

Convex Optimization of Nonlinear Feedback Controllers via Occupation Measures

Anirudha Majumdar, Ram Vasudevan, Mark M. Tobenkin, and Russ Tedrake
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139
Email: {anirudha,ramv,mmt,russt}@mit.edu

Abstract—In this paper, we present an approach for designing feedback controllers for polynomial systems that maximize the size of the time-limited backwards reachable set (BRS). We rely on the notion of *occupation measures* to pose the synthesis problem as an infinite dimensional linear program (LP) and provide finite dimensional approximations of this LP in terms of semidefinite programs (SDPs). The solution to each SDP yields a polynomial control policy and an outer approximation of the largest achievable BRS. In contrast to traditional Lyapunov based approaches which are non-convex and require feasible initialization, our approach is convex and does not require any form of initialization. The resulting time-varying controllers and approximated reachable sets are well-suited for use in a trajectory library or feedback motion planning algorithm. We demonstrate the efficacy and scalability of our approach on four nonlinear systems.

I. INTRODUCTION

Dynamic robotic tasks such as flying, running, or walking demand controllers that push hardware platforms to their physical limit while managing input saturation, nonlinear dynamics, and underactuation. Though motion planning algorithms have begun addressing several of these tasks [16], the constructed open loop motion plans are typically insufficient due to their inability to correct for deviations from a planned path. Despite the concerted effort of several communities, the design of feedback control laws for underactuated nonlinear systems with input saturation remains challenging.

Popular techniques for control synthesis rely either on feedback linearization [27] or on linearizing the dynamics about a nominal operating point in order to make Linear Quadratic Regulator based techniques or Linear Model Predictive Control [3] applicable. Unfortunately, feedback linearization is generally untenable for underactuated systems especially in the presence of actuation limits, and those techniques that rely on linearizations lead to controllers that are valid only locally around the operating point. Dynamic Programming and Hamilton-Jacobi Bellman Equation based techniques [8], [21] have also been used for feedback control design. However, these methods suffer from the curse of dimensionality, and can require exorbitant grid resolution for even low dimensional systems [23].

A. Our Contributions

In this paper, we attempt to address these issues and present an approach for designing feedback controllers that maximize

the time-limited *backward reachable set* (BRS), i.e. the set of points that reach a given target set at a specified finite time. Our approach is inspired by the method presented in [12], which describes a framework based on *occupation measures* for computing the BRS for polynomial systems. In this paper, we extend this method to the control synthesis problem. Our contributions are three-fold. First, in Section II, we formulate the design of the feedback controller that generates the largest BRS as an infinite dimensional linear program (LP) over the space of nonnegative measures. Second, in Section III-A, we construct a sequence of finite dimensional relaxations to our infinite dimensional LP in terms of semidefinite programs (SDPs). Finally, in Section III-B, we prove two convergence properties of our sequence of finite dimensional approximations: first that each solution to the sequence of SDPs is an *outer approximation* to the largest possible BRS with asymptotically vanishing conservatism; and second, that there exists a subsequence of the SDP solutions that weakly converges to an optimizing solution of our original infinite dimensional LP¹.

The result of our analysis is a method capable of designing feedback controllers for nonlinear underactuated robotic systems in the presence of input saturation without resorting to linear analysis. This is valuable for systems with degenerate linearizations, and can result in considerable improvements in performance for many practical robotic systems. Our method could also be used to augment existing feedback motion planning algorithms such as the LQR-Trees approach presented in [20], [30], which computes and sequences together BRSs in order to drive a desired set of initial conditions to some target set. Our approach could be substituted for the local, linear control synthesis employed by the aforementioned papers with the benefit of selecting control laws that maximize the size of the BRS in the presence of input saturations. As a result, the number of trajectories required in a library in order to fill the space of possible initial conditions could be significantly reduced. In some cases, a *single* nonlinear feedback controller could stabilize an entire set of initial conditions that previously required a library of locally-linear controllers. We illustrate the performance of our approach in Section IV on four examples,

¹The proofs of the presented arguments can be found in our technical report on arXiv.

whose source code we make available.

B. Relationship to Lyapunov-Based Techniques

Our approach is most comparable to those that use Lyapunov's criteria for stability in order to synthesize a controller that maximizes the region of attraction (ROA) of a particular target set. These criteria can be checked for polynomial systems by employing sums-of-squares (SOS) programming. However, the computation of the ROA and subsequent controller design are typically non-convex programs in this formulation [26]. The resulting optimization programs are bilinear in the decision variables and are generally solved by employing some form of bilinear alternation [14], [19]. Such methods are not guaranteed to converge to global optima (or necessarily even local optima) and require feasible initializations.

The relationship between our approach and the Lyapunov-based approaches can be understood by examining the dual of our infinite dimensional LP, which is posed on the space of nonnegative continuous functions. This dual program certifies that a certain set cannot reach the target set within a pre-specified time for any valid control law. The complement of this set is an outer approximation of the BRS. This subtle change transforms the non-convex feedback control synthesis problem written in terms of Lyapunov's criteria into the convex synthesis problem which we present herein. The convexity of the control synthesis problem we present also has parallels to the convexity observed in [26] during the design of controllers to achieve global almost-everywhere asymptotic stability. However, this method is not easily extended to provide regional certificates, which are of greater practical utility in robotics since robotic systems are generally not globally stabilizable.

II. PROBLEM FORMULATION

In this section, we formalize our problem of interest, construct an infinite dimensional linear program (LP), and prove that the solution of this LP is equivalent to solving our problem of interest. We make substantial use of measure theory, and the unfamiliar reader may wish to consult [10] for an introduction.

