UC Berkeley

UC Berkeley Previously Published Works

Title

Stochastic model predictive control with joint chance constraints

Permalink

<https://escholarship.org/uc/item/13r9m609>

Journal

International Journal of Control, 93(1)

ISSN

0020-7179

Authors

Paulson, JA
Buehler, EA
Braatz, RD
[et al.](#)

Publication Date

2020-01-02

DOI

10.1080/00207179.2017.1323351

Peer reviewed

Stochastic model predictive control with joint chance constraints

Joel A. Paulson^a, Edward A. Buehler^a, Richard D. Braatz^b and Ali Mesbah^a

^aDepartment of Chemical and Biomolecular Engineering, University of California, Berkeley, CA, USA; ^bDepartment of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA

ABSTRACT

This article investigates model predictive control (MPC) of linear systems subject to arbitrary (possibly unbounded) stochastic disturbances. An MPC approach is presented to account for hard input constraints and joint state chance constraints in the presence of unbounded additive disturbances. The Cantelli–Chebyshev inequality is used in combination with risk allocation to obtain computationally tractable but accurate surrogates for the joint state chance constraints when only the mean and variance of the arbitrary disturbance distributions are known. An algorithm is presented for determining the optimal feedback gain and optimal risk allocation by iteratively solving a series of convex programs. The proposed stochastic MPC approach is demonstrated on a continuous acetone–butanol–ethanol fermentation process, which is used in the production of biofuels.

ARTICLE HISTORY

Received 25 March 2016
Accepted 23 April 2017

KEYWORDS

Stochastic systems;
predictive control;
disturbance feedback
parametrization; risk
allocation

1. Introduction

Recent years have witnessed significant developments in the area of robust model predictive control (MPC) with the aim to devise optimal control approaches that enable systematic handling of system uncertainties (Mayne, 2014). In general, robust MPC approaches consider bounded, deterministic descriptions of uncertainties. The deterministic approaches to robust MPC commonly use a min–max optimal control formulation in which the control policy is designed with respect to the worst-case performance and system constraints are satisfied for all possible uncertainty realisations (Bemporad & Morari, 1999). These approaches can lead to overly conservative or possibly infeasible control designs.

In practice, system uncertainties are often considered to be of stochastic nature. When a stochastic description of uncertainties is available, a natural approach to robust MPC involves explicitly accounting for the likelihood of occurrence of these uncertainties in designing the robust control policy. This consideration has led to the emergence of stochastic MPC (SMPC). A core component of SMPC is (state) chance constraints that allow for constraint satisfaction in a probabilistic sense. Chance constraints enable SMPC to trade off robustness to uncertainties (in terms of constraint satisfaction) with control performance in a systematic manner, possibly resulting in less conservative control performance.

A recent review on different SMPC approaches and their applications is given in Mesbah (2016).

Stochastic tube approaches to SMPC (e.g. Kouvaritakis & Cannon, 2015) use probabilistic tubes with fixed or variable cross sections to replace chance constraints with linear constraints on the nominal state predictions as well as to construct terminal sets for guaranteeing recursive feasibility. These approaches use a prestabilising feedback controller to ensure closed-loop stability. However, stochastic tube approaches cannot handle hard input constraints as the prestabilising state feedback controller is determined offline. SMPC approaches based on an *affine parametrisation* of the feedback control law have been extensively investigated (e.g. Hokayem, Cinquemani, Chatterjee, Ramponi, & Lygeros, 2012; Oldewurtel, Jones, & Morari, 2008; Paulson, Streif, & Mesbah, 2015). Such control law parametrisations allow for obtaining convex SMPC algorithms while solving the stochastic optimal control problem over the feedback gains as well as the open-loop control actions. The notion of the affine disturbance parametrisation of the feedback control laws originates from the fact that the disturbance realisations can be reconstructed from state measurements that will be available to the controller in the future (Goulart, Kerrigan, & Maciejowski, 2006). Therefore, the controller can use this information when determining the future control inputs over the horizon. A key challenge in using such parametrisations, however, arises from handling hard input constraints in the presence of unbounded stochastic uncertainties (e.g. Gaussian noise), as unbounded uncertainties almost surely lead to excursions of states from any bounded set. To address this

challenge, the inclusion of *saturation functions* into the affine feedback control policy has been proposed (Hokayem et al., 2012). Saturation functions render the feedback control policy nonlinear to enable direct handling of hard input constraints without relaxing the hard input constraints to probabilistic input constraints.

Extensive work has also been reported on SMPC methods that use the so-called *sample-based approaches* (also known as scenario-based approaches). These algorithms represent the stochastic system using a finite set of realisations of the uncertainties sampled from the continuous distribution, which are used to solve the optimal control problem in one shot (e.g. Blackmore, Ono, Bektassov, & Williams, 2010; Calafiore and Fagiano, 2013). This class of SMPC approaches typically does not rely on any convexity requirements; however, establishing the recursive feasibility and closed-loop stability of these algorithms is generally challenging, particularly for the case of unbounded uncertainties.

This paper considers the MPC problem for stochastic linear systems with arbitrary (possibly unbounded) disturbances. An SMPC approach is presented that handles both joint state chance constraints and hard input constraints with feedback accounted for in the predictions. The key contribution of this work lies in using risk allocation (Ma, Vichik, & Borrelli, 2012; Vitus & Tomlin, 2011) in combination with the Cantelli–Chebyshev inequality (Marshall & Olkin, 1979) to obtain computationally tractable surrogates for the joint state chance constraints when only the first two moments of an arbitrary disturbance distribution are known (i.e. the full distribution is unknown). An algorithm is presented for solving the SMPC problem to determine the optimal feedback gain and optimal risk allocation iteratively. The problem setup is similar to that in Hokayem et al. (2012) and Farina, Giulioni, Magni, and Scattolini (2015). In contrast to Hokayem et al. (2012), the proposed SMPC approach accounts for both hard input constraints and state chance constraints. What distinguishes this work from Farina et al. (2015) is the direct handling of hard input constraints without relaxation as well as the convexity of the optimisation program. The proposed SMPC approach is demonstrated on a continuous acetone–butanol–ethanol (ABE) fermentation process (Haus et al., 2011), which is used in the production of high value-added drop-in biofuels from lignocellulosic biomass. The performance of the proposed approach is evaluated with respect to that of a certainty equivalence MPC algorithm and an MPC algorithm with fixed uniform risk allocation.

Notation: Hereafter, $\mathbb{R}$ and $\mathbb{N} = \{1, 2, \dots\}$ are the sets of real and natural numbers, respectively; $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. $\mathbb{S}_+^n$ and $\mathbb{S}_{++}^n$ are the sets of positive semidefinite and

definite matrices, respectively. I_N denotes the N by N identity matrix and $\mathbf{1}_N$ denotes a column vector of ones of length N . $\text{tr}(\cdot)$ denotes the trace of a square matrix. $\|\cdot\|_p$ denotes the standard p -norms. $\otimes$ denotes the Kronecker product. For given random vectors X and Y , $\mathbb{E}[X]$ denotes the expected value, $\sigma[X, Y] = \mathbb{E}[(X - \mathbb{E}[X])(Y - \mathbb{E}[Y])^\top]$ denotes the cross covariance matrix, and $\Sigma[X] = \sigma[X, X]$ denotes the covariance matrix. $\mathbb{P}(A)$ denotes the probability of event A .

2. Problem statement

Consider a discrete-time stochastic linear system

$$x^+ = Ax + Bu + Gw, \quad (1)$$

where $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$, and $w \in \mathbb{R}^p$ are the system states, inputs, and disturbances at the current time, respectively; x^+ denotes the system states at the next time; and A , B , and G are the known system matrices. It is assumed that the states x are observed exactly at all times, and the disturbances are mutually independent such that system (1) is a Markov process. The disturbances w can have an arbitrary (unbounded) distribution that is unknown to the controller; the mean $\mathbb{E}[w]$ and covariance $\Sigma[w] \in \mathbb{S}_+^p$ are, however, assumed to be known.

