont, Massachusetts, 2005.

[CA06]   M. Carrión and J. M. Arroyo. A computationally efficient mixed-integer linear formulation for the thermal unit commitment problem. *IEEE Transactions on Power Systems*, 21(3):1371–1378, 2006.

[Cam94]   J. F. Campbell. Integer programming formulations of discrete hub location problems. *European Journal of Operational Research*, 72(2):387–405, 1994.

[CCH89]   V. Chvátal, W. Cook, and M. Hartmann. On cutting-plane proofs in combinatorial optimization. *Linear Algebra and its Applications*, 114:455–499, 1989.

[CCZ14]   M. Conforti, G. Cornuejols, and G. Zambelli. *Integer Programming*. Graduate Texts in Mathematics. Springer International Publishing, 2014.

[Cha12]   R. Chartrand. Nonconvex splitting for regularized low-rank + sparse decomposition. *IEEE Transactions on Signal Processing*, 60(11):5810–5819, 2012.

[CPDB13]   E. Chu, N. Parikh, A. Domahidi, and S. Boyd. Code generation for embedded second-order cone programming. In *Proceedings of the 2013 European Control Conference*, pages 1547–1552, 2013.

[CPL09]   IBM ILOG CPLEX. User's manual for CPLEX. *International Business Machines Corporation*, 46(53):157, 2009.

[CPM10]   J. P. S. Catalão, H. M. I. Pousinho, and V. M. F. Mendes. Scheduling of head-dependent cascaded hydro systems: Mixed-integer quadratic programming approach. *Energy Conversion and Management*, 51(3):524–530, 2010.

[CW13]   R. Chartrand and B. Wohlberg. A nonconvex ADMM algorithm for group sparsity with sparse groups. In *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6009–6013. IEEE, 2013.

[DBEY13]   N. Derbinsky, J. Bento, V. Elser, and J. S. Yedidia. An improved three-weight message-passing algorithm. *arXiv preprint arXiv:1305.1961*, 2013.

[DCB13]   A. Domahidi, E. Chu, and S. Boyd. ECOS: An SOCP solver for embedded systems. In *Proceedings of the 12th European Control Conference*, pages 3071–3076. IEEE, 2013.

[DY12]   W. Deng and W. Yin. On the global and linear convergence of the generalized alternating direction method of multipliers. *Journal of Scientific Computing*, pages 1–28, 2012.

[FDM15]   D. Frick, A. Domahidi, and M. Morari. Embedded optimization for mixed logical dynamical systems. *Computers and Chemical Engineering*, 72:21–33, 2015.

[FGL05]   M. Fischetti, F. Glover, and A. Lodi. The feasibility pump. *Mathematical Programming*, 104(1):91–104, 2005.

[Flo95]   C. A. Floudas. *Nonlinear and Mixed-Integer Optimization: Fundamentals and Applications*. Oxford University Press, 1995.

[G+58]   R. E. Gomory et al. Outline of an algorithm for integer solutions to linear programs. *Bulletin of the American Mathematical society*, 64(5):275–278, 1958.

[GARK02]   F. Glover, B. Alidaee, C. Rego, and G. Kochenberger. One-pass heuristics for large-scale unconstrained binary quadratic problems. *European Journal of Operational Research*, 137(2):272–287, 2002.

[GB14a]   P. Giselsson and S. Boyd. Diagonal scaling in Douglas-Rachford splitting and ADMM. In *53rd Annual IEEE Conference on Decision and Control (CDC)*, pages 5033–5039, 2014.

[GB14b]   P. Giselsson and S. Boyd. Monotonicity and restart in fast gradient methods. In *53rd Annual IEEE Conference on Decision and Control (CDC)*, pages 5058–5063, 2014.

[GB14c]   P. Giselsson and S. Boyd. Preconditioning in fast dual gradient methods. In *53rd Annual IEEE Conference on Decision and Control (CDC)*, pages 5040–5045, 2014.

[GG10]   I. Glover and P. M. Grant. *Digital Communications*. Pearson Education, 2010.

[Gis14]   P. Giselsson. Improved fast dual gradient methods for embedded model predictive control. In *International Federation of Automatic Control (IFAC)*, pages 2303–2309, 2014.

[GJ02]   M. R. Garey and D. S. Johnson. *Computers and Intractability*, volume 29. Freeman, 2002.

[GM72]   A. M. Geoffrion and R. E. Marsten. Integer programming algorithms: a framework and state-of-the-art survey. *Management Science*, 18(9):465–491, 1972.

[GSS05]   M. Guignard-Spielberg and K. Spielberg. *Integer Programming: State of the Art and Recent Advances*, volume 140. Springer, 2005.

[GTSJ15]   E. Ghadimi, A. Teixeira, I. Shames, and M. Johansson. Optimal parameter selection for the alternating direction method of multipliers (ADMM): Quadratic problems. *IEEE Transactions on Automatic Control*, 60(3):644–658, 2015.