A. Notation

Given an element $y \in \mathbb{R}^{n \times m}$, let $[y]_{ij}$ denote the (i, j) -th component of y . We use the same convention for elements belonging to any multidimensional vector space. By \mathbb{N} we denote the non-negative integers, and \mathbb{N}_k^n refers to those $\alpha \in \mathbb{N}^n$ with $|\alpha| = \sum_{i=1}^n [\alpha]_i \leq k$. Let $\mathbb{R}[y]$ denote the ring of real polynomials in the variable y . For a compact set K , let $\mathcal{M}(K)$ denote the space of signed Radon measures supported on K . The elements of $\mathcal{M}(K)$ can be identified with linear functionals acting on the space of continuous functions $C(K)$, that is, as elements of the dual space $C(K)'$ [10, Corollary 7.18]. The duality pairing of a measure $\mu \in (\mathcal{M}(K))^p$ on a test function $v \in (C(K))^p$ is:

$$\langle \mu, v \rangle = \sum_{i=1}^p \int_K [v]_i(z) d[\mu]_i(z). \quad (1)$$

B. Problem Statement

Consider the control-affine system with feedback control

$$\dot{x}(t) = f(t, x(t)) + g(t, x(t)) u(t, x), \quad (2)$$

with state $x(t) \in \mathbb{R}^n$ and control action $u(t, x) \in \mathbb{R}^m$, such that the components of the vector f and the matrix g are polynomials. Our goal is to find a feedback controller, $u(t, x)$, that maximizes the BRS for a given target set while respecting the *input constraint*

$$u(t, x) \in U = [a_1, b_1] \times \dots \times [a_m, b_m], \quad (3)$$

where $\{a_j\}_{j=1}^m, \{b_j\}_{j=1}^m \subset \mathbb{R}$. Define the *bounding set*, and *target set* as:

$$\begin{aligned} X &= \{x \in \mathbb{R}^n \mid h_{X_i}(x) \geq 0, \forall i = \{1, \dots, n_X\}\}, \\ X_T &= \{x \in \mathbb{R}^n \mid h_{T_i}(x) \geq 0, \forall i = \{1, \dots, n_T\}\}, \end{aligned} \quad (4)$$

respectively, for given polynomials $h_{X_i}, h_{T_i} \in \mathbb{R}[x]$.

Given a finite final time $T > 0$, let the BRS for a particular control policy $u \in L^1([0, T] \times X, U)$, be defined as:

$$\begin{aligned} \mathcal{X}(u) &= \left\{ x_0 \in \mathbb{R}^n \mid \dot{x}(t) = f(t, x(t)) + g(t, x(t)) u(t, x(t)) \right. \\ &\quad \text{a.e. } t \in [0, T], \quad x(0) = x_0, \quad x(T) \in X_T, \\ &\quad \left. x(t) \in X \quad \forall t \in [0, T] \right\}. \end{aligned} \quad (5)$$

$\mathcal{X}(u)$ is the set of initial conditions for solutions² to Equation (2) that remain in the bounding set and arrive in the target set at the final time when control law u is applied. Our aim is to find a controller $u^* \in L^1([0, T] \times X, U)$, that maximizes the volume of the BRS:

$$\lambda(\mathcal{X}(u^*)) \geq \lambda(\mathcal{X}(u)), \quad \forall u \in L^1([0, T] \times X, U), \quad (6)$$

where λ is the Lebesgue measure. u^* need not be unique. We denote the BRS corresponding to u^* by \mathcal{X}^* . To solve this problem, we make the following assumptions:

Assumption 1. X and X_T are compact sets.

Remark 1. Without loss of generality, we assume that $U = \{u \in \mathbb{R}^m \mid -1 \leq u_j \leq 1 \quad \forall j \in \{1, \dots, m\}\}$ (since g can be arbitrarily shifted and scaled). Assumption 1 ensures the existence of a polynomial $h_{X_i}(x) = C_X - \|x\|_2^2$ for a large enough $C_X > 0$.

C. Liouville's Equation

We solve this problem by defining measures over $[0, T] \times X$ whose supports' model the evolution of *families of trajectories*. An initial condition and its relationship with respect to the terminal set can be understood via Equation (2), but the relationship between a family of trajectories and the terminal set must be understood through a different lens. First, define the linear operator $\mathcal{L}_f : C^1([0, T] \times X) \rightarrow C([0, T] \times X)$ on a test function v as:

$$\mathcal{L}_f v = \frac{\partial v}{\partial t} + \sum_{i=1}^n \frac{\partial v}{\partial x_i} [f]_i(t, x), \quad (7)$$

²Solutions in this context are understood in the Carathéodory sense, that is, as absolutely continuous functions whose derivatives satisfy the right hand side of Equation (2) almost everywhere [2, Chapter 10].




Fig. 1: An illustration (left) of trajectories (blue) transforming according to Equation (2) and their corresponding occupation measures (red) at times 0, τ , and T (purple) transforming according to Equation (18).

and its adjoint operator $\mathcal{L}'_f : C([0, T] \times X)' \rightarrow C^1([0, T] \times X)'$ by the adjoint relation:

$$\langle \mathcal{L}'_f \mu, v \rangle = \langle \mu, \mathcal{L}_f v \rangle = \int_{[0, T] \times X} \mathcal{L}_f v(t, x) d\mu(t, x) \quad (8)$$

for all $\mu \in \mathcal{M}([0, T] \times X)$ and $v \in C^1([0, T] \times X)$. Define the linear operator $\mathcal{L}_g : C^1([0, T] \times X) \rightarrow C([0, T] \times X)^m$ as:

$$[\mathcal{L}_g v]_j = \sum_{i=1}^n \frac{\partial v}{\partial x_i} [g]_{ij}(t, x), \quad (9)$$

for each $j \in \{1, \dots, m\}$ and define its adjoint operator $\mathcal{L}'_g : (C([0, T] \times X)^m)' \rightarrow C^1([0, T] \times X)'$ according to its adjoint relation as in Equation (8). Note that $\mathcal{L}_f v(t, x) + (\mathcal{L}_g v(t, x))u(t, x)$ is the time-derivative \dot{v} of a function v .

Given a test function, $v \in C^1([0, T] \times X)$, and an initial condition, $x(0) \in X$, it follows that:

$$v(T, x(T)) = v(0, x(0)) + \int_0^T \dot{v}(t, x(t|x_0)) dt. \quad (10)$$

The traditional approach to designing controllers that stabilize the system imposes Lyapunov conditions on the test functions. However, simultaneously searching for a controller and Lyapunov function results in a nonconvex optimization problem [26]. Instead we examine conditions on the space of measures—the dual to the space of functions—in order to arrive at a convex formulation.

