Computers and Chemical Engineering 114 (2018) 158–170

Contents lists available at ScienceDirect

Computers and Chemical Engineering

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

Stochastic model predictive control — how does it work?

Tor Aksel N. Heirung, Joel A. Paulson, Jared O'Leary, Ali Mesbah*

Department of Chemical and Biomolecular Engineering, University of California, Berkeley, CA 94720, USA

ARTICLE INFO

Article history:
Received 16 June 2017
Received in revised form 24 October 2017
Accepted 25 October 2017
Available online 5 November 2017

Keywords:
Model predictive control
Stochastic optimal control
Chance constraints
Disturbance modeling

ABSTRACT

Stochastic model predictive control (SMPC) provides a probabilistic framework for MPC of systems with stochastic uncertainty. A key feature of SMPC is the inclusion of chance constraints, which enables a systematic trade-off between attainable control performance and probability of state constraint violations in a stochastic setting. This paper presents an overview of core concepts in SMPC in relation to MPC and stochastic optimal control, with numerical illustrations on a typical chemical process. Estimation of stochastic disturbances as well as the impact of estimation quality of stochastic disturbances on the SMPC performance are discussed. Some avenues for future research in SMPC are suggested.

1. Introduction

This article presents a tutorial overview of stochastic model predictive control (SMPC). After introducing the concept of stochastic optimal control, the connections between SMPC and both stochastic optimal control and MPC are explained in order to illustrate how receding-horizon control in a stochastic setting requires only a minor modification to the standard formulation of MPC. In particular, the article argues that the basic ideas of MPC and classical linear quadratic control of Gaussian systems provide the necessary foundation for SMPC.

This section first briefly reviews MPC, then discusses a standard formulation of stochasticity in linear systems, and finally outlines the rest of the paper and a case study used throughout.

Notation

The identity matrix of dimension $\mathbb{R}^{n\times n}$ is denoted by I_n. $M > 0$ ($M \geq 0$) denotes a positive definite (semi-definite) matrix. v_k denotes the value of a variable v at time k, with $v_{k+i|k}$ denoting the value of v at future time $k+i$ predicted from time k. p_v denotes the probability distribution of the variable v (i.e., $v \sim p_v$). $\Pr[A]$ and $\mathrm{E}[v]$ denote the probability of the event A and the expected value of the random variable v, respectively. $\Pr_k[A]$ and $\mathrm{E}_k[v]$ denote, respectively, the probability and expected value conditioned on the system information available at time k.

1.1. Model predictive control

MPC, also known as receding-horizon control, is widely used for advanced control of multivariable systems with constraints on states and control inputs (Mayne et al., 2000; Morari and Lee, 1999). Well-established applications of MPC include chemical process control (Qin and Badgwell, 2003; Forbes et al., 2015), building climate control (Oldewurtel et al., 2012; Zhang et al., 2014; Ma et al., 2015), networked controlled systems (Camponogara et al., 2002; Scattolini, 2009), and vehicle path following (Falcone et al., 2007; Faulwasser et al., 2009). We briefly review the mathematical foundations of MPC.

Consider a linear time-invariant (LTI) system of the discrete-time form

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

$$y_k = Cx_k, \tag{1b}$$

where $x_k \in \mathbb{R}^{n_x}$, $u_k \in \mathbb{R}^{n_u}$, and $y_k \in \mathbb{R}^{n_y}$ denote the state, control input, and measured output at sampling instant k, respectively; $A \in \mathbb{R}^{n_x\times n_x}$, $B \in \mathbb{R}^{n_x\times n_u}$, and $C \in \mathbb{R}^{n_y\times n_x}$ are the system matrices.

