Automatica 75 (2017) 249–259

Contents lists available at ScienceDirect

Automatica

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

Robust optimal control with adjustable uncertainty sets[☆]

Xiaojing Zhang [a], Maryam Kamgarpour [a], Angelos Georghiou [b], Paul Goulart [c], John Lygeros [a]

[a] Automatic Control Laboratory, ETH Zurich, Switzerland

[b] Desautels Faculty of Management, McGill University, Montreal, Quebec, Canada

[c] Department of Engineering Science, Oxford University, United Kingdom

ARTICLE INFO

Article history:
Received 15 November 2015
Received in revised form
7 June 2016
Accepted 12 August 2016

Keywords:
Robust optimal control
Adjustable uncertainty sets
Affine policies
Robust optimization

ABSTRACT

In this paper, we develop a unified framework for studying constrained robust optimal control problems with *adjustable uncertainty sets*. In contrast to standard constrained robust optimal control problems with known uncertainty sets, we treat the uncertainty sets in our problems as additional decision variables. In particular, given a finite prediction horizon and a metric for adjusting the uncertainty sets, we address the question of determining the optimal size and shape of the uncertainty sets, while simultaneously ensuring the existence of a control policy that will keep the system within its constraints for all possible disturbance realizations inside the adjusted uncertainty set. Since our problem subsumes the classical constrained robust optimal control design problem, it is computationally intractable in general. Nevertheless, we demonstrate that by restricting the families of admissible uncertainty sets and control policies, the problem can be formulated as a tractable convex optimization problem. We show that our framework captures several families of (convex) uncertainty sets of practical interest, and illustrate our approach on a demand response problem of providing control reserves for a power system.

1. Introduction

Robust finite-horizon optimal control of constrained linear systems subject to additive uncertainty has been studied extensively in the literature, both in the control (Bemporad & Morari, 1999; Löfberg, 2003; Skaf & Boyd, 2010) and operations research community (Ben-Tal, Goryashko, Guslitzer, & Nemirovski, 2004; Bertsimas & Thiele, 2006; Postek & den Hertog, 2016). Apart from issues such as stability and recursive feasibility that arise in the context of Model Predictive Control, significant amount of research is concerned with the approximation and efficient computation of the optimal control policies associated with such problems (Camacho & Bordons, 2013; Mayne, Rawlings, Rao, & Scokaert, 2000; Wang, Ong, & Sim, 2010).

Commonly, robust control problems of constrained systems over a finite horizon deal with uncertainty sets that are known a priori. In this paper, we add another layer of complexity to these problems by allowing the uncertainty sets to be decision variables of our problems, and refer to such problems as constrained robust optimal control problems with *adjustable uncertainty sets*. For example, if the uncertainty sets are interpreted as a system's resilience against disturbance, then our framework can be used in a robustness analysis setup for determining the limits of robustness of a given system. The goal then is to determine the optimal size and shape of the uncertainty sets which maximize a given metric, while ensuring the existence of a control policy that will keep the system within its constraints. Unfortunately, such problems are computationally intractable in general, since they subsume the standard robust optimal control problem with fixed uncertainty set. The aim of this paper is to propose a systematic method for finding approximate solutions in a computationally efficient way.

Our work is motivated by reserve provision problems, where the adjustable uncertainty set is interpreted as a *reserve capacity*, which a system can offer to third parties and for which it receives (financial) reward. In this case, the maximum reserve capacity can be computed by maximizing the size of the uncertainty set.

☆ This research was partially supported by Nano-Tera.ch under the project HeatReserves and by the European Commission under the project Local4Global (grant agreement number 611538). The material in this paper was partially presented at the 53rd IEEE Conference on Decision and Control, December 15–17, 2014, Los Angeles, CA, USA. This paper was recommended for publication in revised form by Associate Editor Riccardo Scattolini under the direction of Editor Ian R. Petersen.

E-mail addresses: xiaozhan@control.ee.ethz.ch (X. Zhang), mkargam@control.ee.ethz.ch (M. Kamgarpour), angelos.georghiou@mcgill.ca (A. Georghiou), paul.goulart@eng.ox.ac.uk (P. Goulart), jlygeros@control.ee.ethz.ch (J. Lygeros).

http://dx.doi.org/10.1016/j.automatica.2016.09.016

Moreover, the reserve capacity is to be chosen such that for every admissible reserve demand, i.e. for every realization within the reserve capacity set, our system is indeed able to provide this reserve without violating its constraints. Reserve provision problems of this kind were first formulated and studied in Zhang, Kamgarpour, Goulart, and Lygeros (2014) where it was shown that for uncertainty sets described by norm balls, the problems can be reformulated as tractable convex optimization problems. Another problem that admits the interpretation as a robust control problem with adjustable uncertainty set is *robust input tracking* (Gorecki, Bitlislioglu, Stathopoulos, & Jones, 2015; Vrettos, Oldewurtel, Zhu, & Andersson, 2014), where the aim is to determine the largest set of inputs that can be tracked by a system without violating its constraints. Reserve provision and input tracking problems have recently received increased attention in demand response applications of control reserves for electrical power grids (Vrettos, Oldewurtel, & Andersson, in press; Wu, Chen, Zhang, & Su, 2014; Zhang, Vrettos, Kamgarpour, Andersson, & Lygeros, 2015).