Let $N \in \mathbb{N}$ denote the prediction horizon of the predictive control problem. The states, inputs, and disturbances over the prediction horizon are defined, respectively, by

$$\begin{aligned} \mathbf{x} &= [x_0^\top, x_1^\top, \dots, x_N^\top]^\top \\ \mathbf{u} &= [u_0^\top, u_1^\top, \dots, u_{N-1}^\top]^\top \\ \mathbf{w} &= [w_0^\top, w_1^\top, \dots, w_{N-1}^\top]^\top, \end{aligned}$$

where $x_k = Ax_{k-1} + Bu_{k-1} + Gw_{k-1}$ is the predicted states k steps ahead from the known current states $x_0 = x$; and u_k and w_k are the inputs and disturbances k time steps into the future, respectively. Using this compact notation, the system model is written as

$$\mathbf{x} = \mathbf{A}x_0 + \mathbf{B}\mathbf{u} + \mathbf{D}\mathbf{G}\mathbf{w}, \quad (2)$$

where the matrices $\mathbf{A}$, $\mathbf{B}$, $\mathbf{D}$, and $\mathbf{G}$ can be straightforwardly derived, e.g. see Hokayem, Chatterjee, and Lygeros (2009) for details on the explicit construction of these matrices.

The control inputs are assumed to be constrained to a convex feasible region F_U described by a finite set of N_U linear inequalities

$$F_U = \{\mathbf{u} \mid \mathbf{H}\mathbf{u} \leq \mathbf{h}\}, \quad (3)$$

where $\mathbf{H} \in \mathbb{R}^{N_U \times N_m}$ and $\mathbf{h} \in \mathbb{R}^{N_U}$. The system states are also restricted to lie in a convex region F_X , which is defined by a collection of N_X linear inequality constraints

$$F_X = \bigcap_{i=1}^{N_X} \{\mathbf{x} \mid \mathbf{a}_i^\top \mathbf{x} \leq b_i\}, \quad (4)$$

with $\mathbf{a}_i \in \mathbb{R}^{(N+1)n}$ and $b_i \in \mathbb{R}$. For linear systems, the cost function is typically chosen to be quadratic, i.e.

$$V_N(x_0, \mathbf{u}, \mathbf{w}) = \mathbf{x}^\top \mathbf{Q} \mathbf{x} + \mathbf{u}^\top \mathbf{R} \mathbf{u}, \quad (5)$$

where $\mathbf{Q} \in \mathbb{S}_{++}^{(N+1)n}$ and $\mathbf{R} \in \mathbb{S}_{++}^{N_m}$ are specified weight matrices. As the distribution of the disturbances is unknown and could be unbounded, it cannot be guaranteed that there always exists a control action such that hard state constraints $\mathbf{x} \in F_X$ are satisfied. Therefore, $\mathbf{x} \in F_X$ is replaced with a *joint state chance constraint* of the form

$$\mathbb{P}(\mathbf{x} \notin F_X) \leq \delta, \quad (6)$$

where $\delta \in (0, 1)$ is the maximum probability of constraint violation. The SMPC problem is now stated as

SMPC Problem (P1)

$$\begin{aligned} & \min_{\mathbf{u}} \mathbb{E}[V_N(x, \mathbf{u}, \mathbf{w})] \\ & \text{subject to: } \mathbf{x} = \mathbf{A}\mathbf{x}_0 + \mathbf{B}\mathbf{u} + \mathbf{D}\mathbf{G}\mathbf{w} \\ & \quad \mathbf{u} \in F_U, \quad \mathbb{P}(\mathbf{x} \notin F_X) \leq \delta, \quad x_0 = x \end{aligned}$$

Problem P1 is solved online, given the most recently observed states x . Let $\mathbf{u}^*(x)$ be the optimal feedback control policy that solves P1 as a function of the initial states. The receding-horizon implementation of P1 indicates that only the first element of this policy, u_0^* , is applied to system (1). Note that $V_N(x_0, \mathbf{u}, \mathbf{w})$ is a random variable with an unknown distribution (as it is a function of $\mathbf{w}$). Here, the expected value of the value function is optimised to obtain a convex program.

There are two main difficulties that prevent the direct solution of P1. First, the control input $\mathbf{u}$ should be a causal feedback policy that is some function of the current and past states. In general, solving P1 over arbitrary functions of states is impractical using available optimal control approaches. Second, the distribution of the disturbances is unknown and possibly unbounded, which makes handling the hard input constraints and joint state chance constraints challenging.

To address these challenges, a certain class of causal feedback policies is adopted to define the control

policy $\mathbf{u}$. The adopted feedback policy allows for building feedback into the prediction to reduce uncertainty in the state predictions as well as directly handling the input constraints in the face of unbounded disturbances. In addition, the joint state chance constraints are approximated for arbitrary disturbance distributions using distributionally robust bounds that are only a function of the mean and variance of the stochastic disturbance (i.e. bounds do not depend on the full distribution of the disturbance, which is usually not known in practical applications).

3. Feedback parametrisation

A natural approach to obtaining a computationally tractable surrogate for P1 is to adopt an affine state feedback parametrisation for the control policy $\mathbf{u}$. Affine state feedback is in fact the solution to the linear-quadratic-Gaussian (LQG) problem, which minimises (5) in the absence of the input and state constraints. Solving P1 over an affine state feedback control policy, however, results in a nonconvex optimisation due to the product of the gains over time. An alternative parametrisation is an affine function of the sequence of past disturbances (Goulart et al., 2006)

$$u_i = \sum_{j=0}^{i-1} M_{i,j} G w_j + v_i, \quad \forall i = 0, \dots, N-1, \quad (7)$$

where $M_{i,j} \in \mathbb{R}^{m \times n}$ and $v_i \in \mathbb{R}^m$. This parametrisation yields convex optimisations, and is shown to be equivalent to the class of feedback control policies that are affine in the states (Goulart et al., 2006, Theorem 9). Using (7), the control policy $\mathbf{u}$ can be written as

$$\mathbf{u} = \mathbf{M}\mathbf{G}\mathbf{w} + \mathbf{v}, \quad (8)$$

where the block lower triangular matrix $\mathbf{M} \in \mathbb{R}^{mN \times nN}$ and stacked vector $\mathbf{v} \in \mathbb{R}^{mN}$ are given by

$$\mathbf{M} = \begin{bmatrix} 0 & \cdots & \cdots & 0 \\ M_{1,0} & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ M_{N-1,0} & \cdots & M_{N-1,N-2} & 0 \end{bmatrix}, \quad (9)$$

$$\mathbf{v} = (v_0, v_1, \dots, v_{N-1}). \quad (10)$$

The pair $(\mathbf{M}, \mathbf{v})$ then comprises the decision variables in P1.

A key challenge in using the feedback control policy (8) arises from guaranteeing the hard input constraints (3) in the presence of unbounded disturbances.

SMPC algorithms commonly overcome this difficulty by relaxing the hard input constraints to expectation-type constraints (Primbs & Hwang Sung, 2009) or probabilistic/chance constraints (Farina et al., 2015). These approaches, however, suffer from the fact that the computed inputs may not be feasible in practice, which will cause the inputs to saturate unknowingly to the controller. In this work, a saturated disturbance affine parametrisation of the form (Hokayem et al., 2009)

$$\mathbf{u} = \mathbf{M}\mathbf{G}\boldsymbol{\varphi}(\mathbf{w}) + \mathbf{v}, \quad (11)$$

is used, where, for any vector $\mathbf{z} = [z_1, \dots, z_n]^\top$, $\boldsymbol{\varphi}(\mathbf{z}) = [\varphi(z_1), \dots, \varphi(z_n)]^\top$ with $\varphi: \mathbb{R} \rightarrow \mathbb{R}$ denoting any function with the property $\sup_{a \in \mathbb{R}} |\varphi(a)| \leq \varphi_{\max}$ for some $\varphi_{\max} > 0$. The functions $\varphi(\cdot)$ are known as *saturation functions* (Hokayem et al., 2009). This definition implies that $\|\boldsymbol{\varphi}(\mathbf{w})\|_\infty \leq \varphi_{\max}$, which can be written as a polytope of the form

$$F_W = \{\mathbf{w} \mid \mathbf{S}\mathbf{w} \leq \mathbf{s}\},$$

with $\mathbf{S} \in \mathbb{R}^{N_W \times N_p}$ and $\mathbf{s} \in \mathbb{R}^{N_W}$, and allows the hard input constraints $\mathbf{u} \in F_U$ to be rewritten as

$$\mathbf{H}\mathbf{v} + \max_{\boldsymbol{\varphi}(\mathbf{w}) \in F_W} (\mathbf{H}\mathbf{M}\mathbf{G}\boldsymbol{\varphi}(\mathbf{w})) \leq \mathbf{h}, \quad (12)$$

where the maximisation is row-wise (i.e. maximum of each element in the vector). The following lemma, which is inspired by the work of Ben-Tal and Goryashko (2004) and Goulart et al. (2006), indicates that (12) can be defined by a set of linear inequalities.