When perfect knowledge of the state x_k is available, MPC involves solving the following optimal control problem (OCP) at every sampling time k

* Corresponding author.
E-mail address: mesbah@berkeley.edu (A. Mesbah).

https://doi.org/10.1016/j.compchemeng.2017.10.026
0098-1354/

$$\min_{\mathbf{u}} \; J_N(x_k, \mathbf{u}) \tag{2a}$$

$$\text{subject to} \quad x_{k+i+1|k} = Ax_{k+i|k} + Bu_{k+i|k} \tag{2b}$$

$$Hx_{k+i+1|k} \leq h \tag{2c}$$

$$Du_{k+i|k} \leq d \tag{2d}$$

$$i = 0, 1, \ldots, N-1 \tag{2e}$$

$$x_{k|k} = x_k, \tag{2f}$$

where N is the prediction horizon and the subscript $k+i|k$ denotes the value of a variable at the future time $k+i$, predicted based on the knowledge of system at time k. In (2c)–(2d), $H \in \mathbb{R}^{n_s \times n_x}$ and $D \in \mathbb{R}^{n_i \times n_u}$ are state and input constraint matrices, respectively, with $h \in \mathbb{R}^{n_s}$ and $d \in \mathbb{R}^{n_i}$ denoting the corresponding constraint values for the n_s state constraints and n_i input constraints. Here, the cost function, commonly chosen as a regularization cost for driving the state and input to zero, is

$$J_N(x_k, \mathbf{u}) = \sum_{i=0}^{N-1} (x_{k+i|k}^\top Q_x x_{k+i|k} + u_{k+i|k}^\top R_u u_{k+i|k}) + x_{k+N|k}^\top Q_N x_{k+N|k} \tag{3}$$

where $\mathbf{u} := \{u_{k|k}, u_{k+1|k}, \ldots, u_{k+N-1|k}\}$ is a sequence of control inputs, and the matrices $Q_x \geq 0$, $R_u > 0$, and $Q_N \geq 0$ are weight matrices.

Note that the cost function (3) can be modified to include a term that penalizes the input rate of change $\Delta u_{k+i|k} = u_{k+i|k} - u_{k+i-1|k}$. Because the constraints (2b)–(2d) are linear, the cost $J_N(x_k, \mathbf{u})$ is quadratic, and the weight matrices are positive (semi) definite, the OCP (2) is a convex quadratic programming (QP) problem. An optimization problem of this type always has a unique minimum, and can be solved efficiently using standard techniques (Nocedal and Wright, 2006).

The minimizer for (2) at sampling time k is the open-loop optimal control sequence $\mathbf{u}^*(x_k)$. Since the OCP (2) is solved in a receding-horizon manner, only the first element of the optimal control input sequence, $u_{k|k}^*(x_k)$, is applied to the system. Thus, at every k the MPC computes an implicit feedback control law $u_k = \kappa_N(x_k)$, where

$$\kappa_N(x_k) = u_{k|k}^*(x_k) \tag{4}$$

is determined by solving (2).

In the absence of state and input constraints, the solution to the OCP (2) takes the form of a linear feedback controller $u_{k+i|k} = -K_i x_{k+i|k}$, known as the linear quadratic regulator (LQR). The time-varying feedback gain K_i is given by

$$K_i = (R_u + B^\top P_{i+1} B)^{-1} B^\top P_{i+1} A. \tag{5}$$

LQR is the optimal control for the linear system (1) under mild assumptions, and results in the stable closed-loop dynamics $x_{k+1} = (A - BK_k)x_k$.

For any time i, the matrix $P_i = P_i^\top \geq 0$ in (5) is computed through iteration of the discrete-time Riccati equation

$$P_{i-1} = Q_x + A^\top P_i A - A^\top P_i B (R_u + B^\top P_i B)^{-1} B^\top P_i A \tag{6}$$

backward in time, starting from $P_N = Q_N$. The matrix P_i generally converges rapidly to its steady-state value P, leading to a constant feedback gain K after only a few iterations. At steady state, P does not depend on the time step i. As a result, the expression (6) can be rewritten as the discrete-time algebraic Riccati equation (DARE)

$$P = (A - BK)^\top P (A - BK) + Q_x + K^\top R_u K,$$

which can be derived from (6) using the Sherman–Morrison–Woodbury formula for matrix inversion.