The purpose of this paper is two-fold: First, we generalize the work of Gorecki et al. (2015) and Zhang et al. (2014) by considering a larger class of adjustable uncertainty sets based on techniques of conic convex optimization. Second, we provide a unified framework for studying reserve provision, input tracking and robustness analysis problems under the umbrella of constrained robust optimal control with adjustable uncertainty sets. The main contributions of this paper with respect to the existing literature can be summarized as follows:

- We show that if (i) the uncertainty sets are restricted to those that can be expressed as affine transformations of properly selected *primitive* convex sets, and (ii) the control policies are restricted to be affine with respect to the elements in these primitive sets, then the problems admit convex reformulations that can be solved efficiently, and whose size grows polynomially in the decision parameters. In particular, we extend the results of Zhang et al. (2014) and Gorecki et al. (2015) in two ways: First, we show that any convex set can be used as a primitive set, allowing us to target a much larger class of uncertainty sets. Second, by allowing the primitive sets to be defined on higher dimensional spaces than those of the uncertainty sets, we are able to design more flexible uncertainty sets.
- We identify families of uncertainty sets of practical interest, including norm-balls, ellipsoids and hyper-rectangles, and show that they can be adjusted efficiently. Extending the work of Gorecki et al. (2015) and Zhang et al. (2014), we also show that by choosing the primitive set as the simplex, our framework enables us to efficiently optimize over compact polytopes with a predefined number of vertices. Furthermore, we prove that if the primitive sets are polytopes (e.g. the simplex), then our policy approximation gives rise to continuous piece-wise affine controllers.
- We study a reserve provision problem that arises in power systems, and show how it can be formulated as a robust optimal control problem with an adjustable uncertainty set. The problem is addressed using the developed tools, and we show that it can be formulated as a linear optimization problem of modest size that can be solved efficiently within 0.3 s, making it also practically applicable.

This paper is organized as follows: Section 2 introduces the general problem setup. Section 3 focuses on the problem of adjusting the uncertainty sets, while in Section 4, we return to the original problem and restrict the family of control policies to obtain tractable instances thereof. Section 5 illustrates our approach on a demand response problem, while Section 6 demonstrates the usefulness of allowing the uncertainty sets to be projections of high-dimensional primitive convex sets. Finally, Section 7 concludes the paper. The Appendix contains auxiliary results needed to prove the main results of the paper.

Notation

For given matrices $(A_1, \ldots, A_n)$, we define $A := \mathrm{diag}(A_1, \ldots, A_n)$ as the block-diagonal matrix with elements $(A_1, \ldots, A_n)$ on its diagonal. A_{ij} denotes the (i, j)th element of the matrix A, while $A_{\cdot j}$ denotes the jth column of A. Given a cone $\mathcal{K} \subset \mathbb{R}^l$ and two vectors $a, b \in \mathbb{R}^l$, $a \preceq_{\mathcal{K}} b$ implies $(b-a) \in \mathcal{K}$. For a matrix $B \in \mathbb{R}^{m\times l}$, $B \succeq_{\mathcal{K}} 0$ denotes row-wise inclusion in $\mathcal{K}$. For a symmetric matrix $C \in \mathbb{R}^{n\times n}$, $C \succeq 0$ denotes positive semi-definiteness of C. Given vectors $(v_1, \ldots, v_m)$, $v_i \in \mathbb{R}^l$, we denote their convex hull as $\mathrm{conv}(v_1, \ldots, v_m)$. Moreover, $[v_1, \ldots, v_m] := [v_1^\top \ \ldots \ v_m^\top]^\top \in \mathbb{R}^{lm}$ denotes their vector concatenation.