Lemma 3.1 (Robust linear constraints with polytopic disturbance sets): *The input constraint (12) is represented exactly by linear inequalities $\mathbf{H}\mathbf{v} + \mathbf{Z}^\top \mathbf{s} \leq \mathbf{h}$ and $\mathbf{Z} \geq 0$ (element-wise) for any $\mathbf{Z}$ satisfying $\mathbf{Z}^\top \mathbf{S} = \mathbf{H}\mathbf{M}\mathbf{G}$.*

Proof: The proof follows from the concept of the dual norm as shown in, e.g. Bertsimas and Sim (2006). Let the i th row of the maximisation in (12) be the primal linear program. The corresponding dual linear program is

$$\min \mathbf{s}^\top \mathbf{z}_i, \quad \text{subject to: } \mathbf{S}^\top \mathbf{z}_i = (\mathbf{H}\mathbf{M}\mathbf{G})_{(i)}^\top, \quad \mathbf{z}_i \geq 0,$$

where $(\mathbf{H}\mathbf{M}\mathbf{G})_{(i)}$ denotes the i th row of $\mathbf{H}\mathbf{M}\mathbf{G}$ and $\mathbf{z}_i \in \mathbb{R}^{N_W}$ denotes the dual variables. By the strong duality theorem, it is known that

$$\max_{\boldsymbol{\varphi}(\mathbf{w}) \in F_W} (\mathbf{H}\mathbf{M}\mathbf{G})_{(i)} \boldsymbol{\varphi}(\mathbf{w}) \leq \mathbf{s}^\top \mathbf{z}_i$$

holds for any $\mathbf{z}_i$ satisfying the dual linear program constraints. Stacking the dual variables into a matrix $\mathbf{Z} \triangleq$

$[\mathbf{z}_1, \dots, \mathbf{z}_{N_U}]$ yields the inequality

$$\max_{\boldsymbol{\varphi}(\mathbf{w}) \in F_W} (\mathbf{H}\mathbf{M}\mathbf{G}\boldsymbol{\varphi}(\mathbf{w})) \leq \mathbf{Z}^\top \mathbf{s}$$

for any $\mathbf{Z} \geq 0$ satisfying $\mathbf{Z}^\top \mathbf{S} = \mathbf{H}\mathbf{M}\mathbf{G}$. Hence, the assertion of the lemma directly follows. $\blacksquare$

Remark 3.1: The saturation functions φ can be chosen in any way that ensures $|\varphi(z)| \leq \varphi_{\max}$ for all $z \in \mathbb{R}$. Common examples include sigmoid functions such as $\frac{z}{1+|z|}$, $\tanh(z)$, and $\frac{z}{\sqrt{1+z^2}}$ as well as the standard saturation function $\text{sat}(z) \triangleq \text{sign}(z) \min\{|z|, 1\}$. An interesting route for future research could involve developing a method (based on theory or heuristics) for choosing φ based on the system characteristics. In addition, exploring the tradeoff between control performance and optimisation cost for different saturation functions φ may be of interest.

Remark 3.2: The value of $\varphi_{\max}$ is selected offline, and can be chosen as any value greater than zero. The larger the value of $\varphi_{\max}$, the closer the saturated control law (11) is to the unsaturated control law (8); however, the size of the saturated disturbance set F_W will become larger. In other words, there is a tradeoff between selecting a large $\varphi_{\max}$ to mimic (8) and selecting a small $\varphi_{\max}$ to reduce the uncertainty in predicted satisfaction of the input constraints. A rule-of-thumb is to select $\varphi_{\max}$ to cover two or three standard deviations of the disturbances away from the mean.

4. Joint state chance constraints

It is generally impractical to ensure that the system states lie in the feasible region $\mathbf{x} \in F_X$ when the disturbances $\mathbf{w}$ are unbounded. Hence, the hard state constraints (4) must be replaced with the joint chance constraint (6) in P1. However, joint chance constraints are generally intractable and nonconvex (Ben-Tal, El Ghaoui, & Nemirovski, 2009; Bertsimas, Brown, & Caramanis, 2011). This is because evaluating joint chance constraints requires solving a multivariate integrate over a known distribution of the uncertainty.

To obtain a tractable deterministic surrogate for (6), this work uses Boole's inequality to bound the probability of violation of the joint chance constraint

$$\begin{aligned} \mathbb{P}(\mathbf{x} \notin F_X) &= \mathbb{P}\left(\mathbf{x} \in \bigcup_{i=1}^{N_X} \{\mathbf{x} \mid \mathbf{a}_i^\top \mathbf{x} > b_i\}\right) \\ &\leq \sum_{i=1}^{N_X} \mathbb{P}(\mathbf{a}_i^\top \mathbf{x} > b_i). \end{aligned} \quad (13)$$

This expression implies that the joint chance constraint (6) can be replaced with N_X individual chance constraints

of the form

$$\mathbb{P}(\mathbf{a}_i^\top \mathbf{x} > b_i) \leq \epsilon_i, \quad i = 1, \dots, N_X, \quad (14)$$

where $\epsilon_i \in [0, \delta]$ denotes the violation probability for the i th individual chance constraint. When the so-called *risk allocation* ϵ_i is chosen such that

$$\sum_{i=1}^{N_X} \epsilon_i \leq \delta,$$

is satisfied, then the joint chance constraint (6) will be satisfied according to (13). Two main approaches exist for defining the risk allocation. The first assumes a fixed risk allocation in which the values of ϵ_i are fixed a priori, usually using a uniform allocation $\epsilon_i = \delta/N_X$ (Nemirovski & Shapiro, 2006). Although this approach simplifies the optimisation problem, it may lead to significant conservatism in many situations as the prespecified risk may be better allocated to near-active constraints. To address this shortcoming, the second approach optimises the risk allocation by treating ϵ_i as decision variables in the optimisation (Blackmore & Ono, 2009).

The knowledge of the cumulative distribution function (cdf) of the disturbances $\mathbf{w}$ is required to exactly evaluate the individual chance constraints (14). Once this knowledge is available, the cdf of $\mathbf{x}$ can be straightforwardly determined using the linear relationship (2). Then, the probability of violating an individual chance constraint is given by

$$\mathbb{P}(\mathbf{a}_i^\top \mathbf{x} > b_i) = 1 - \text{cdf}_{\mathbf{a}_i^\top \mathbf{x}}(b_i).$$

This expression is, however, difficult to evaluate for general disturbances as their cdfs do not necessarily have a convex form. More importantly, the distribution of disturbances is not known in many practical applications. Hence, the Cantelli-Chebyshev inequality is used in this work to evaluate the individual chance constraints (14) for arbitrary distributions of disturbances when only their first two moments are known.

Lemma 4.1 (Cantelli-Chebyshev inequality) (Marshall & Olkin, 1979): *Let Z be a scalar random variable with finite variance. For every $c \geq 0$, it holds that*

$$\mathbb{P}(Z \geq \mathbb{E}[Z] + c) \leq \frac{\Sigma[Z]}{\Sigma[Z] + c^2}.$$

To apply the result of Lemma 4.1 to (14), it is assumed that some $\Delta b_i \geq 0$ exists such that

$$\mathbf{a}_i^\top \mathbb{E}[\mathbf{x}] + \Delta b_i \leq b_i. \quad (15)$$

The goal is to derive a lower bound on Δb_i . Notice that

$$\begin{aligned} \mathbb{P}(\mathbf{a}_i^\top \mathbf{x} > b_i) &\leq \mathbb{P}(\mathbf{a}_i^\top \mathbf{x} \geq \mathbf{a}_i^\top \mathbb{E}[\mathbf{x}] + \Delta b_i), \\ &\leq \frac{\mathbf{a}_i^\top \Sigma[\mathbf{x}] \mathbf{a}_i}{\mathbf{a}_i^\top \Sigma[\mathbf{x}] \mathbf{a}_i + \Delta b_i^2}. \end{aligned}$$

When this upper bound is less than or equal to ϵ_i , then the individual chance constraint (14) must be satisfied. This inequality implies that

$$\sqrt{\frac{1 - \epsilon_i}{\epsilon_i}} \sqrt{\mathbf{a}_i^\top \Sigma[\mathbf{x}] \mathbf{a}_i} \leq \Delta b_i. \quad (16)$$

Combining (15) with (16), the individual chance constraint (14) can be (conservatively) approximated by the deterministic constraint

$$\mathbf{a}_i^\top \mathbb{E}[\mathbf{x}] + \sqrt{\frac{1 - \epsilon_i}{\epsilon_i}} \sqrt{\mathbf{a}_i^\top \Sigma[\mathbf{x}] \mathbf{a}_i} \leq b_i, \quad (17)$$

which is guaranteed to hold for any distribution of the states $\mathbf{x}$. This result has been derived previously in Calafiore and El Ghaoui (2006, Theorem 3.1). The key contribution of this work is to combine this result with optimal risk allocation to substantially reduce the conservatism of (17). In the next section, we present an MPC formulation that incorporates these robust constraints while also including feedback in the predictions so that the controller can more effectively shape the state variance.