For a fixed control policy $u \in L^1([0, T] \times X, U)$ and an initial condition $x_0 \in \mathbb{R}^n$, let $x(\cdot|x_0) : [0, T] \rightarrow X$ be a solution to Equation (2). Define the *occupation measure* as:

$$\mu(A \times B|x_0) = \int_0^T I_{A \times B}(t, x(t|x_0)) dt, \quad (11)$$

for all subsets $A \times B$ in the Borel σ -algebra of $[0, T] \times X$, where $I_{A \times B}(\cdot)$ denotes the indicator function on a set $A \times B$. This computes the amount of time the graph of the solution, $(t, x(t|x_0))$, spends in $A \times B$. Equation (10) then becomes:

$$v(T, x(T)) = v(0, x(0)) + \int_{[0, T] \times X} (\mathcal{L}_f v(t, x) + \mathcal{L}_g v(t, x)u(t, x)) d\mu(t, x|x_0). \quad (12)$$

When the initial state is not a single point, but is a distribution modeled by an *initial measure*, $\mu_0 \in \mathcal{M}(X)$, we define the *average occupation measure*, $\mu \in \mathcal{M}([0, T] \times X)$ by:

$$\mu(A \times B) = \int_X \mu(A \times B|x_0) d\mu_0(x_0), \quad (13)$$

and the *final measure*, $\mu_T \in \mathcal{M}(X_T)$ by:

$$\mu_T(B) = \int_X I_B(x(T|x_0)) d\mu_0(x_0). \quad (14)$$

Integrating with respect to μ_0 and introducing the initial, average occupation, and final measures, Equation (12) becomes:

$$\int_{X_T} v(T, x) d\mu_T(x) = \int_X v(0, x) d\mu_0(x) + \int_{[0, T] \times X} (\mathcal{L}_f v(t, x) + \mathcal{L}_g v(t, x)u(t, x)) d\mu(t, x). \quad (15)$$

It is useful to think of the measures μ_0 , μ and μ_T as *unnormalized* probability distributions. The support of μ_0 models the set of initial conditions, the support of μ models the flow of trajectories, and the support of μ_T models the set of states at time T .

Next, we subsume $u(t, x)$ into a *signed* measure $\sigma^+ - \sigma^-$ defined by *nonnegative* measures³ $\sigma^+, \sigma^- \in (\mathcal{M}([0, T] \times X))^m$ such that:

$$\int_{A \times B} u_j(t, x) d\mu(t, x) = \int_{A \times B} d[\sigma^+]_j(t, x) - \int_{A \times B} d[\sigma^-]_j(t, x) \quad (16)$$

for all subsets $A \times B$ in the Borel σ -algebra of $[0, T] \times X$ and for each $j \in \{1, \dots, m\}$. This key step allows us to pose an infinite dimensional LP over measures without explicitly parameterizing a control law while allowing us to “back out” a control law using Equation (16). Equation (15) becomes:

$$\langle \mu_T, v(T, \cdot) \rangle = \langle \mu_0, v(0, \cdot) \rangle + \langle \mu, \mathcal{L}_f v \rangle + \langle \sigma^+ - \sigma^-, \mathcal{L}_g v \rangle \quad (17)$$

for all test functions $v \in C^1([0, T] \times X)$. Notice that this substitution renders Equation (17) linear in its measure components. Let δ_t denote the Dirac measure at a point t and \otimes denote the product of measures. Since Equation (17) must hold for all test functions, we obtain a linear operator equation:

$$\mathcal{L}'_f \mu + \mathcal{L}'_g \sigma^+ - \mathcal{L}'_g \sigma^- = \delta_T \otimes \mu_T - \delta_0 \otimes \mu_0, \quad (18)$$

called Liouville’s Equation, which is a classical result in statistical physics that describes the evolution of a density of particles within a fluid [1]. Figure 1 illustrates the evolution of densities according to Liouville’s Equation. This equation is satisfied by families of admissible trajectories starting from the initial distribution μ_0 . The converse statement is true for control affine systems with a convex admissible control set, as we have assumed. We refer the reader to [12, Appendix A] for an extended discussion of Liouville’s Equation.

³Note that we can always decompose a signed measure into unsigned measures as a result of the Jordan Decomposition Theorem [10, Theorem 3.4].

D. BRS via an Infinite Dimensional LP

The goal of this section is to use Liouville's Equation to formulate an infinite dimensional LP, P , that maximizes the size of the BRS, modeled by $\text{spt}(\mu_0)$, for a given target set, modeled by $\text{spt}(\mu_T)$, where $\text{spt}(\mu)$ denotes the support of a measure μ . Slack measures (denoted with "hats") are used to impose the constraints $\lambda \geq \mu_0$ and $\mu \geq [\sigma^+]_j + [\sigma^-]_j$ for each $j \in \{1, \dots, m\}$, where λ is the Lebesgue measure. The former constraint ensures that the optimal value of P is the Lebesgue measure of the largest achievable BRS (see Theorem 2). The latter constraint ensures that we are able to extract a bounded control law by applying Equation (16) (see Theorem 3). Define P as:

$$\begin{aligned} \sup \quad & \mu_0(X) & (P) \\ \text{s.t.} \quad & \mathcal{L}'_f \mu + \mathcal{L}'_g (\sigma^+ - \sigma^-) = \delta_T \otimes \mu_T - \delta_0 \otimes \mu_0, \\ & [\sigma^+]_j + [\sigma^-]_j + [\hat{\sigma}]_j = \mu & \forall j \in \{1, \dots, m\}, \\ & \mu_0 + \hat{\mu}_0 = \lambda, \\ & [\sigma^+]_j, [\sigma^-]_j, [\hat{\sigma}]_j \geq 0 & \forall j \in \{1, \dots, m\}, \\ & \mu, \mu_0, \mu_T, \hat{\mu}_0 \geq 0, \end{aligned}$$

where the given data are f, g, X, X_T and the supremum is taken over a tuple of measures $(\sigma^+, \sigma^-, \hat{\sigma}, \mu, \mu_0, \hat{\mu}_0, \mu_T) \in (\mathcal{M}([0, T] \times X))^m \times (\mathcal{M}([0, T] \times X))^m \times (\mathcal{M}([0, T] \times X))^m \times \mathcal{M}([0, T] \times X) \times \mathcal{M}(X) \times \mathcal{M}(X) \times \mathcal{M}(X_T)$. Given measures that achieve the supremum, the control law that maximizes the size of the BRS is then constructed by finding the $u \in L^1([0, T] \times X, U)$ whose components each satisfy Equation (16) for all subsets in the Borel σ -algebra of $[0, T] \times X$. Before proving that this two-step procedure computes $u^* \in L^1([0, T] \times X, U)$ as in Equation (6), define the dual program to P denoted D as:

$$\begin{aligned} \inf \quad & \int_X w(x) d\lambda(x) & (D) \\ \text{s.t.} \quad & \mathcal{L}_f v + \sum_{i=1}^m [p]_i \leq 0, \\ & [p]_i \geq 0, \quad [p]_i \geq |[\mathcal{L}_g v]_i| & \forall i = \{1, \dots, m\}, \\ & w \geq 0, \\ & w(x) \geq v(0, x) + 1 & \forall x \in X, \\ & v(T, x) \geq 0 & \forall x \in X_T \end{aligned}$$

where the given data are f, g, X, X_T and the infimum is taken over $(v, w, p) \in C^1([0, T] \times X) \times C(X) \times (C([0, T] \times X))^m$. The dual allows us to obtain approximations of the BRS \mathcal{X}^* (see Theorem 4).

Theorem 1. *There is no duality gap between P and D .*

Theorem 2. *The optimal value of P is equal to $\lambda(\mathcal{X}^*)$, the Lebesgue measure of the BRS of the controller defined by Equation (16).*

The solution to P can be used in order to construct the control law that maximizes the BRS:

Theorem 3. *There exists a control law, $\tilde{u} \in L^1([0, T] \times X, U)$, that satisfies Equation (16) when substituting in the vector of measures that achieves the supremum of P ,*

$(\sigma^{+*}, \sigma^{-*}, \hat{\sigma}^*, \mu^*, \mu_0^*, \hat{\mu}_0^*, \mu_T^*)$, and is the control law that maximizes the size of the BRS, i.e. $\lambda(\mathcal{X}(\tilde{u})) \geq \lambda(\mathcal{X}(u))$, $\forall u \in L^1([0, T] \times X, U)$. Moreover, any two control laws constructed by applying Equation (16) to the vector of measures that achieves the supremum of P are equal μ^* -almost everywhere.

Next, we note that the w -component to a feasible point of D is an outer approximation to \mathcal{X}^* .

Theorem 4. *\mathcal{X}^* is a subset of $\{x \mid w(x) \geq 1\}$, for any feasible w of the D . Furthermore, there is a sequence of feasible solutions to D such that the w -component converges from above to $I_{\mathcal{X}^*}$ in the L^1 norm and almost uniformly.*

III. NUMERICAL IMPLEMENTATION

The infinite dimensional problems P and D are not directly amenable to computation. However, a sequence of finite dimensional approximations in terms of semidefinite programs (SDPs) can be obtained by characterizing measures in P by their moments, and restricting the space of functions in D to polynomials. The solutions to each of the SDPs in this sequence can be used to construct controllers and outer approximations that converge to the solution of the infinite dimensional LP. A comprehensive introduction to such *moment relaxations* can be found in [15].

Measures on the set $[0, T] \times X$ are completely determined by their action (via integration) on a dense subset of the space $C^1([0, T] \times X)$ [10]. Since $[0, T] \times X$ is compact, the Stone-Weierstrass Theorem [10, Theorem 4.45] allows us to choose the set of polynomials as this dense subset. Every polynomial on \mathbb{R}^n , say $p \in \mathbb{R}[x]$ with $x = (x_1, \dots, x_n)$, can be expanded in the monomial basis via

$$p(x) = \sum_{\alpha \in \mathbb{N}^n} p_\alpha x^\alpha,$$

where $\alpha = (\alpha_1, \dots, \alpha_n)$ ranges over vectors of non-negative integers, $x^\alpha = x_1^{\alpha_1} \dots x_n^{\alpha_n}$, and $\text{vec}(p) = (p_\alpha)_{\alpha \in \mathbb{N}^n}$ is the vector of coefficients of p . By definition, the p_α are real and only finitely many are non-zero. We define $\mathbb{R}_k[x]$ to be those polynomials such that p_α is non-zero only for $\alpha \in \mathbb{N}_k^n$. The degree of a polynomial, $\text{deg}(p)$, is the smallest k such that $p \in \mathbb{R}_k[x]$.

The moments of a measure μ defined over a real n -dimensional space are given by:

$$y_\mu^\alpha = \int x^\alpha d\mu(x). \quad (19)$$

Integration of a polynomial with respect to a measure ν can be expressed as a linear functional of its coefficients:

$$\langle \mu, p \rangle = \int p(x) d\mu(x) = \sum_{\alpha \in \mathbb{N}^n} p_\alpha y_\mu^\alpha = \text{vec}(p)^T y_\mu. \quad (20)$$

Integrating the square of a polynomial $p \in \mathbb{R}_k[x]$, we obtain:

$$\int p(x)^2 d\mu(x) = \text{vec}(p)^T M_k(y_\mu) \text{vec}(p), \quad (21)$$

where $M_k(y_\mu)$ is the *truncated moment matrix* defined by

$$[M_k(y_\mu)]_{(\alpha, \beta)} = y_\mu^{\alpha+\beta} \quad (22)$$

for $\alpha, \beta \in \mathbb{N}_k^n$. Note that for any positive measure μ , the matrix $M_k(y_\mu)$ must be positive semidefinite. Similarly, given $h \in \mathbb{R}[x]$ with $(h_\gamma)_{\gamma \in \mathbb{N}^n} = \text{vec}(h)$ one has

$$\int p(x)^2 h(x) d\mu(x) = \text{vec}(p)^T M_k(h, y_\mu) \text{vec}(p), \quad (23)$$

where $M_k(h, y)$ is a *localizing matrix* defined by

$$[M_k(h, y_\mu)]_{(\alpha, \beta)} = \sum_{\gamma \in \mathbb{N}^n} h_\gamma y_\mu^{\alpha + \beta} \quad (24)$$

for all $\alpha, \beta \in \mathbb{N}_k^n$. The localizing and moment matrices are symmetric and linear in the moments y .