[HL12]   M. Hong and Z. Luo. On the linear convergence of the alternating direction method of multipliers. *arXiv preprint arXiv:1208.3922*, 2012.

[Hoc82]   D. S. Hochbaum. Approximation algorithms for the set covering and vertex cover problems. *SIAM Journal on Computing*, 11(3):555–556, 1982.

[Iba76]   T. Ibaraki. Integer programming formulation of combinatorial optimization problems. *Discrete Mathematics*, 16(1):39–52, 1976.

[Jer80]   R. G. Jeroslow. Representations of unbounded optimization problems as integer programs. *Journal of Optimization Theory and Applications*, 30(3):339–351, 1980.

[Jer87]   R. G. Jeroslow. Representability in mixed-integer programming: Characterization results. *Discrete Applied Mathematics*, 17(3):223–243, 1987.

[Jer89]   R. G. Jeroslow. Representability of functions. *Discrete Applied Mathematics*, 23(2):125–137, 1989.

[JGR+14]   J. L. Jerez, P. J. Goulart, S. Richter, G. Constantinides, E. C. Kerrigan, M. Morari, et al. Embedded online optimization for model predictive control at megahertz rates. *IEEE Transactions on Automatic Control*, 59(12):3238–3251, 2014.

[JL85]   R. G. Jeroslow and J. K. Lowe. Experimental results on the new techniques for integer programming formulations. *Journal of the Operational Research Society*, pages 393–403, 1985.

[JLN+09]   M. Jünger, T. M. Lieblingand, D. Naddef, G. L. Nemhauser, W. R. Pulleyblank, G. Reinelt, G. Rinaldi, and L. A. Wolsey. *50 Years of Integer Programming 1958-2008: From the Early Years to the State of the Art.* Springer Science and Business Media, 2009.

[JNS00]   E. L. Johnson, G. L. Nemhauser, and M. W. P. Savelsbergh. Progress in linear programming-based algorithms for integer programming: an exposition. *INFORMS Journal on Computing*, 12(1):2–23, 2000.

[JW90]   R. G. Jeroslow and J. Wang. Solving propositional satisfiability problems. *Annals of Mathematics and Artificial Intelligence*, 1(1-4):167–187, 1990.

[Kar72]   R. M. Karp. *Reducibility Among Combinatorial Problems.* Springer, 1972.

[KN01]   K. Katayama and H. Narihisa. Performance of simulated annealing-based heuristic for the unconstrained binary quadratic programming problem. *European Journal of Operational Research*, 134(1):103–119, 2001.

[lAW98]   l. A. Wolsey. *Integer Programming*, volume 42. Wiley New York, 1998.

[Laz82]   R. Lazimy. Mixed-integer quadratic programming. *Mathematical Programming*, 22(1):332–349, 1982.

[LD10]   A. H. Land and A. G. Doig. An automatic method for solving discrete programming problems. In *50 Years of Integer Programming 1958-2008*, pages 105–132. Springer, 2010.

[Ley94]   R. Fletcherand S. Leyffer. Solving mixed-integer nonlinear programs by outer approximation. *Mathematical Programming*, 66(1-3):327–349, 1994.

[Ley01]   S. Leyffer. Integrating SQP and branch-and-bound for mixed-integer nonlinear programming. *Computational Optimization and Applications*, 18(3):295–309, 2001.

[LS06]   D. Li and X. Sun. *Nonlinear Integer Programming*, volume 84. Springer Science and Business Media, 2006.

[LW66]   E. L. Lawler and D. E. Wood. Branch-and-bound methods: a survey. *Operations Research*, 14(4):699–719, 1966.

[LZD04]   R. Li, D. Zhou, and D. Du. Satisfiability and integer programming as complementary tools. In *Proceedings of the 2004 Asia and South Pacific Design Automation Conference*, pages 879–882, 2004.

[LZW$^+$07] Z. Li, S. Zhang, Y. Wang, X. Zhang, and L. Chen. Alignment of molecular networks by integer quadratic programming. *Bioinformatics*, 23(13):1631–1639, 2007.

[MB10]   J. Mattingley and S. Boyd. Automatic code generation for real-time convex optimization. *Convex Optimization in Signal Processing and Communications*, pages 1–41, 2010.

[MB12]   J. Mattingley and S. Boyd. CVXGEN: a code generator for embedded convex optimization. *Optimization and Engineering*, 13(1):1–27, 2012.

[Mey75]   R. R. Meyer. Integer and mixed-integer programming models: general properties. *Journal of Optimization Theory and Applications*, 16(3-4):191–206, 1975.

[Mey76]   R. R. Meyer. Mixed-integer minimization models for piecewise-linear functions of a single variable. *Discrete Mathematics*, 16(2):163–171, 1976.