5. SMPC algorithm

This section uses the saturated affine disturbance parametrisation of the control inputs in conjunction with the risk allocation method for bounding the joint state chance constraint to obtain a tractable formulation for the SMPC problem P1. To this end, explicit expressions are first derived for the mean and covariance of the states. Using the system model (2), the dynamics for $\mathbb{E}[\mathbf{x}]$ and $\Sigma[\mathbf{x}]$ are described by

$$\mathbb{E}[\mathbf{x}] = \mathbf{A}\mathbf{x}_0 + \mathbf{B}\mathbb{E}[\mathbf{u}] + \mathbf{D}\mathbf{G}\mathbb{E}[\mathbf{w}], \quad (18a)$$

$$\begin{aligned} \Sigma[\mathbf{x}] &= \mathbf{B}\Sigma[\mathbf{u}]\mathbf{B}^\top + \mathbf{D}\mathbf{G}\Sigma[\mathbf{w}]\mathbf{G}^\top\mathbf{D}^\top + \mathbf{B}\sigma[\mathbf{u}, \mathbf{w}]\mathbf{G}^\top\mathbf{D}^\top \\ &\quad + \mathbf{D}\mathbf{G}\sigma[\mathbf{u}, \mathbf{w}]^\top\mathbf{B}^\top, \end{aligned} \quad (18b)$$

where $\mathbb{E}[\mathbf{w}] = \mathbf{1}_N \otimes \mathbb{E}[w]$ and $\Sigma[\mathbf{w}] = I_N \otimes \Sigma[w]$ are assumed to be known. The statistics of the control inputs $\mathbf{u}$ are derived from (11) as

$$\mathbb{E}[\mathbf{u}] = \mathbf{M}\mathbf{G}\mathbb{E}[\varphi(\mathbf{w})] + \mathbf{v}, \quad (19a)$$

$$\Sigma[\mathbf{u}] = \mathbf{M}\mathbf{G}\Sigma[\varphi(\mathbf{w})]\mathbf{G}^\top\mathbf{M}^\top, \quad (19b)$$

$$\sigma[\mathbf{u}, \mathbf{w}] = \mathbf{M}\mathbf{G}\sigma[\boldsymbol{\varphi}(\mathbf{w}), \mathbf{w}]. \quad (19c)$$

For any chosen saturation function $\boldsymbol{\varphi}(\cdot)$, the statistics $\mathbb{E}[\boldsymbol{\varphi}(\mathbf{w})]$, $\Sigma[\boldsymbol{\varphi}(\mathbf{w})]$, and $\sigma[\boldsymbol{\varphi}(\mathbf{w}), \mathbf{w}]$ can be straightforwardly computed by applying the saturation function to the data used to estimate the mean and covariance of the disturbances $\mathbf{w}$. The mean and variance equations (18) and (19) are used to recast P1 to the following program.

Deterministic surrogate for SMPC problem (P2)

$$\begin{aligned} \min_{\mathbf{M}, \mathbf{v}, \boldsymbol{\epsilon}} \quad & \mathbb{E}[\mathbf{x}]^\top \mathbf{Q} \mathbb{E}[\mathbf{x}] + \mathbb{E}[\mathbf{u}]^\top \mathbf{R} \mathbb{E}[\mathbf{u}] + \text{tr}(\mathbf{Q} \Sigma[\mathbf{x}]) + \text{tr}(\mathbf{R} \Sigma[\mathbf{u}]) \\ \text{subject to:} \quad & (\mathbb{E}[\mathbf{x}], \Sigma[\mathbf{x}]) \quad \text{given by (18)} \quad (P2.1) \\ & (\mathbb{E}[\mathbf{u}], \Sigma[\mathbf{u}], \sigma[\mathbf{u}, \mathbf{w}]) \quad \text{given by (19)} \quad (P2.2) \\ & \mathbf{M} \quad \text{satisfies (9)} \quad (P2.3) \\ & (\mathbf{M}, \mathbf{v}) \quad \text{satisfy Lemma 3.1} \quad (P2.4) \\ & \beta_i \quad = \sqrt{(1 - \epsilon_i)/\epsilon_i} \quad (P2.5) \\ & v_i \quad = \sqrt{\mathbf{a}_i^\top \Sigma[\mathbf{x}] \mathbf{a}_i} \quad (P2.6) \\ & \mathbf{a}_i^\top \mathbb{E}[\mathbf{x}] + \beta_i v_i \leq b_i \quad (P2.7) \\ & \epsilon_i \geq 0 \quad (P2.8) \\ & \sum_{i=1}^{N_X} \epsilon_i \leq \delta \quad (P2.9) \\ & x_0 = x \quad (P2.10) \\ & \forall i = 1, \dots, N_X \end{aligned}$$

In P2, the risk allocation is defined in terms of $\boldsymbol{\epsilon} = (\epsilon_1, \dots, \epsilon_{N_X})$.

Let $(\mathbf{M}^*(x), \mathbf{v}^*(x))$ be the optimal control policy (i.e. solution to P2) for any initial condition x . Receding-horizon implementation of the optimal control policy results in a time-invariant control law $\mu_N : \mathbb{R}^n \rightarrow \mathbb{R}^m$ defined as the first element of the optimal control sequence

$$\mu_N(x) = v_0^*(x). \quad (20)$$

Hence, the closed-loop response of system (1) is given by

$$x^+ = Ax + B\mu_N(x) + Gw. \quad (21)$$

Note that at every sampling time, the state x is measured and P2 is subsequently solved to obtain the control input $\mu_N(x)$, which defines the recursive SMPC algorithm (i.e. $\mathbf{M}^*(x)$ and $\mathbf{v}^*(x)$ are solved at every sampling time, and are only functions of the initial states).

5.1 Convexity analysis

Program P2 is nonconvex due to the multiplication of v_i and β_i in (P2.7), which makes simultaneous optimisation over the feedback gain $\mathbf{M}$ and risk allocation $\boldsymbol{\epsilon}$ a non-convex problem. An iterative strategy can be devised to solve P2 by taking advantage of the fact that the optimisation problem is convex when either $\mathbf{M}$ or $\boldsymbol{\epsilon}$ is fixed. To

prove this, notice that $\mathbb{E}[\mathbf{u}]$ and $\mathbb{E}[\mathbf{x}]$ are linear functions of the decision variables $\mathbf{v}$ and $\mathbf{M}$, while $\Sigma[\mathbf{u}]$ and $\Sigma[\mathbf{x}]$ are quadratic functions of $\mathbf{M}$. Thus, the objective function is quadratic in $\mathbf{v}$ and $\mathbf{M}$, and is a convex function of the decision variables since $\mathbf{Q}$ and $\mathbf{R}$ are assumed to be positive semidefinite and definite matrices, respectively.

Requiring $\mathbf{M}$ to be lower block triangular can be represented by linear equality constraints so that (P2.3) be convex. The hard input constraints (P2.4) are exactly

represented by a set of linear inequalities that are convex in $\mathbf{v}$ and $\mathbf{M}$ (see Lemma 3.1). Clearly, (P2.8)–(P2.10) are linear inequalities or equalities, which are convex.

Now, let us consider the surrogate expressions for the chance constraints (P2.7). When the feedback gain $\mathbf{M}$ is fixed, $\Sigma[\mathbf{x}]$ and $\Sigma[\mathbf{u}]$ must be constant matrices from (18) and (19), respectively. Therefore, v_i from (P2.6) will be constant for $i = 1, \dots, N_X$. Since $\boldsymbol{\epsilon}$ are still decision variables, (P2.7) reduces to $\mathbf{a}_i^\top \mathbb{E}[\mathbf{x}] + v_i \sqrt{(1 - \epsilon_i)/\epsilon_i} \leq b_i$. The first term is an affine function of $\mathbf{v}$. The second term is convex for any $\epsilon_i \in [0, 0.75]$, which can be verified by observing that the second derivative of $\sqrt{(1 - \epsilon_i)/\epsilon_i}$ is positive on this range. Since the sum of convex functions is a convex function, (P2.7) will be convex for any fixed $\mathbf{M}$ and any choice of $\delta \leq 0.75$.