A. Approximating Problems

Finite dimensional SDPs approximating P can be obtained by replacing constraints on measures with *constraints on moments*. All of the equality constraints of P can be expressed as an infinite dimensional linear system of equations which the moments of the measures appearing in P must satisfy. This linear system is obtained by restricting to polynomial test functions (which we note are sufficient given our discussion above): $v(t, x) = t^\alpha x^\beta$, $[p]_j(t, x) = t^\alpha x^\beta$, and $w(x) = x^\beta$, $\forall \alpha \in \mathbb{N}, \beta \in \mathbb{N}^n$. For example, the equality constraint corresponding to Liouville's Equation is obtained by examining:

$$\begin{aligned} 0 = & \int_{[0, T] \times X} \mathcal{L}_f(t^\alpha x^\beta) d\mu(t, x) + \int_{[0, T] \times X} \mathcal{L}_g(t^\alpha x^\beta) d[\sigma^+]_j(t, x) \\ & - \int_{[0, T] \times X} \mathcal{L}_g(t^\alpha x^\beta) d[\sigma^-]_j(t, x) - \int_{X_T} T^\alpha x^\beta d\mu_T(x) + \int_X x^\beta d\mu_0(x). \end{aligned}$$

A *finite dimensional* linear system is obtained by truncating the degree of the polynomial test functions to $2k$. Let $\Gamma = \{\sigma^+, \sigma^-, \hat{\sigma}, \mu, \mu_0, \hat{\mu}_0, \mu_T\}$, then let $\mathbf{y}_k = (y_{k, \gamma}) \subset \mathbb{R}$ be a vector of sequences of moments truncated to degree $2k$ for each $\gamma \in \Gamma$. The finite dimensional linear system is then represented by the linear system:

$$A_k(\mathbf{y}_k) = b_k. \quad (25)$$

Constraints on the support of the measures also need to be imposed (see [15] for details). Let the k -th relaxed SDP representation of P , denoted P_k , be defined as:

$$\begin{aligned} \sup & y_{k, \mu_0}^0 & (P_k) \\ \text{s.t.} & A_k(\mathbf{y}_k) = b_k, \\ & M_k(y_{k, \gamma}) \succeq 0 & \forall \gamma \in \Gamma, \\ & M_{k_{X_i}}(h_{X_i}, y_{k, \gamma}) \succeq 0 & \forall (i, \gamma) \in \{1, \dots, n_X\} \times \Gamma \setminus \mu_T, \\ & M_{k_{T_i}}(h_{T_i}, y_{k, \mu_T}) \succeq 0 & \forall i \in \{1, \dots, n_T\}, \\ & M_{k-1}(h_\tau, y_{k, \gamma}) \succeq 0 & \forall \gamma \in \Gamma \setminus \{\mu_0, \mu_T, \hat{\mu}_0\}, \end{aligned}$$

where the given data are f, g, X, X_T and the supremum is taken over the sequence of moments, $\mathbf{y}_k = (y_{k, \gamma})$, $h_\tau = t(T - t)$, $k_{X_i} = k - \lceil \deg(h_{X_i})/2 \rceil$, $k_{T_i} = k - \lceil \deg(h_{T_i})/2 \rceil$, and $\succeq 0$ denotes positive semi-definiteness. For each $k \in \mathbb{N}$, let \mathbf{y}_k^* denote the optimizer of P_k , with components $y_{k, \gamma}^*$ where $\gamma \in \Gamma$ and let p_k^* denote the supremum of P_k .

The dual of P_k is a sums-of-squares (SOS) program denoted D_k for each $k \in \mathbb{N}$, which is obtained by first restricting the optimization space in the D to the polynomial functions with degree truncated to $2k$ and by then replacing the non-negativity constraint D with a *sums-of-squares* constraint [25]. Define $Q_{2k}(h_{X_1}, \dots, h_{X_{n_X}}) \subset \mathbb{R}_{2k}[x]$ to be the set of polynomials $q \in \mathbb{R}_{2k}[x]$ (i.e. of total degree less than $2k$) expressible as

$$q = s_0 + \sum_{i=1}^{n_X} s_i h_{X_i}, \quad (26)$$

for some polynomials $\{s_i\}_{i=0}^{n_X} \subset \mathbb{R}_{2k}[x]$ that are sums of squares of other polynomials. Every such polynomial is clearly non-negative on X . Define $Q_{2k}(h_\tau, h_{X_1}, \dots, h_{X_{n_X}}) \subset \mathbb{R}_{2k}[t, x]$ and $Q_{2k}(h_{T_1}, \dots, h_{T_{n_T}}) \subset \mathbb{R}_{2k}[x]$, similarly. Employing this notation, the k -th relaxed SDP representation of D , denoted D_k , is defined as:

$$\begin{aligned} \inf & l^T \text{vec}(w) & (D_k) \\ \text{s.t.} & -\mathcal{L}_f v - \mathbf{1}^T p \in Q_{2k}(h_\tau, h_{X_1}, \dots, h_{X_{n_X}}), \\ & p - (\mathcal{L}_g v)^T \in (Q_{2k}(h_\tau, h_{X_1}, \dots, h_{X_{n_X}}))^m, \\ & p + (\mathcal{L}_g v)^T \in (Q_{2k}(h_\tau, h_{X_1}, \dots, h_{X_{n_X}}))^m, \\ & w \in Q_{2k}(h_{X_1}, \dots, h_{X_{n_X}}), \\ & w - v(0, \cdot) - 1 \in Q_{2k}(h_{X_1}, \dots, h_{X_{n_X}}), \\ & v(T, \cdot) \in Q_{2k}(h_{T_1}, \dots, h_{T_{n_T}}), \end{aligned}$$

where the given data are f, g, X, X_T , the infimum is taken over the vector of polynomials $(v, w, p) \in \mathbb{R}_{2k}[t, x] \times \mathbb{R}_{2k}[x] \times (\mathbb{R}_{2k}[t, x])^m$, and l is a vector of moments associated with the Lebesgue measure (i.e. $\int_X w d\lambda = l^T \text{vec}(w)$ for all $w \in \mathbb{R}_{2k}[x]$). For each $k \in \mathbb{N}$, let d_k^* denote the infimum of D_k .

