Automatica 77 (2017) 143–154

Contents lists available at ScienceDirect

Automatica

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

Simultaneous nonlinear model predictive control and state estimation ☆

David A. Copp, João P. Hespanha

University of California, Santa Barbara, CA 93106-9560, USA

ARTICLE INFO

Article history:
Received 30 March 2016
Received in revised form
30 August 2016
Accepted 7 November 2016

Keywords:
Model predictive control
Output feedback control
Control of constrained systems
Optimal control
Optimal estimation
Algorithms and software

ABSTRACT

An output-feedback approach to model predictive control that combines state estimation and control into a single min–max optimization is introduced for discrete-time nonlinear systems. Specifically, a criterion that involves finite forward and backward horizons is minimized with respect to control input variables and is maximized with respect to the unknown initial state as well as disturbance and measurement noise variables. Under appropriate assumptions that encode controllability and observability, we show that the state of the closed-loop remains bounded and that a bound on tracking error can be found for trajectory-tracking problems. We also introduce a primal–dual interior-point method that can be used to efficiently solve the min–max optimization problem and show in simulation examples that the method succeeds even for severely nonlinear and non-convex problems.

1. Introduction

Online optimization has become a ubiquitous approach for solving control and estimation problems in both academia and industry. This is largely due to the ability to explicitly accommodate hard state and input constraints in online optimization techniques. Because of this, an especially popular online optimization control technique called model predictive control (MPC) is used in numerous industrial applications (Qin & Badgwell, 2003), and, consequently, much effort has been devoted to developing a stability theory for MPC (see e.g. Camacho & Bordons, 2004; Grüne & Pannek, 2011; Morari & H Lee, 1999; Rawlings, 2000; Rawlings & Mayne, 2009). An overview of recent developments can be found in Mayne (2014).

MPC involves the solution of an open-loop optimal control problem at each sampling time. Each of these optimizations results in a sequence of future optimal control actions and a sequence of corresponding future states. The first control action in the sequence is applied to the plant, and then the optimization is solved again at the next sampling time. MPC has historically been popular for problems in which the plant dynamics are sufficiently slow so that the optimization can be solved between consecutive sampling times. However, as available computational power increases and optimization algorithms improve in terms of speed, MPC can be applied to broader application areas.

MPC is often formulated assuming that the full state of the process to be controlled can be measured. However, this is not possible in many practical cases, so the use of independent algorithms for state-estimation, including observers, filters, and moving horizon estimation (MHE), as discussed, i.e., in Rawlings and Bakshi (2006), is required. Of these methods, MHE is especially attractive for use with MPC because it can be formulated as a similar online optimization problem that explicitly handles constraints. Solving the MHE problem produces a state estimate that is compatible with a set of past measurements that recedes as the current time advances. This estimate is optimal in the sense that it maximizes a criterion that captures the likelihood of the measurements. By receding the set of measurements considered in the MHE optimization, one maintains a constant computational cost for the optimization.

In this paper, we propose an approach to combine MPC and MHE into a single optimization that is solved online to construct an output-feedback controller. To account for the uncertainty that results from unmeasured disturbances and measurement noise, we replace the minimization that is used in classical MPC by a min–max optimization. In this case, the minimization is carried out with respect to future control actions, and the maximization

☆ 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: dacopp@engr.ucsb.edu (D.A. Copp), hespanha@ece.ucsb.edu (J.P. Hespanha).

http://dx.doi.org/10.1016/j.automatica.2016.11.041

is taken with respect to the variables that cannot be measured, namely the system's initial state, the unmeasured disturbances, and the output measurement noise. The criterion for this min–max optimization combines a term that captures the control objective and a term that captures the likelihood of the uncertain variables, resulting in essentially the summation of an MPC criterion with an MHE criterion.

The main technical contribution of this paper addresses the stability of the proposed combined MPC/MHE approach. We show that the proposed output-feedback controller results in closed-loop trajectories along which the state of the process remains bounded, and, for tracking problems, our results provide explicit bounds on the tracking error. These results rely on three key assumptions: The first assumption requires the existence of saddle-point equilibria for the min–max optimization, or equivalently, that the min and max commute. For linear systems and quadratic costs, this assumption is satisfied if the system is observable and weights in the cost function are chosen appropriately (Copp & Hespanha, 2016b). The second assumption requires the optimization criterion to include a terminal cost that is a control ISS-Lyapunov function with respect to the disturbance input. This type of assumption is common in classical state-feedback robust MPC. The final observability assumption essentially requires that the backwards horizon is sufficiently large so that enough information about the initial state is obtained in order to find past estimates that are compatible with the dynamics.

A second contribution of this paper is a new primal–dual interior-point algorithm that can be used to compute the saddle-point equilibrium that needs to be solved for online at each sampling time. This algorithm relies on the use of Newton's method to solve a relaxed version of the Karush–Kuhn–Tucker (KKT) conditions associated with the coupled optimizations that define the saddle-point equilibrium. As in classical primal–dual methods, we replace the equality to zero of the complementary slackness conditions by equality to a positive constant μ that we force to converge to zero as the Newton iterations progress. In practice, the algorithm will stop with a positive value for μ, but we show that this still leads to an ϵ-saddle-point, where ϵ can be explicitly computed and made arbitrarily small through the selection of an appropriate stopping criterion.