On the other hand, when the risk allocation $\boldsymbol{\epsilon}$ is fixed, β_i from (P2.5) will be constant for $i = 1, \dots, N_X$. In this case, (P2.7) reduces to $\mathbf{a}_i^\top \mathbb{E}[\mathbf{x}] + \beta_i v_i \leq b_i$ where the first term is linear in $\mathbf{v}$ and the second term is linear in v_i . By substituting the expression for $\Sigma[\mathbf{x}]$ in (P2.6), this constraint can be rewritten as a second-order cone constraint

$$v_i = \left\| \begin{bmatrix} \Sigma[\boldsymbol{\varphi}(\mathbf{w})] & \sigma[\boldsymbol{\varphi}(\mathbf{w}), \mathbf{w}] \\ \star & \Sigma[\mathbf{w}] \end{bmatrix}^{1/2} \begin{bmatrix} \mathbf{G}^\top \mathbf{M}^\top \mathbf{B}^\top \mathbf{a}_i \\ \mathbf{G}^\top \mathbf{D}^\top \mathbf{a}_i \end{bmatrix} \right\|_2.$$

Substituting this expression into (P2.7) then renders P2 a convex second-order cone program for fixed ϵ .

5.2 Iterative optimisation strategy

The optimal control problem in P2 can be solved by optimising both the risk allocation ϵ and the control feedback gain $\mathbf{M}$. An iterative two-stage optimisation strategy is presented in Vitus and Tomlin (2011) to bisect the uniform risk allocation in the upper stage and to optimise the feedback gain with fixed uniform risk allocation in the lower stage. On the other hand, Ma et al. (2012) proposed optimising the risk allocation and feedback gain simultaneously using a tailored interior point method that exploits the sparse multistage structure of the nonconvex optimisation. Although these approaches were developed under different disturbance assumptions and control law parametrisations, they can be applied for solving P2 owing to the similar structure of the optimisation problems.

In this work, a simple iterative approach is proposed for solving P2, as summarised in Algorithm 1. The primary notion of Algorithm 1 is to solve for the optimal risk allocation given a fixed feedback gain and then solve for the optimal feedback gain given a fixed risk allocation. This approach is similar to the well-known DK iteration used in μ -synthesis problems (Balas, Doyle, Glover, Packard, & Smith, 1994). This technique is known as a (block-)coordinate descent algorithm, and has been applied more broadly to optimisation problems subject to bilinear matrix inequality (BMI) constraints (Simon, R-Ayerbe, Stoica, Dumur, & Wertz, 2011). Although this algorithm is not guaranteed to converge to a local optimum (as each iteration provides a solution that is optimal in the ‘directions’ of one subset of variables, but not in all directions), it is a commonly applied heuristic that performs well in practice.

Algorithm 1: Coordinate descent for SMPC

Require Initial feedback gain $\mathbf{M}^{(0)}$ and maximum number of iterations $I_{\max}$.

```

1: for  $i = 0$  to  $I_{\max} - 1$  do
2:   Solve convex optimisation P2 with fixed  $\mathbf{M} \leftarrow \mathbf{M}^{(i)}$  for the optimal risk allocation  $\epsilon^*$ 
3:   Set  $\epsilon^{(i+1)} \leftarrow \epsilon^*$ 
4:   Solve convex optimisation P2 with fixed  $\epsilon \leftarrow \epsilon^{(i+1)}$  for the optimal feedback gain  $\mathbf{M}^*$ 
5:   Set  $\mathbf{M}^{(i+1)} \leftarrow \mathbf{M}^*$ 
6: end for
```

Two choices have been made in Algorithm 1: (1) initialising the algorithm with a fixed feedback gain $\mathbf{M}^{(0)}$ (instead of a fixed risk allocation $\epsilon^{(0)}$ and switching the order of the optimisation problems), and (2) running the

algorithm for a fixed number of iterations instead of running until a prespecified tolerance has been met. The initial feedback gain $\mathbf{M}^{(0)}$ can be designed optimally without explicitly considering constraints using any of the numerous existing robust control methods (e.g. Kouvaritakis, Rossiter, & Schuurmans, 2000). Since there has been a plethora of work on offline feedback control design, initialising the algorithm based on a nearly optimal feedback gain is likely to yield better performance than initialising the algorithm using a fixed uniform risk allocation, which will rarely be optimal in practice. In addition, since adequate closed-loop performance can often be obtained with just a few iterations from a near-optimal choice of $\mathbf{M}^{(0)}$, it is best to run Algorithm 1 for a fixed number of iterations so as to ensure that the control inputs can be computed within a reasonable computation time. This idea has been widely used in the fast MPC literature to significantly reduce the cost of solving MPC problems online (Wang & Boyd, 2010).

5.3 Feasibility and stability

Due to the inclusion of input and state constraints, the region of attraction $\mathcal{X}_N$ for P2 (i.e. the set of initial conditions for which there exists a feasible solution to the optimisation problem) will be a subset of $\mathbb{R}^n$. When the disturbances lie in a compact set, recursive feasibility (as well as closed-loop stability) of the MPC problem can be guaranteed by defining terminal constraints and/or terminal penalties (Goulart et al., 2006; Mayne, Rawlings, Rao, & Scokaert, 2000).

The proposed SMPC approach, however, considers arbitrary stochastic disturbances with a (possibly) unbounded support. Hence, it is impractical to ensure that the states remain inside $\mathcal{X}_N$ in the presence of input constraints (Chatterjee & Lygeros, 2015). One approach for guaranteeing recursive feasibility for SMPC problems with unbounded disturbances is to choose between a closed-loop and open-loop initialisation strategy online (Farina et al., 2015). The key idea in this approach is to choose the closed-loop strategy when the problem is feasible and to choose the open-loop strategy (whose feasibility is guaranteed through a proper selection of terminal constraints) when the SMPC problem is infeasible for the most recently observed states. Although this approach guarantees recursively feasibility, it disregards the most recent state measurements, which may degrade closed-loop performance when the states are not in the region of attraction of the controller.

Alternatively, a backup controller can be applied when the system states leave the region of attraction of P2. In this case, a natural choice is to soften the state constraints in P2 since this will enable driving the states back

into $\mathcal{X}_N$ (Oldewurtel et al., 2008). To this end, the *exact penalty function* method can be used to ensure that the backup controller yields the same solution as the fully constrained MPC problem when it is feasible (Kerrigan & Maciejowski, 2000). This approach allows for solving a single optimisation problem instead of having to verify feasibility and decide which MPC problem to solve accordingly, as is the case in Farina et al. (2015).

Stability of stochastic linear systems (in a mean-square boundedness sense) in the presence of unbounded disturbances and bounded control inputs has been explored by Chatterjee, Ramponi, Hokayem, and Lygeros (2012). If the eigenvalues of the system matrix A lie inside the unit disc, the variance of the states is shown to be bounded as long as the disturbance has bounded variance. When A has eigenvalues on the unit disc (with equal geometric and algebraic multiplicities), the variance of states will be bounded provided that $\|u\|_2 \leq R$ for a sufficiently large R . However, if A has even one unstable eigenvalue and the system is subjected to unbounded stochastic disturbances along the directions of the unstable eigen-subspace of A , the linear system cannot be stabilised by means of bounded control inputs (Chatterjee et al., 2012). Inspired by these results, the feasibility and stability properties of the proposed SMPC approach outlined in Algorithm 1 are summarised below.