Theorem 5. *For each $k \in \mathbb{N}$, there is no duality gap between P_k and D_k .*

Next, we construct a technique to extract a polynomial control law from the solution \mathbf{y}_k of P_k . Given moment sequences truncated to degree $2k$, one can choose an approximate control law u_k with components $[u_k]_j \in \mathbb{R}_k[t, x]$ so that the truncated analogue of Equation (16) is satisfied. That is, by requiring:

$$\int_{[0, T] \times X} t^{\alpha_0} x^\alpha [u_k]_j(t, x) d\mu(t, x) = \int_{[0, T] \times X} t^{\alpha_0} x^\alpha d[\sigma^+ - \sigma^-]_j, \quad (27)$$

for (α_0, α) satisfying $\sum_{i=0}^n \alpha_i \leq k$. When constructing a polynomial control law from the solution of P_k , these linear equations written with respect to the coefficients of $[u_k]_j$ are expressible in terms of y_{k, σ^+}^* , y_{k, σ^-}^* , and $y_{k, \mu}^*$. Direct calculation shows the linear system of equations is:

$$M_k(y_{k, \mu}^*) \text{vec}([u_k]_j) = y_{k, [\sigma^+]_j}^* - y_{k, [\sigma^-]_j}^*. \quad (28)$$

B. Convergence of Approximating Problems

Next, we prove the convergence properties of P_k and D_k and the corresponding controllers. We begin by proving that the polynomial w approximates the indicator function of the set \mathcal{X}^* . As we increase k , this approximation gets tighter. The following theorem makes this statement precise.

Theorem 6. For each $k \in \mathbb{N}$, let $w_k \in \mathbb{R}_{2k}[x]$ denote the w -component of the solution to D_k , and let $\bar{w}_k(x) = \min_{i \leq k} w_i(x)$. Then, w_k converges from above to $I_{\mathcal{X}^*}$ in the L^1 norm, and $\bar{w}_k(x)$ converges from above to $I_{\mathcal{X}^*}$ in the L^1 norm and almost uniformly.

Corollary 1. $\{d_k^*\}_{k=1}^\infty$ and $\{p_k^*\}_{k=1}^\infty$ converge monotonically from above to the optimal value of D and P .

Next, we prove that the 1-superlevel set of the polynomial w converges in Lebesgue measure to the largest achievable BRS \mathcal{X}^* .

Theorem 7. For each $k \in \mathbb{N}$, let $w_k \in \mathbb{R}_{2k}[x]$ denote the w -component of the solution to D_k , and let $\mathcal{X}_k := \{x \in \mathbb{R}^n : w_k(x) \geq 1\}$. Then, $\lim_{k \rightarrow \infty} \lambda(\mathcal{X}_k \setminus \mathcal{X}^*) = 0$.

Finally, we prove a convergence result for the sequence of controllers generated by (28). For each $k \in \mathbb{N}$, let u_k^* denote the controller constructed by Equation (28) using the optimizers y_k of P_k . Let $y_{k,\mu}^*$ be the optimizing moment sequence corresponding to μ .

Theorem 8. Let $\{\mu_k^*\}_{k=1}^\infty$ be any sequence of measures such that the truncated moments of μ_k^* match $y_{k,\mu}^*$. Then, there exists an optimizing vector of measures $(\sigma^{+*}, \sigma^{-*}, \hat{\sigma}^*, \mu^*, \mu_0^*, \hat{\mu}_0^*, \mu_T^*)$ for P , a $u^* \in L^1([0, T] \times X)$ generated using σ^{+*}, σ^{-*} , and μ^* according to Equation (16), and a subsequence $\{k_i\}_{i=1}^\infty \subset \mathbb{N}$ such that:

$$\int_{[0, T] \times X} v(t, x) ([u_{k_i}^*]_j(t, x) d\mu_{k_i}^*(t, x) - [u^*]_j(t, x) d\mu^*(t, x)) \xrightarrow{i \rightarrow \infty} 0, \quad (29)$$

for all $v \in C^1([0, T] \times X)$, and each $j \in \{1, \dots, m\}$.

IV. EXAMPLES

This section provides a series of numerical experiments on example systems of increasing complexity. SDPs were prepared using a custom software toolbox and the modeling tool YALMIP [17]. For simulations, control laws are taken to be the saturation of the polynomial law derived by the proposed method. The programs, whose source code is available for download⁴, are solved using SeDuMi 1.3 [29], for the first three examples, and the SDPT3 solver [31], for the last example, on a machine with 8 Intel Xeon processors with a clock speed of 3.1 GHz and 32 GB RAM.

Additionally, for several of the examples we examine a different objective wherein we look to drive initial conditions starting in X to X_T at any time $t \in [0, T]$ (commonly referred to as a “free final time” problem). The analogous BRS approximation problem is addressed in [12] and the control synthesis problem follows using our approach in a straightforward manner. We make clear when we employ this different objective while describing each of our examples.

A. Double Integrator

The double integrator is a two state, single input system given by $\dot{x}_1 = x_2$, $\dot{x}_2 = u$, with u restricted to the interval $U = [-1, 1]$. Setting the target set to the origin, $X_T = \{0\}$,




Fig. 2: An illustration of the convergence of outer approximations and the performance of controllers designed using our approach for increasing truncation degree, k , for the double integrator. Solid lines indicate the outer approximations, defined by $w_k = 1$, for $k = 2$ (red), $k = 3$ (green), $k = 4$ (blue), and the boundary of the true BRS (solid black). Points indicate terminal states $x(T)$ of controlled solutions with $x(0)$ inside the BRS using our generated feedback control laws u_k (colors match the outer approximations).

the optimal BRS \mathcal{X}^* can be computed analytically based on a minimum time “bang-bang controller” [4, pp. 136]. Note that this is a challenging system for grid based optimal control methods, since they require high resolution near the switching surface of the “bang-bang” control law.

We take the bounding set to be $X = \{x \mid \|x\|^2 \leq 1.6^2\}$. Figure 2 compares the outer approximations of \mathcal{X}^* for $k = 2, 3, 4$. The quality of the approximations increases quickly. Figure 2 also evaluates the performance of the control laws u_k by plotting the terminal states $x(T)$ for controlled solutions starting in \mathcal{X}^* . Even for $k = 3$, reasonable performance is achieved. The running times for $k = 2, 3, 4$ are 0.3s, 0.7s, and 4.2s, respectively.