2. Problem formulation

In this section, we formulate the robust optimal control problem with adjustable uncertainty set. We consider uncertain linear systems of the form

$$x_{k+1} = Ax_k + Bu_k + Ew_k, \tag{1}$$

where $x_k \in \mathbb{R}^{n_x}$ is the state at time step k given an initial state $x_0 \in \mathbb{R}^{n_x}$, $u_k \in \mathbb{R}^{n_u}$ is the control input and $w_k \in \mathbb{W}_k \subset \mathbb{R}^{n_w}$ is an uncertain disturbance. We consider compact polytopic state and input constraints

$$\begin{aligned} &x_k \in \mathbb{X} := \{x \in \mathbb{R}^{n_x} : F_x x \le f_x\}, \quad k = 1, \ldots, N, \\ &u_k \in \mathbb{U} := \{u \in \mathbb{R}^{n_u} : F_u u \le f_u\}, \quad k = 0, \ldots, N-1, \end{aligned} \tag{2}$$

where $F_x \in \mathbb{R}^{n_f\times n_x}$, $f_x \in \mathbb{R}^{n_f}$, $F_u \in \mathbb{R}^{n_g\times n_u}$, $f_u \in \mathbb{R}^{n_g}$, and n_f (n_g) is the number of state (input) constraints. Given a planning horizon N, we denote by $\phi_k(\mathbf{u}, \mathbf{w})$ the predicted state after k time steps resulting from the input sequence $\mathbf{u} := [u_0, \ldots, u_{N-1}] \in \mathbb{R}^{Nn_u}$ and disturbance sequence $\mathbf{w} := [w_0, \ldots, w_{N-1}] \in \mathbb{R}^{Nn_w}$.

In contrast to classical robust control problem formulations, we assume that the uncertainty set $\mathbb{W}_k$ is not fixed and needs to be adjusted according to some objective function $\varrho : \mathcal{P}(\mathbb{R}^{n_w}) \to \mathbb{R}$, where $\mathcal{P}(\mathbb{R}^{n_w})$ denotes the power set of $\mathbb{R}^{n_w}$. For example, we may think of $\varrho(\mathbb{W}_k)$ as the volume of $\mathbb{W}_k$, although depending on the application, it can represent other qualities such as the diameter or circumference of $\mathbb{W}_k$. Our objective is to maximize $\varrho(\mathbb{W}_k)$, while simultaneously minimizing some operating cost and ensuring satisfaction of input and state constraints. Hence, the cost to be minimized is given by

$$\max_{\mathbf{w}\in\mathcal{W}} \{J(\mathbf{u}, \mathbf{w})\} - \lambda \sum_{k=0}^{N-1} \varrho(\mathbb{W}_k), \tag{3}$$

where $J(\mathbf{u}, \mathbf{w}) := \ell_f(\phi_N(\mathbf{u}, \mathbf{w})) + \sum_{k=0}^{N-1} \ell(\phi_{k+1}(\mathbf{u}, \mathbf{w}), u_k)$ is some "nominal" cost function with $\ell : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}$ and $\ell_f : \mathbb{R}^{n_x} \to \mathbb{R}$ linear, and $\lambda \ge 0$ is a user-defined weighting factor. Note that convex quadratic cost can also be incorporated in our framework by taking the certainty-equivalent cost $J(\mathbf{u}, \bar{\mathbf{w}})$ instead of the min-max cost in (3), where $\bar{\mathbf{w}}$ is some fixed (or expected) uncertainty. Due to the presence of the uncertainties $\mathbf{w}$, we consider the design of a causal disturbance feedback policy $\boldsymbol{\pi}(\cdot) := [\pi_0(\cdot), \ldots, \pi_{N-1}(\cdot)]$, with each $\pi_k : \mathbb{W}_0 \times \cdots \times \mathbb{W}_k \to \mathbb{R}^{n_u}$, such that the control input at each time step is given by $u_k = \pi_k(w_0, \ldots, w_k)$.[1] Combining (1)–(3), we express the optimal control problem compactly as

$$\begin{aligned} \min \quad & \max_{\mathbf{w}\in\mathcal{W}} \left\{\mathbf{c}^\top \boldsymbol{\pi}(\mathbf{w})\right\} - \lambda\boldsymbol{\varrho}(\mathcal{W}) \\ \text{s.t.} \quad & \boldsymbol{\pi}(\cdot) \in \mathcal{C}, \ \mathcal{W} \in \mathcal{P}(\mathbb{R}^{Nn_w}), \\ & \mathbf{C}\boldsymbol{\pi}(\mathbf{w}) + \mathbf{D}\mathbf{w} \le \mathbf{d}, \quad \forall \mathbf{w} \in \mathcal{W}, \end{aligned} \tag{4}$$

[1] Strictly causal policies can be incorporated by restricting $\pi_k(\cdot)$ to depend on $(w_0, \ldots, w_{k-1})$ only. For simplicity, this paper considers causal policies. However, all subsequent results apply to strictly causal policies with minor modifications.