Systems & Control Letters 86 (2015) 34–40

Contents lists available at ScienceDirect

Systems & Control Letters

journal homepage: www.elsevier.com/locate/sysconle

A perspective-based convex relaxation for switched-affine optimal control

Nicholas Moehle [a,*], Stephen Boyd [b]

[a] Mechanical Engineering Department, Stanford University, United States

[b] Electrical Engineering Department, Stanford University, United States

ARTICLE INFO

Article history:
Received 8 December 2014
Received in revised form
13 May 2015
Accepted 2 September 2015
Available online 3 November 2015

Keywords:
Switched-affine systems
Optimal control
Mixed-integer convex programming
Hybrid systems
Disjunctive programming

ABSTRACT

We consider the switched-affine optimal control problem, *i.e.*, the problem of selecting a sequence of affine dynamics from a finite set in order to minimize a sum of convex functions of the system state. We develop a new reduction of this problem to a mixed-integer convex program (MICP), based on perspective functions. Relaxing the integer constraints of this MICP results in a convex optimization problem, whose optimal value is a lower bound on the original problem value. We show that this bound is at least as tight as similar bounds obtained from two other well-known MICP reductions (via conversion to a mixed logical dynamical system, and by generalized disjunctive programming), and our numerical study indicates it is often substantially tighter. Using simple integer-rounding techniques, we can also use our formulation to obtain an upper bound (and corresponding sequence of control inputs). In our numerical study, this bound was typically within a few percent of the optimal value, making it attractive as a stand-alone heuristic, or as a subroutine in a global algorithm such as branch and bound. We conclude with some extensions of our formulation to problems with switching costs and piecewise affine dynamics.

1. Switched-affine control

A *switched-affine system* has the form

$$x_{t+1} = A^{u_t} x_t + b^{u_t}, \quad t = 0, 1, \ldots,$$

where $x_t \in \mathbf{R}^n$ is the state at time t, $u_t \in \{1, \ldots, K\}$ is the control input at time t, and $A^1, \ldots, A^K$ and $b^1, \ldots, b^K$ are given matrices and vectors. At each time period, the control input selects from a given finite set of affine dynamics. We assume, without loss of generality, that $(A^i, b^i) \neq (A^j, b^j)$ for $i \neq j$. Switched-affine systems arise in various engineering applications, for example as models of switched-mode power supplies and power conversion circuits.

The switched-affine control problem is

$$\begin{array}{ll} \text{minimize} & \displaystyle\sum_{t=0}^{T} g_t(x_t) \\ \text{subject to} & x_{t+1} = A^{u_t} x_t + b^{u_t} \\ & u_t \in \{1, \ldots, K\}, \end{array} \tag{1}$$

where the constraints must hold for $t = 0, \ldots, T-1$. The problem variables are the system states $x_0, \ldots, x_T \in \mathbf{R}^n$ and the control inputs $u_0, \ldots, u_{T-1}$. The problem parameters are the dynamics (A^i, b^i) for $i = 1, \ldots, K$ and the stage cost functions $g_0, \ldots, g_T$. We assume the stage cost functions $g_t : \mathbf{R}^n \to \mathbf{R} \cup \{\infty\}$ are convex and extended valued, which allows us to represent convex state constraints in the stage cost function. We define the state constraint set as $\mathcal{X}_t = \{x \mid g_t(x) < \infty\}$, so the objective is infinite unless $x_t \in \mathcal{X}_t$ holds for $t = 0, \ldots, T$. We can use g_0 to encode a given initial condition, so that $\mathcal{X}_0 = \{x_{\text{init}}\}$, for some $x_{\text{init}} \in \mathbf{R}^n$.

The switched-affine control problem (1) is NP-hard in general (this is proven by Egerstedt and Blondel [1] for a special case), and can be solved globally only at great computational cost in the worst-case. However, by reformulating it as a mixed-integer convex program (MICP), lower bounds on the optimal value can be obtained by relaxing the integer constraints, and upper bounds can be obtained by applying an integer-rounding heuristic to the relaxed solution. These bounds can be used as the basis for a global solver (using, *e.g.*, branch and bound), or alternatively, the rounding procedure can be used as a heuristic to produce a good, if not optimal, sequence of control inputs. The success of both methods (*i.e.*, the run-time of a global solution algorithm, or the quality of the heuristic control sequence) depends crucially on the MICP reformulation (and the tightness of the bounds it produces).

* Corresponding author.
E-mail addresses: moehle@stanford.edu (N. Moehle), boyd@stanford.edu (S. Boyd).

http://dx.doi.org/10.1016/j.sysconle.2015.09.002
0167-6911/

In this paper, we give a new MICP formulation that achieves better bounds than those obtained from other popular reformulation techniques. Although we focus on the specific problem given in (1), we give some extensions of our approach to some related problems in Section 6.

1.1. Previous work

1.1.1. Switched-affine control

Many approaches exist for optimal control of switched systems; a summary can be found in [2]. Here we mention some particularly relevant techniques.