B. Ground Vehicle Model

The “Dubin’s car” [9] is a popular model for autonomous ground and air vehicles and has been employed in a wide variety of applications [5], [6], [11]. Its dynamics are:

$$\dot{a} = v \cos(\theta), \quad \dot{b} = v \sin(\theta), \quad \dot{\theta} = \omega, \quad (30)$$

where the states are the x-position (a), y-position (b) and yaw angle (θ) of the vehicle and the control inputs are the forward speed (v) and turning rate (ω). A change of coordinates can be applied to this system in order to make the dynamics polynomial [7]. The rewritten dynamics are given by:

$$\dot{x}_1 = u_1, \quad \dot{x}_2 = u_2, \quad \dot{x}_3 = x_2 u_1 - x_1 u_2. \quad (31)$$

This system is also known as the Brockett integrator and is a popular benchmark since it is prototypical of many nonholonomic systems. Notice that the system has an uncontrollable linearization and does not admit a smooth time-invariant control law that makes the origin asymptotically stable [7]. Hence, this example illustrates the advantage of our method when compared to linear control synthesis techniques. We solve the “free final time” problem to construct a time-varying control law that drives the initial conditions in $X = \{x \mid \|x\|^2 \leq 4\}$ to the target set $X_T = \{x \mid \|x\|^2 \leq 0.1^2\}$ by time $T = 4$. In

⁴<https://groups.csail.mit.edu/locomotion/software.html>




Fig. 3: The boundary of X (black line) and the outer approximation of the BRS (blue line) are shown in the (x_1, x_3) and (x_1, x_2) planes. In each plane, black points indicate initial conditions of controlled solutions with $x(T) \in X_T$ (i.e. $\|x(T)\|^2 \leq 0.1^2$), and grey points indicate initial conditions of solutions ending near the target set (specifically $\|x(T)\|^2 \leq 0.2^2$).




Fig. 4: A pair of sample trajectories drawn in red generated by our algorithm for the Dubin's car system. Each drawn time sample of the car is colored black with blue-colored tires and grey forward-facing direction. The origin is marked by a green dot. Each trajectory is initialized at the 'O' mark and terminates at the state where the target set is first reached, which is marked with an 'X' mark. The trajectory starting in the upper right hand corner executes a three-point turn to arrive at the desired position and orientation.

the Dubin's car coordinates, the target set is a neighborhood of the origin while being oriented in the positive a -direction. The control is restricted to $u_1, u_2 \in [-1, 1]$. Figure 3 plots outer approximations of the BRS for $k = 5$. Figure 4 illustrates two sample trajectories generated using a feedback controller designed by our algorithm after transforming back to the original coordinate system. Solving the SDP took 599 seconds.

C. Torque Limited Simple Pendulum

Next, we consider the torque limited simple pendulum, described by the equations

$$\dot{x}_1 = x_2, \quad I\dot{x}_2 = mgl \sin(x_1) - bx_2 + u, \quad (32)$$

where x_1 represents the angle θ from upright, x_2 represents the angular rate $\dot{\theta}$, and u represents a torque source at the pivot constrained to take values in $U = [-3, 3]$. We take $m = 1$, $l = 0.5$, $I = ml^2$, $b = 0.1$ and $g = 9.8$. The bounding set is defined by $x_1 \in [-\pi, \pi]$ and $x_2 \in [-8, 8]$. Our method can be readily adapted to handle dynamics with trigonometric dependence on a state, x , so long as the dynamics are polynomial in $\sin(x)$ and $\cos(x)$. This is accomplished by introducing indeterminates c and s identified with $\sin(x)$ and $\cos(x)$ and modifying the approach to work over the quotient ring associated with the equation $1 = c^2 + s^2$ [24].

For this example, we solve the "free final time" problem by taking $T = 1.5$, and defining the target set as $X_T =$




Fig. 5: A depiction of the controller performance for $k = 5$ for the torque limited simple pendulum. Black regions indicate sampled initial conditions whose controlled solutions pass through the target set (yellow square). Three sample solutions are also plotted (red) each with terminal conditions in the target set. Note the solution starting near $(-2, 3.2)$ passes through zero velocity — the solution "pumps" to reach the upright position.

$\{(x_1, x_2) \mid \cos(x_1) \geq 0.95, x_2^2 \leq 0.05\}$. The running time for the SDP is 11 mins 20 secs for $k = 5$. Figure 5 plots sample solutions and summarizes the initial conditions that reach the target set. Notice that the controller is able to "swing-up" states close to the downright position to the upright configuration despite the stringent actuator limits and a short time-horizon.

D. Planar Quadrotor

Finally, we demonstrate the scalability of our approach on a six state, two input planarized quadrotor model used in a various robotic applications [13], [18]. The dynamics are defined by [28]:

$$\begin{aligned} \ddot{x}_1 &= -(u_1 + u_2) \sin(\theta)/m \\ \ddot{x}_2 &= -g + (u_1 + u_2) \cos(\theta)/m \\ \ddot{\theta} &= L(u_2 - u_1)/I \end{aligned} \quad (33)$$

where x_1, x_2 , and θ are the horizontal and vertical positions, and the attitude of the quadrotor, respectively. The control inputs u_1 and u_2 are the force produced by the left and right rotors, respectively, and are bounded to have a thrust to weight ratio of 2.5. Further, $L = 0.25$ is the length of the rotor arm, $m = 0.486$ is the mass, $I = 0.00383$ is the moment of inertia and $g = 9.8$ is the acceleration due to gravity. Using a time horizon of $T = 4$, we solve a "free final time" problem and require trajectories reach the target set $X_T = \{x \mid \|x\|^2 \leq 0.1\}$. The bounding set is $X = \{x \mid \|x\|^2 \leq 1\}$. We apply the proposed control design method with $k = 2$, and handle trigonometric terms in the same manner as the pendulum example. The SDP takes 49 minutes to solve. The resulting controller is able to stabilize a large set of initial conditions. Figure 6 illustrates a few representative trajectories of the closed-loop system.

V. CONCLUSION

We presented an approach for designing feedback controllers that maximize the size of the BRS by posing an infinite dimensional LP over the space of non-negative measures. Finite dimensional approximations to this LP in terms of SDPs can be used to obtain outer approximations of the largest




Fig. 6: Four trajectories drawn in red generated by our algorithm for the planar quadrotor. Each drawn time sample of the car is colored black with props on the upward-facing direction. The origin is marked by a green dot.

achievable BRS and polynomial control laws that approximate the optimal control law. In contrast to previous approaches relying on Lyapunov’s stability criteria, our method is inherently convex and does not require feasible initialization. The proposed method can be used to augment existing feedback motion planning techniques that rely on sequencing together BRSs in order to drive some desired set of initial conditions to a given target set. The number of distinct controllers required by such algorithms could be significantly reduced (potentially down to a single feedback law) by using our algorithm. By reasoning about the nonlinear dynamics of a robotic system, our algorithm should also be able to obtain improved performance during dynamic tasks while maintaining robustness.