When $u_k = -Kx_k$ is implemented as the control law for the system (2), the Riccati matrix P has the property that

$$\sum_{k=0}^{\infty} (x_k^\top Q_x x_k + u_k^\top R_u u_k) = \sum_{k=0}^{\infty} x_k^\top (Q_x + K^\top R_u K) x_k = x_0^\top P x_0 .$$

This implies that when setting $Q_N = P_N = P$ in (3), the *terminal cost* $x_{k+N|k}^\top P x_{k+N|k}$ captures the cost from $k = N$ to $k = \infty$ under the assumption that $u_{k+i|k} = -Kx_{k+i|k}$ for $i \geq N$. This formulation is known as the dual-mode paradigm (Mayne et al., 2000), which refers to the mode $0 \leq i < N$ in which the inputs are free decision variables over the finite horizon N and the mode $i \geq N$ where the state feedback control law is used over the subsequent infinite horizon. The idea of dual-mode MPC is that as long as N is sufficiently large, $x_{k+i|k}$ for all $i \geq N$ will be sufficiently close to the origin for the constraints not to be active. When the constraints are inactive, the LQR law $u_{k+i|k} = -Kx_{k+i|k}$ is the optimal control. For $0 \leq i \leq N-1$, the inputs $u_{k+i|k}$ are decision variables that minimize the cost (3) such that the predicted states and inputs are feasible with respect to the constraints (2c)–(2d). Hence, setting $Q_N = P$ implies that $u_{k+i|k} = -Kx_{k+i|k}$ for $i \geq N$, and enables optimization over an infinite horizon.

1.2. System uncertainty

In system (1), it is assumed that the state x_k evolves in a deterministic manner and that there are no errors in the state measurements y_k. In practice, however, our knowledge of the system dynamics is uncertain. The system uncertainty generally manifests itself in terms of uncertain model structure and/or parameters, uncertain initial conditions, unmeasured disturbances, and measurement error. A common way to incorporate uncertainty into the system description (1) is to modify the model to

$$x_{k+1} = Ax_k + Bu_k + Gw_k \tag{7a}$$

$$y_k = Cx_k + v_k, \tag{7b}$$

where $w_k \in \mathbb{R}^{n_w}$ and $v_k \in \mathbb{R}^{n_v}$ denote system disturbances and the measurement noise, respectively, and $G \in \mathbb{R}^{n_x \times n_w}$ models the effects of w_k on the system state. Here the disturbances w_k can be seen as variables that capture the combined effect of model uncertainty and exogenous disturbances on the evolution of the state. When the disturbances and measurement noise are described as random variables, (7) becomes a stochastic model of the system dynamics. That is, even though the true system (1) evolves deterministically, our understanding of the system dynamics is described in a probabilistic manner. Note that this concept is in contrast to having a system that is intrinsically stochastic in that the system naturally exhibits random behavior, for example, due to Brownian motion. Although the stochastic optimal control methods discussed in this paper can be adopted for naturally stochastic systems, we restrict our discussion to the problem of probabilistic system uncertainty based on the stochastic system model (7). In the remainder of the paper, both w_k and v_k are assumed to be sequences of independent and identically distributed (i.i.d.) variables with known probability distributions p_w and p_v, respectively, and $\mathrm{E}[w_k w_k^\top] = Q_w$, $\mathrm{E}[v_k v_k^\top] = Q_v$, and $\mathrm{E}[w_k v_k^\top] = 0$.

When a stochastic model (7) is used to describe the dynamics of the uncertain system, the OCP (2) must be modified to account for the probabilistic nature of the model predictions. The cost $J_N(x_k, \mathbf{u})$ will become a random quantity, which must be replaced with some statistic such as its expected value $\mathrm{E}_k[J_N(x_k, \mathbf{u})]$. The conditional expectation $\mathrm{E}_k[J_N(x_k, \mathbf{u})]$ relies on the measurement information used for deducing the system state (i.e., through state estimation in the absence of complete state measurements), the knowledge of which is required for initializing the OCP (2) at every sampling time k. Similarly, the state constraints must be replaced to reflect that the