Theorem 5.1: *Let $\mathcal{X}_N^s$ be the domain of attraction for the softened version of P2 with positive slack variables added to the chance constraints (P2.7). Let $\mu_N^s(x)$ be the receding-horizon control policy defined similarly to (20) based on the softened P2. Let $x_{k+1} = Ax_k + B\mu_N^s(x_k) + Gw_k$ for $k \in \mathbb{N}_0$ describe the evolution of the closed-loop system from any initial condition $x_0 \in \mathcal{X}_N^s$. Then, $\mathcal{X}_N^s = \mathbb{R}^n$ such that the softened version of P2 will have a feasible solution for all initial states $x \in \mathbb{R}^n$, ensuring recursive feasibility of the optimisation problem. In addition, the evolution of the closed-loop system is mean-square bounded such that the sequence $(x_k)_{k \in \mathbb{N}_0}$ satisfies*

$$\sup_{k \in \mathbb{N}_0} \mathbb{E} [\|x_k\|_2^2] < \infty, \quad \forall x_0 \in \mathbb{R}^n, \quad (22)$$

if either of the following conditions is met:

- (i) A is Schur-stable (eigenvalues inside of the unit circle) and $\Sigma[w_k]$ is bounded;
- (ii) A is discrete-time, Lyapunov stable (eigenvalues on the unit disc with equal geometric and algebraic multiplicities), $\Sigma[w_k]$ is bounded, and $F_U \supseteq \{\mathbf{u} \in \mathbb{R}^{N_m} : \|\mathbf{u}\|_\infty \leq U_{\max}^*\}$ where $U_{\max}^*$ is defined based on reachability of the orthogonal decomposition of the state space and a bound on the fourth moment

of the disturbances as given in Cherukuri, Chatterjee, Hokayem, and Lygeros (2011).

Proof: The fact that $\mathcal{X}_N^s = \mathbb{R}^n$ follows directly from softening constraints in P2. In this case, the remaining hard constraints are related to the system dynamics (P2.1)–(P2.2) and the control inputs (P2.3)–(P2.4). These constraints are trivially satisfied with $\mathbf{M} = 0$ and some $\mathbf{v} \in F_U \neq \emptyset$.

The proof of condition (i) follows from Hokayem et al. (2009), and is summarised here for completeness. When A is Schur-stable, there exists a positive definite matrix P that satisfies $A^\top P A - P \leq -I_n$. From the closed-loop system dynamics, we have

$$\begin{aligned} \mathbb{E}_{x_k} [x_{k+1}^\top P x_{k+1}] &= x_k^\top A^\top P A x_k + 2x_k^\top P B \mu_N(x_k) + 2x_k^\top A^\top P G \mathbb{E}[w_k] \\ &\quad + \mu_N^\top(x_k) B^\top P B \mu_N(x_k) + 2\mu_N^\top(x_k) B^\top P G \mathbb{E}[w_k] \\ &\quad + \mathbb{E}[w_k^\top G^\top P G w_k]. \end{aligned}$$

Since the set F_U is bounded, there must exist some finite $U_b > 0$ for which $\mu_N(x_k) \in \mathbb{U}_\infty = \{u \in \mathbb{R}^m \mid \|u\|_\infty \leq U_b\}$. Taking advantage of Hölder's inequality and well-known norm properties, we can bound terms involving the control inputs on the right-hand side of the above inequality

$$\begin{aligned} 2x_k^\top P B \mu_N(x_k) &\leq 2\|B^\top P A x_k\|_\infty \|\mu_N(x_k)\|_1 \\ &\leq 2m\|B^\top P A\|_\infty U_b \|x_k\|_\infty, \\ \mu_N^\top(x_k) B^\top P B \mu_N(x_k) &\leq \|B^\top P B \mu_N(x_k)\|_\infty \|\mu_N(x_k)\|_1 \\ &\leq m\|B^\top P B\|_\infty U_b^2. \end{aligned}$$

Therefore, there exist positive constants $c_1 > 0$ and $c_2 > 0$ such that

$$\begin{aligned} \mathbb{E}_{x_k} [x_{k+1}^\top P x_{k+1}] &\leq x_k^\top A^\top P A x_k + 2c_1 \|x_k\|_\infty^2 + c_2 \\ &\leq x_k^\top P x_k - \|x_k\|_2^2 + 2c_1 \|x_k\|_\infty + c_2, \end{aligned}$$

where the Lyapunov inequality $A^\top P A - P \leq -I_n$ is used to derive the second inequality. Defining the compact set $\mathbb{D} = \{x \in \mathbb{R}^n \mid \|x\|_\infty \leq r\}$ for $r = \frac{1}{\theta}(c_1 + \sqrt{c_1^2 + c_2\theta})$, it can be verified that the following inequality must hold for all $x_k \notin \mathbb{D}$:

$$\begin{aligned} 2c_1 \|x_k\|_\infty + c_2 &\leq \theta \|x_k\|_\infty^2 \leq \theta \|x_k\|_2^2 \\ &\implies \mathbb{E}_{x_k} [x_{k+1}^\top P x_{k+1}] \\ &\leq x_k^\top P x_k - (1 - \theta) \|x_k\|_2^2. \end{aligned}$$

Combining the above expression with $x_k^\top P x_k \leq \lambda_{\max}(P) \|x_k\|_2^2$ leads to

$$\mathbb{E}_{x_k} [x_{k+1}^\top P x_{k+1}] \leq \left(1 - \frac{1 - \theta}{\lambda_{\max}(P)}\right) \|x_k\|_2^2, \quad \forall x_k \notin \mathbb{D}.$$

To ensure the multiplying constant above is positive, we must select some $\theta \in (\max\{0, 1 - \lambda_{\max}(P)\}, 1)$, which is guaranteed to exist due to the fact that A is Schur-stable. Then, from Hokayem et al. (2009, Lemma 8), the value function $V(x) = x^\top P x$ satisfies a *geometric drift condition* outside of the compact set $\mathbb{D}$ such that $\mathbb{E}_{x_0} [x_k^\top P x_k]$ is bounded for all $x_0 \in \mathbb{R}^n$. Since $\lambda_{\min}(P) \|x_k\|_2^2 \leq x_k^\top P x_k$, assertion (i) directly follows.

This proof does not hold for (ii) since there will not exist a θ for which the same geometric drift condition is satisfied. A detailed proof of condition (ii) is provided in Chatterjee et al. (2012), which involves a similar geometric drift condition, but requires a more complex argument based on the reachability and higher order moments of the system. ■

6. Case study

6.1 Problem description

The performance of the proposed SMPC approach is evaluated on a continuous *clostridial* ABE fermentation. The model of Haus et al. (2011) and Buehler and Mesbah (2016) is linearised around a desired steady-state operating point to obtain the system description in the form of (1) consisting of 12 states and 2 inputs. The state vector is defined by

$$x = [C_{AC}, C_A, C_{En}, C_{AaC}, C_{Aa}, C_{BC}, C_B, C_{An}, C_{Bn}, C_{Ad}, C_{Cf}, C_{Ah}]^\top,$$

where C denotes the concentration (mM) of species acetyl-CoA (AC), acetate (A), ethanol (En), acetoacetate-CoA (AaC), acetoacetate (Aa), butyryl-CoA (BC), butyrate (B), acetone (An), butanol (Bn), enzyme *adc* (Ad), enzyme *ctfA/B* (Cf), and enzyme *adhE* (Ah). The inputs are

$$u = [D, G_0]^\top,$$

where D is the dilution rate (h^{-1}) and G_0 is the inlet glucose concentration (mM). The system matrices A , B , and G are given in the Appendix.

The control problem is defined in terms of setpoint tracking for the ABE products (ethanol, acetone, and

butanol) concentrations. Hard input constraints

$$\begin{aligned} 0.005 \text{ h}^{-1} &\leq D \leq 0.145 \text{ h}^{-1} \\ 0 \text{ mM} &\leq G_0 \leq 80 \text{ mM} \end{aligned}$$

as well as joint chance constraints on acetate and butyrate (which are two key intermediate species in the metabolic pathway) with a maximum allowed violation of 20%

$$\begin{aligned} \mathbb{P}(13.83 \text{ mM} \leq C_A \leq 15.68 \text{ mM}, 10.55 \text{ mM} \leq C_B \\ \leq 12.30 \text{ mM}) \geq 0.8 \end{aligned}$$

are enforced. The control problem is converted into a regularisation problem by defining the states/inputs in terms of deviation variables with respect to the setpoint

$$\begin{aligned} x_{ss} &= [1.04, 14.83, 7.35, 1.66, 0, 2.95\text{e-}7, 11.55, \\ &\quad 43.51, 56.94, 1.46, 14.14, 37.98]^\top \\ u_{ss} &= [0.075, 40.0]^\top. \end{aligned}$$

The initial conditions (in deviation variables) are then given by

$$x_0 - x_{ss} = [0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ -5.16 \ 0 \ 0 \ 0]^\top.$$

Since butanol is the main product of interest, the weight matrix in the cost function of MPC is selected such that butanol has a weight that is 1000 times higher than that for ethanol and acetone

$$Q = \text{diag}(0, 0, 0.01, 0, 0, 0, 0, 0.01, 10, 0, 0, 0),$$

while the input penalty weight is chosen to be $R = \text{diag}(0.1, 0.1)$. The sampling time for this process is 1 h, and a horizon of $N = 10$ is chosen for subsequent simulations (unless otherwise noted).