We are presently pursuing convergence results that guarantee the set-wise convergence of the BRS of the controllers generated via (28) to the largest achievable BRS, which is stronger than the result in Theorem 8. We are also working to extend our method to hybrid dynamical systems. Our approach potentially can address the difficulties that linearization based approaches face due to the inherent nonlinearities associated with hybrid systems such as walking robots.

ACKNOWLEDGEMENTS

The authors are grateful to Milan Korda for many helpful discussions. This work was supported by ONR MURI grant N00014-09-1-1051, NSF Contract IIS-1161679 and the Siebel Foundation.

REFERENCES

- [1] V. Arnold. *Mathematical Methods of Classical Mechanics*, volume 60. Springer, 1989.
- [2] J. Aubin and H. Frankowska. *Set-Valued Analysis*. Birkhauser Boston, 2008.
- [3] A. Bemporad, F. Borrelli, and M. Morari. Model predictive control based on linear programming—the explicit solution. *IEEE Transactions on Automatic Control*, 47(12):1974–1985, 2002.
- [4] D. P. Bertsekas. *Dynamic Programming & Optimal Control*, volume I. 3rd edition, 2005.
- [5] A. Bhatia, M. Graziano, S. Karaman, R. Naldi, and E. Frazzoli. Dubins trajectory tracking using commercial off-the-shelf autopilots. In *AIAA Guidance, Navigation, and Control Conference*, Honolulu, Hawaii, August 2008.
- [6] Q. Chen and U. Ozguner. Intelligent off-road navigation algorithms and strategies of team desert buckeyes in the darpa grand challenge05. In *The 2005 DARPA Grand Challenge*, pages 183–203. Springer, 2007.
- [7] D. DeVon and T. Bretl. Kinematic and dynamic control of a wheeled mobile robot. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4065–4070. IEEE, 2007.

- [8] J. Ding and C. Tomlin. Robust reach-avoid controller synthesis for switched nonlinear systems. In *49th IEEE Conference on Decision and Control*, pages 6481–6486. IEEE, 2010.
- [9] L. Dubins. On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents. *American Journal of Mathematics*, 79(3):497–516, 1957.
- [10] G. B. Folland. *Real Analysis: Modern Techniques and Their Applications*, volume 2. John Wiley & Sons, second edition, 1999.
- [11] A. Gray, Y. Gao, T. Lin, K. J. Hedrick, H. E. Tseng, and F. Borrelli. Predictive control for agile semi-autonomous ground vehicles using motion primitives. In *American Control Conference (ACC)*, 2012, pages 4239–4244, June 2012.
- [12] D. Henrion and M. Korda. Convex computation of the region of attraction of polynomial control systems. *arXiv preprint arXiv:1208.1751*, 2012.
- [13] G. M. Hoffmann, H. Huang, S. L. Waslander, and C. J. Tomlin. Quadrotor helicopter flight dynamics and control: Theory and experiment. In *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, pages 1–20, 2007.
- [14] Z. Jarvis-Wloszek, R. Feeley, W. Tan, K. Sun, and A. Packard. Control applications of sum of squares programming. In D. Henrion and A. Garulli, editors, *Positive Polynomials in Control*, volume 312 of *Lecture Notes in Control and Information Sciences*, pages 3–22. Springer Berlin / Heidelberg, 2005.
- [15] J. B. Lasserre. *Moments, Positive Polynomials and Their Applications*, volume 1. World Scientific, 2010.
- [16] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, 2006.
- [17] J. Löfberg. YALMIP : A toolbox for modeling and optimization in MATLAB. In *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004.
- [18] S. Lupashin, A. Schollig, M. Sherback, and R. D’Andrea. A simple learning strategy for high-speed quadcopter multi-flips. In *2010 IEEE International Conference on Robotics and Automation*, pages 1642–1648. IEEE, 2010.
- [19] A. Majumdar, A. A. Ahmadi, and R. Tedrake. Control design along trajectories with sums of squares programming. In *Proceedings of the 2013 IEEE International Conference on Robotics and Automation (ICRA)*, 2013.
- [20] A. Majumdar and R. Tedrake. Robust online motion planning with regions of finite time invariance. In *Proceedings of the Workshop on the Algorithmic Foundations of Robotics*, 2012.
- [21] I. Mitchell, A. Bayen, and C. Tomlin. A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games. *IEEE Transactions on Automatic Control*, 50(7):947–957, July 2005.
- [22] J. Moore and R. Tedrake. Powerline perching with a fixed-wing UAV. In *Proceedings of the AIAA Infotech@Aerospace Conference*, Seattle, WA, April 2009. AIAA.
- [23] R. Munos and A. Moore. Variable resolution discretization in optimal control. *Machine Learning*, 49(2/3):291–323, November/December 2002.
- [24] P. Parrilo. Exploiting structure in sum of squares programs. In *42nd IEEE Conference on Decision and Control*, volume 5, pages 4664–4669, December 2003.
- [25] P. A. Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, California Institute of Technology, May 18 2000.
- [26] S. Prajna, P. Parrilo, and A. Rantzer. Nonlinear control synthesis by convex optimization. *IEEE Transactions on Automatic Control*, 49(2):310–314, February 2004.
- [27] S. Sastry. *Nonlinear Systems: Analysis, Stability, and Control*, volume 10. Springer, 1999.
- [28] J. Steinhardt and R. Tedrake. Finite-time regional verification of stochastic nonlinear systems. *International Journal of Robotics Research*, 31(7):901–923, June 2012.
- [29] J. F. Sturm. Using SeDuMi 1.02, a Matlab toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 11(1-4):625 – 653, 1999.
- [30] R. Tedrake, I. R. Manchester, M. M. Tobenkin, and J. W. Roberts. LQR-Trees: Feedback motion planning via sums of squares verification. *International Journal of Robotics Research*, 29:1038–1052, July 2010.
- [31] R. Tutunçu, K. Toh, and M. Todd. Solving semidefinite-quadratic-linear programs using sdpt3. *Mathematical programming*, 95(2):189–217, 2003.