1.1. Related work

State-feedback MPC is a mature field with numerous contributions. Particularly relevant to the results in this paper is the work on the so-called robust or min–max MPC, which considers model uncertainty, input disturbances, and noise (Bemporad & Morari, 1999; Campo & Morari, 1987; Lee & Yu, 1997; Magni, De Nicolao, Scattolini, & Allgöwer, 2003). Min–max MPC for constrained linear systems was considered by Bemporad, Borrelli, and Morari (2003) and Scokaert and Mayne (1998), and a game theoretic approach for robust constrained nonlinear MPC was proposed by Chen, Scherer, and Allgöwer (1997). More recent studies of input-to-state stability of min–max MPC can be found in Lazar, Muñoz de la Peña, Heemels, and Alamo (2008), Limon, Alamo, Raimondo, Muñoz de la Peña, Bravo, Ferramosca, and Camacho (2009) and Raimondo, Limon, Lazar, Magni, and Camacho (2009). These works focused on state-feedback MPC and did not consider robustness with respect to errors in state estimation. A novelty of the work presented in this paper is the reliance on saddle-point equilibria, rather than a simple min–max optimal, which we found instrumental in establishing our stability results.

Fewer results are available for output-feedback MPC. An overview of nonlinear output-feedback MPC is given by Findeisen, Imsland, Allgöwer, and Foss (2003) and the references therein. Many of these output-feedback approaches involve designing separate state estimator and MPC schemes. Several of the observers, estimators, and filters that have been proposed for use with nonlinear output-feedback MPC include an extended Kalman filter (Huang, Patwardhan, & Biegler, 2009), optimization based moving horizon observers (Michalska & Mayne, 1995), high gain observers (Imsland, Findeisen, Bullinger, Allgöwer, & Foss, 2003), extended observers (Roset, Lazar, Nijmeijer, & Heemels, 2006), and robust MHE (Zhang & Liu, 2013). In contrast to solving the estimation and control problems separately, the formulation of our combined MPC/MHE approach as a single optimization facilitates the stability analysis of the closed-loop without the need for a separation principle for nonlinear systems.

Results on robust output-feedback MPC for constrained, linear, discrete-time systems with bounded disturbances and measurement noise can be found in Mayne, Raković, Findeisen, and Allgöwer (2006, 2009), where a stable Luenberger observer is employed for state estimation and robustly stabilizing tube-based MPC is performed to control the state of the observer. Alternatively, in Sui, Feng, and Hovd (2008), MHE is employed for state estimation and is combined with a similar tube-based MPC approach. These approaches first solve the estimation problem and show convergence of the state estimate to a bounded set and then take the uncertainty of the state estimate into account when solving the robust MPC problem. The work of Löfberg (2002) combines an estimation scheme, which provides a guaranteed ellipsoidal error bound on the state estimate, with a min–max MPC scheme for estimation and control of linear systems with bounded disturbances and measurement noise.

During the same time that many important results on MPC were developed, parallel work began on MHE. The work of Allgöwer, Badgwell, Qin, Rawlings, and Wright (1999) gives a tutorial overview and background of both MPC and MHE as well as methods that can be used to solve these optimization problems. Useful overviews of constrained linear and nonlinear MHE can be found in Rao, Rawlings, and Lee (2001); Rao, Rawlings, and Mayne (2003) where, with appropriate assumptions regarding observability, continuity, and an approximate arrival cost, the authors prove asymptotic stability as well as bounded stability in the presence of bounded noise.

More recent results regarding MHE for discrete-time nonlinear systems are given by Alessandri, Baglietto, and Battistelli (2008), in which the authors minimize a quadratic cost that includes the standard output error term as well as a term penalizing the distance of the current estimated state from its prediction. The authors prove boundedness of the estimation error, when considering bounded disturbances and measurement noise, and convergence of the state estimate to the true value in the noiseless case. Even more recent work on robust MHE for nonlinear systems appeared in Liu (2013), where first a high-gain observer is used to bound the estimation error, and then that bound is used to design a constraint for incorporation in an MHE problem. This formulation seems to reduce the sensitivity of the performance of MHE to the accuracy of the approximate arrival cost, and boundedness of the state estimate is proven when the noise is bounded.

The optimization algorithm proposed here is heavily inspired by primal–dual interior-point methods (Wright, 1997b) that have been very successful in solving convex optimizations (Boyd & Vandenberghe, 2004). The use of interior-point algorithms to solve MPC problems is discussed by Rao, Wright, and Rawlings (1998), and additional early work on efficient numerical methods for solving MPC problems can be found in Biegler (2000), Biegler and Rawlings (1991) and Wright (1997a). An overview of the numerical methods available for solving the optimization problems that arise in nonlinear MPC and MHE is given by Diehl, Ferreau, and Haverbeke (2009), whereas the more recent work (Wang & Boyd,