[Mey81]   R. R. Meyer. A theoretical and computational comparison of equivalent mixed-integer formulations. *Naval Research Logistics Quarterly*, 28(1):115–131, 1981.

[MF02]   P. Merz and B. Freisleben. Greedy and local search heuristics for unconstrained binary quadratic programming. *Journal of Heuristics*, 8(2):197–213, 2002.

[MF13]   R. Misener and C. A. Floudas. GloMIQO: global mixed-integer quadratic optimizer. *Journal of Global Optimization*, 57(1):3–50, 2013.

[MJSB12]   N. Murgovski, L. Johannesson, J. Sjöberg, and B.Egardt. Component sizing of a plug-in hybrid electric powertrain via convex optimization. *Mechatronics*, 22(1):106–120, 2012.

[MK87]   K. G. Murty and S. N. Kabadi. Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39(2):117–129, 1987.

[MTH80]   R. R. Meyer, M. V. Thakkar, and W. P. Hallman. Rational mixed-integer and polyhedral union minimization models. *Mathematics of Operations Research*, 5(1):135–146, 1980.

[MWB11]   J. Mattingley, Y. Wang, and S. Boyd. Receding horizon control: Automatic generation of high-speed solvers. *IEEE Control Systems Magazine*, 31(3):52–65, 2011.

[O$^+$12]   Gurobi Optimization et al. Gurobi optimizer reference manual. *URL: http://www.gurobi.com*, 2012.

[OCP07]   D. Oulai, S. Chamberland, and S. Pierre. A new routing-based admission control for MPLS networks. *IEEE Communications Letters*, 11(2):216–218, 2007.

[O'k87]   M. E. O'kelly. A quadratic integer program for the location of interacting hub facilities. *European Journal of Operational Research*, 32(3):393–404, 1987.

[OSB13]   B. O'Donoghue, G. Stathopoulos, and S. Boyd. A splitting method for optimal control. *IEEE Transactions on Control Systems Technology*, 21(6):2432–2442, 2013.

[Pad73]   M. W. Padberg. On the facial structure of set packing polyhedra. *Mathematical Programming*, 5(1):199–215, 1973.

[PB13]   N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):123–231, 2013.

[PB14]   N. Parikh and S. Boyd. Block splitting for distributed optimization. *Mathematical Programming Computation*, 6(1):77–102, 2014.

[PF07]   L. G. Papageorgiou and E. S. Fraga. A mixed-integer quadratic programming formulation for the economic dispatch of generators with prohibited operating zones. *Electric Power Systems Research*, 77(10):1292–1296, 2007.

[PS98]   C. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Courier Corporation, 1998.

[PU04]   M. Propato and J. G. Uber. Booster system design using mixed-integer quadratic programming. *Journal of Water Resources Planning and Management*, 130(4):348–352, 2004.

[Slu69]   A. V. D. Sluis. Condition numbers and equilibration of matrices. *Numerische Mathematik*, 14(1):14–23, 1969.

[SLY$^+$14]   W. Shi, Q. Ling, K. Yuan, G. Wu, and W. Yin. On the linear convergence of the admm in decentralized consensus optimization. *IEEE Transactions on Signal Processing*, 62(7):1750–1761, 2014.

[SM99]   R. A. Stubbs and S. Mehrotra. A branch-and-cut method for 0-1 mixed convex programming. *Mathematical Programming*, 86(3):515–532, 1999.

[SMFH01] T. Schouwenaars, B. De Moor, E. Feron, and J. How. Mixed-integer program-ming for multi-vehicle path planning. In *European Control Conference*, volume 1, pages 2603–2608. Citeseer, 2001.

[TS04] M. Tawarmalani and N. V. Sahinidis. Global optimization of mixed-integer nonlinear programs: a theoretical and computational study. *Mathematical Pro-gramming*, 99(3):563–591, 2004.

[Ull11] F. Ullmann. FiOrdOs: A Matlab toolbox for C-code generation for first order methods. Master's thesis, ETH Zurich, 2011.

[VAN08] J. P. Vielma, S. Ahmed, and G. L. Nemhauser. A lifted linear programming branch-and-bound algorithm for mixed-integer conic quadratic programs. *IN-FORMS Journal on Computing*, 20(3):438–450, 2008.

[WB10] Y. Wang and S. Boyd. Fast model predictive control using online optimization. *IEEE Transactions on Control Systems Technology*, 18(2):267–278, 2010.

[WFGX08] Z. Wang, S. Fang, D. Y. Gao, and W. Xing. Global extremal conditions for multi-integer quadratic programming. *Journal of Industrial and Management Optimization*, 4(2):213–225, 2008.

[WN99] S. J. Wright and J. Nocedal. *Numerical Optimization*, volume 2. Springer New York, 1999.

[WN14] L. A. Wolsey and G. L. Nemhauser. *Integer and Combinatorial Optimization*. John Wiley & Sons, 2014.