Mixed logical dynamical systems. Switched-affine systems are a special case of *hybrid systems*, *i.e.*, systems involving continuous and logical dynamics. A standard approach to solve (1), proposed by Torrisi and Bemporad [3], is to first convert the switched-affine system into an equivalent *mixed logical dynamical* (MLD) system, which expresses the system using a combination of linear and binary constraints on the original variables and some auxiliary variables (see [4] for details on MLD systems). Minimizing a sum of convex functions of the system states can therefore be expressed as an MICP. We will call this the MLD approach to solving (1), and will briefly describe it in Section 4.1.

Disjunctive programming. Problem (1) can be cast as a *disjunctive program*, *i.e.*, an optimization problem in which the decision variables must lie in the union of some sets (see [5]). Ceria and Soares [6] show that minimizing a convex function over the union of convex sets can be equivalently formulated as an MICP, using lifted variables and perspective functions. This technique has seen much application in process engineering (see, *e.g.*, [7]); for some other applications, see [8]. Several works apply disjunctive programming to switched-affine optimal control; the first appears to be by Stursberg and Panek [9]; we refer to this approach as the GDP formulation, and we describe it in Section 4.2. Oldenburg and Marquardt [10,11] give a detailed account of how to formulate complex switched dynamic constraints using a disjunctive programming framework. Disjunctive programming techniques have also been suggested for deriving mixed logical dynamical systems; see [12]. Several strategies for finding an upper bounds, some with guaranteed suboptimality bounds, can be found in the work of Sager, Jung, and Reinelt [13,14].

Approximate dynamic programming. Wang, O'Donoghue, and Boyd [15] give a method for obtaining relaxations for several hard optimal control problems, including switched-affine systems. The bounds are obtained by maximizing a quadratic approximate value function, evaluated over some initial state distribution, while constraining it to be an under-estimator of the true value function (using a chain of Bellman inequalities).

1.2. Convex optimization

Convex optimization problems can be solved efficiently and reliably using standard techniques [16, Ch. 1]. In practice, this is often done by representing the functions involved in terms of a few standard convex cones, then using a conic optimization solver. Typical cones used in convex optimization include the positive orthant, second-order cone, semidefinite cone, exponential cone, and combinations thereof. Many functions and constraints are representable in terms of these cones; several examples are given in [17–20].

Mixed-integer optimization problems that are convex if the integrality constraints are relaxed are called *mixed-integer convex programs* (MICPs). Although mixed-integer convex programming is NP-hard, these problems can, in principle, be solved using simple branch-and-bound schemes; see [21] for details. Other techniques apply specifically mixed-integer linear programs (MILPs) and, more recently, mixed-integer second-order cone programs (MISOCPs); specialized solvers capable of handling MILPs and MISOCPs include the commercial solvers Mosek, Gurobi, and CPLEX, as well as ECOS-BB, an extension to the open-source, embedded second-order cone programming solver ECOS [22].

1.3. Contributions

In this paper, we give a new formulation of (1) as a mixed-integer convex program, based on perspective functions. We can then obtain a lower bound on (1) by relaxing the integer constraints and solving the resulting convex optimization problem. We show that this lower bound is at least as good as the lower bound obtained by relaxing the integer constraints of either the MLD or GDP formulations; our numerical study suggests that this difference can be substantial. We also show how to combine our formulation with a simple shrinking-horizon heuristic to get upper bounds on (1). Again, our numerical study suggests that this upper bound can be much tighter than the upper bound obtained using the same shrinking-horizon heuristic with the MLD or GDP formulation.

Our formulation is of course related to, and derivable from, several other approaches, although not in simple or obvious ways. Our formulation is derivable from the standard MICP reformulation procedure for (convex) disjunctive programs, as given in [6,7]. However, it differs from the "convex hull" approach followed in [9], which involves minimizing the original objective function over the convex hull of the disjunctive constraints. Instead, our formulation is obtained by first considering an epigraph formulation of (1), then treating all constraints as disjunctive constraints (even if the constraint is the same for all disjunctions); only then do we apply the convex hull relaxation.

Our lower bound can also be derived from the approach of Wang, O'Donoghue, and Boyd [15] (when modified to apply to a finite-horizon problem). In particular, if we take a chain of T Bellman inequalities, and restrict our search to value function under-estimators that are affine (instead of quadratic), then the problem of maximizing the value function under-estimator (evaluated at x_{init}) is the dual of our formulation.

1.4. Outline

In Section 2, we review some properties of perspectives of convex functions. In Section 3, we give an alternate MICP formulation based using perspective functions, and we prove its equivalence to (1). In Section 4, we review three other approaches to solving (1): by the standard conversion to a mixed logical dynamical system, by generalized disjunctive programming, and by approximate dynamic programming. We then compare these methods to our perspective-base formulation. In Section 5, we give an example with numerical results, and in Section 6, we give some extensions of our method to problems similar to (1).

2. Perspective of a function

Recall that the *perspective* of an extended-value convex function $g : \mathbf{R}^n \to \mathbf{R} \cup \{\infty\}$ is the function $p : \mathbf{R}^{n+1} \to \mathbf{R} \cup \{\infty\}$ defined by:

$$p(x, s) = \begin{cases} sg(x/s) & s > 0 \\ 0 & s = 0,\ x = 0 \\ \infty & \text{otherwise.} \end{cases}$$

Crucially, if g is convex, then so is p. (This can be shown by directly checking Jensen's inequality for all cases above.) For more