6.2 Closed-loop performance

The control problem is solved using Algorithm 1, which iteratively solves the deterministic surrogate SMPC problem P2 for the optimal feedback control policy and risk allocation. The performance of the proposed approach is compared to that of a certainty equivalence MPC algorithm (in which the disturbance is set equal to its expected value for the purposes of prediction) and an SMPC algorithm with fixed uniform risk allocation. The fixed gain optimisation problem is solved using IPOPT, whereas the CVX package with the Mosek solver is utilised to solve the fixed risk allocation optimisation problem (Grant & Boyd, 2014).

Figure 1 shows the closed-loop response of butanol, under 300 realisations of the disturbances, obtained

Figure 1. Comparison of closed-loop response of butanol concentration (controlled state) under 300 disturbance realisations for (a) the proposed SMPC approach, (b) certainty equivalence MPC, and (c) SMPC with fixed uniform risk allocation. The setpoint is shown with a black dashed line.

using the three control algorithms. The proposed SMPC approach shows comparable setpoint tracking performance to certainty equivalence MPC wherein both algorithms keep the butanol concentration around its setpoint with minimal variation. On the other hand, SMPC with fixed uniform risk allocation yields the worst performance due to the relatively large variation in butanol concentration. The poor setpoint tracking performance can be attributed to conservative constraint handling since the SMPC algorithm with fixed uniform risk allocation attempts to fulfil each individual chance constraint (decomposed from the joint chance constraint) with equal risk regardless of how close the states are to any particular constraint.

Figure 2 shows the closed-loop response of acetate and butyrate concentrations, under 300 realisations of the disturbances, obtained using the three control algorithms. The proposed SMPC algorithm results in a maximum constraint violation of 14%, which is below the allowed violation of 20%. This result is expected as the proposed algorithm guarantees satisfaction of the joint chance constraint, regardless of the distribution of the

disturbance, while not being overly conservative since the risk allocation is optimised online. On the other hand, certainty equivalence MPC yields 79% constraint violation, which is much larger than the allowed 20%, since it does not explicitly account for uncertainty. The improved setpoint tracking performance observed in Figure 1 for certainty equivalence MPC is due to this large constraint violation. SMPC with fixed uniform risk allocation exhibits 0% constraint violation at all times, indicating very conservative handling of the joint chance constraint leading to the poor setpoint tracking performance. Overall, the proposed SMPC approach provides the best control performance, while effectively satisfying the hard input and joint chance constraints.

6.3 Convergence and optimality

As discussed in Section 5, the convergence of Algorithm 1 has not been established. However, the convergence properties of the proposed SMPC algorithm are evaluated for the case study at hand in order to illustrate possible advantages of Algorithm 1 in practice. For multiple

Figure 2. Comparison of the closed-loop response of acetate and butyrate concentrations (joint-constrained states) under 300 disturbance realisations for (a) the proposed SMPC approach, (b) certainty equivalence MPC, and (c) SMPC with fixed uniform risk allocation. The four state constraints are shown with black dashed lines. The chance constraints are enforced jointly with a maximum allowed violation of 20%.

Figure 3. Value of the cost function as a function of number of iterations of Algorithm 1. Iteration 0 corresponds to the initial condition. The cost function converges within two iterations.

sampling times in a closed-loop simulation, Figure 3 shows the value of the cost function as a function of the number of iterations of Algorithm 1. At all time points, the cost function converges to its minimum value within

at most two iterations (most points require only one iteration). This indicates that, for this case study, only two iterations of Algorithm 1 are sufficient for convergence. Note that, based on this analysis, the results of the proposed SMPC algorithm shown in Section 6.2 utilised two iterations of Algorithm 1.

Now that (practical) convergence of Algorithm 1 has been demonstrated, we look to investigate the optimality of the converged solutions by comparing the results to a nonlinear gradient-based optimisation method. To this end, the general nonlinear Matlab solver `fmincon` is used for comparison. The horizon is decreased from $N = 10$ to $N = 2$ for these simulations due to the large computational cost of `fmincon` since there are more than 2400 decision variables when $N = 10$ and only 36 decision variables when $N = 2$. Note that variants of the gradient-based algorithms utilised by `fmincon` provide guaranteed convergence to a local minimum.

Figure 4 shows the value of the ‘optimal’ cost function over a closed-loop simulation using the proposed SMPC algorithm (with two iterations) and the

Figure 4. Value of the ‘optimal’ cost function obtained with the proposed SMPC algorithm and with a general nonlinear solver over an entire closed-loop simulation. Both optimisation algorithms were initialised with the unconstrained LQR solution at all times for consistency.

nonlinear optimisation solver `fmincon`. For consistency, both optimisation methods were supplied the same initial conditions derived from the unconstrained LQR solution. Figure 4 indicates that the value of the ‘optimal’ cost function in the proposed algorithm is always less than or equal to that of `fmincon`. The proposed algorithm results in cost function values that are almost an order of magnitude lower than `fmincon` (note that the y -axis is displayed in logarithmic scale). Similar results were observed for different initial conditions, including a cold-start of all zeros. These results are likely due to the nonlinear optimiser being trapped near a local solution. This was verified by solving the nonlinear program for different initial conditions wherein the optimiser converged to different solutions (not shown here).

The results indicate that the proposed SMPC approach in Algorithm 1 (that dynamically allocates risk to the constraints) is robust to the initialisation strategy, which makes intuitive sense due to the properties of P2. When the states are far away from their constraints (acetate and butyrate in this case), the problem is insensitive to the chosen risk allocation. In this case, P2 is convex and both the proposed algorithm and `fmincon` reach the global optimum, as seen in Figure 4 where the curves overlap. However, when constraints become active, the proposed algorithm optimises in the directions of ϵ and $(\mathbf{M}, \mathbf{v})$ separately as opposed to a general nonlinear solver that optimises in all directions simultaneously. Hence, the directions in which the decision variables can be updated will be effectively constrained in the proposed algorithm, likely reducing the chance of getting stuck in local solutions (defined as zero gradient with respect to all decision

Table 1. Comparison of CPU time of the proposed SMPC algorithm and a general nonlinear solver. Results were obtained on a desktop computer with a 4.0 GHz Core i7 CPU.

Optimisation method	Mean CPU time per MPC iteration (s)	Total CPU time (s)
Algorithm 1	2 ± 2	1589
Nonlinear solver	807 ± 316	32,278

variables). This is supported by Figure 4, wherein `fmincon` yields a higher ‘optimal’ cost than the proposed algorithm at the sampling times where the constraints are (nearly) active at the beginning and at the end of the simulation.

Finally, the computational cost of the proposed SMPC algorithm (with two iterations) is compared to the nonlinear optimisation solver `fmincon` in Table 1, which shows the mean CPU time per MPC iteration ($\pm$ standard deviation) and the total CPU time for one closed-loop simulation. The proposed algorithm takes significantly less time than `fmincon` (over 400 times faster) due to the fact that only a small number of convex optimisation problems need to be solved in comparison with a general nonlinear optimisation problem. Note that tailored interior point methods such as that proposed by Ma et al. (2012) can provide improved performance over `fmincon`; however, the same applies to the convex optimisation problems in Algorithm 1.

7. Conclusions

This paper presents an MPC approach for linear systems subject to arbitrary (possibly unbounded) stochastic disturbances with known mean and variance. The approach enables: (1) accounting for hard input constraints and joint state chance constraints under a feedback prediction, (2) efficient handling of joint chance constraints by using the Cantelli–Chebyshev inequality in conjunction with risk allocation, and (3) determining the optimal feedback gain and risk allocation by iteratively solving convex optimisations. The proposed SMPC approach is demonstrated on a continuous ABE fermentation process with 12 states, and its performance is compared to certainty equivalence MPC and SMPC with fixed risk allocation. Convergence and optimality properties of the proposed iterative optimisation strategy are demonstrated in the context of this case study.

Disclosure statement

No potential conflict of interest was reported by the authors.

References

- Balas, G.J., Doyle, J.C., Glover, K., Packard, A., & Smith, R. (1994). *μ -Analysis and synthesis toolbox*. Natick, MA: The MathWorks Inc.
- Bemporad, A., & Morari, M. (1999). Robust model predictive control: A survey. In *Robustness in identification and control* (pp. 207–226). London: Springer.
- Ben-Tal, A., El Ghaoui, L., & Nemirovski, A. (2009). *Robust optimization*. Princeton, NJ: Princeton University Press.
- Ben-Tal, A., & Goryashko, A. (2004). Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99, 351–376.
- Bertsimas, D., Brown, D.B., & Caramanis, C. (2009). Theory and applications of robust optimization. *SIAM Review*, 53, 464–501.
- Bertsimas, D., & Sim, M. (2006). Tractable approximations to robust conic optimization problems. *Mathematical Programming*, 107, 5–36.
- Blackmore, L., & Ono, M. (2009). Convex chance constrained predictive control without sampling. In *Proceedings of the AIAA guidance, navigation, and control conference* (pp. 7–21). Chicago, Illinois.
- Blackmore, L., Ono, M., Bektassov, A., & Williams, B.C. (2010). A probabilistic particle-control approximation of chance-constrained stochastic predictive control. *IEEE Transactions on Robotics*, 26, 502–517.
- Buehler, E.A., & Mesbah, A. (2016). Kinetic study of acetone-butanol-ethanol fermentation in continuous culture. *PLOS One*, 0158243.
- Calafiore, G.C., & El Ghaoui, L. (2006). On distributionally robust chance-constrained linear programs. *Journal of Optimization Theory and Applications*, 130, 1–22.
- Calafiore, G.C., & Fagiano, L. (2013). Robust model predictive control via scenario optimization. *IEEE Transactions on Automatic Control*, 58, 219–224.
- Chatterjee, D., & Lygeros, J. (2015). On stability and performance of stochastic predictive control techniques. *IEEE Transactions on Automatic Control*, 60, 509–514.
- Chatterjee, D., Ramponi, F., Hokayem, P., & Lygeros, J. (2012). On mean-square boundedness of stochastic linear systems with bounded controls. *System and Control Letters*, 61, 375–280.
- Cherukuri, A., Chatterjee, D., Hokayem, P., & Lygeros, J. (2011). Stochastic receding horizon control: Stability results. In *Proceedings of the 18th IFAC world congress on automatic control, Milano, Italy* (pp. 150–155). Milano, Italy.
- Farina, M., Giulioni, L., Magni, L., & Scattolini, R. (2015). An approach to output-feedback MPC of stochastic linear discrete-time systems. *Automatica*, 55, 140–149.
- Goulart, P.J., Kerrigan, E.C., & Maciejowski, J.M. (2006). Optimization over state feedback policies for robust control with constraints. *Automatica*, 42, 523–533.
- Grant, M., & Boyd, S. (2014). CVX: Matlab software for disciplined convex programming, version 2.1. March 2014. Retrieved from <http://cvxr.com/cvx>
- Haus, S., Jabbari, S., Millat, T., Janssen, H., Fischer, R.J., Bahl, H., ... Wolkenhauer, O. (2011). A systems biology approach to investigate the effect of pH-induced gene regulation on solvent production by *Clostridium acetobutylicum* in continuous culture. *BMC Systems Biology*, 5(10), 21247470.
- Hokayem, P., Chatterjee, D., & Lygeros, J. (2009). On stochastic receding horizon control with bounded control inputs. In *Proceedings of the 48th IEEE conference on decision and control held jointly with the 28th chinese control conference, Shanghai, China* (pp. 6359–6364). Shanghai, China.
- Hokayem, P., Cinquemani, E., Chatterjee, D., Ramponi, F., & Lygeros, J. (2012). Stochastic receding horizon control with output feedback and bounded controls. *Automatica*, 48, 77–88.
- Kerrigan, E.C., & Maciejowski, J.M. (2000). Soft constraints and exact penalty functions in model predictive control. In *Proceedings of UKACC international conference (Control 2000)*. Cambridge, UK.
- Kouvaritakis, B., & Cannon, M. (2015). Developments in robust and stochastic predictive control in the presence of uncertainty. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering*, 1(2), 021003.
- Kouvaritakis, B., Rossiter, J.A., & Schuurmans, J. (2000). Efficient robust predictive control. *IEEE Transactions on Automatic Control*, 45, 1545–1549.
- Ma, Y., Vichik, S., & Borrelli, F. (2012). Fast stochastic MPC with optimal risk allocation applied to building control systems. In *Proceedings of the 51st IEEE conference on decision and control, Maui, HI, USA* (pp. 7559–7564). Maui, HI.
- Marshall, A., & Olkin, I. (1979). *Inequalities: Theory and majorization and its applications*. New York, NY: Academic Press.
- Mayne, D.Q. (2014). Model predictive control: Recent developments and future promise. *Automatica*, 50, 2967–2986.
- Mayne, D.Q., Rawlings, J.B., Rao, C.V., & Sokaert, P.O. M. (2000). Constrained model predictive control: Stability and optimality. *Automatica*, 36, 798–814.
- Mesbah, A. (2016). Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Systems Magazine*, 36, 30–44.
- Nemirovski, A., & Shapiro, A. (2006). Convex approximations of chance constrained programs. *SIAM Journal on Optimization*, 17, 969–996.
- Oldewurtel, F., Jones, C.N., & Morari, M. (2008). A tractable approximation of chance constrained stochastic MPC based on affine disturbance feedback. In *Proceedings of the 47th IEEE conference on decision and control, Cancun, Mexico* (pp. 4731–4736). Cancun, Mexico.
- Paulson, J.A., Streif, S., & Mesbah, A. (2015). Stability for receding-horizon stochastic model predictive control. In *Proceedings of the american control conference, Chicago, IL, USA* (pp. 937–943). Chicago, IL.
- Primbs, J.A., & Hwan Sung, C. (2009). Stochastic receding horizon control of constrained linear systems with state and control multiplicative noise. *IEEE Transactions on Automatic Control*, 54, 221–230.
- Simon, E., R-Ayerbe, P., Stoica, C., Dumur, D., & Wertz, V. (2011). LMIs-based coordinate descent method for solving BMIs in control design. In *Proceedings of the 18th IFAC world congress on automatic control, Milano, Italy* (pp. 10180–10186). Milano, Italy.
- Vitus, M.P., & Tomlin, C.J. (2011). On feedback design and risk allocation in chance constrained control. In *Proceedings of the 50th IEEE conference on decision and control and European control conference, Orlando, FL, USA* (pp. 734–739). Orlando, FL.
- Wang, Y., & Boyd, S. (2010). Fast model predictive control using online optimization. *IEEE Transactions on Control Systems Technology*, 18, 267–278.

Appendix. ABE fermentation system model

The system matrices are given by

$$A = 10^{-2} \times \begin{bmatrix} 51 & 5.3 & 0 & 29 & 0 & 0 & -2.7 & 0 & 0 & 0 & 3.4 & -10 \\ -2.5 & 85 & 0 & -43 & 0 & 0 & 3.4 & 0 & 0 & 0 & -5.1 & -0.014 \\ 37 & 1.8 & 93 & 12 & 0 & 0 & -0.59 & 0 & 0 & 0 & 1.4 & 1.1 \\ 3.3 & -4.5 & 0 & 18 & 0 & 0 & -5.0 & 0 & 0 & 0 & -8.8 & -0.036 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -0.030 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -2.1 & 2.6 & 0 & -35 & 0 & 0 & 85 & 0 & 0 & 0 & -4.1 & 0.012 \\ 4.6 & 4.9 & 0 & 78 & 93 & 0 & 4.9 & 93 & 0 & 0.030 & 9.1 & -0.027 \\ 2.1 & -2.6 & 0 & 35 & 0 & 93 & 8.3 & 0 & 93 & 0 & 4.1 & -0.012 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 93 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 93 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 93 \end{bmatrix},$$

$$B = \begin{bmatrix} -1.7 & -13 & -7.8 & 0.94 & -1.4\text{e-}4 & 0 & -10 & -45 & -51 & -1.4 & -14 & -37 \\ 57\text{e-}6 & 0 & 16\text{e-}6 & 2\text{e-}6 & 0 & 0 & 0 & 1.5\text{e-}6 & 0 & 0 & 0 & 0 \end{bmatrix}^{\top},$$

and

$$G = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}^{\top}.$$