

Sample-average model predictive control of uncertain linear systems

2016 IEEE Conference on Decision and Control, Las Vegas, NV

Kevin J. Kircher and K. Max Zhang

Abstract—We propose a randomized implementation of stochastic model predictive control (MPC). As a proxy for the expected cost, which may not be efficiently computable, the algorithm minimizes the empirical average cost under N random samples of the uncertain influences on the system. The setting is an imperfectly-observed linear system with multiplicative and additive uncertainty, convex, deterministic control constraints, and convex costs that may include penalties on state constraint violations. In this setting, each sample-average MPC subproblem is a feasible convex program that, under mild regularity conditions, yields consistent estimators of the stochastic MPC subproblem’s optimal value and minimizers. Under stronger assumptions, the full sample-average MPC control trajectory is asymptotically optimal for stochastic MPC as $N \rightarrow \infty$. A numerical example shows that even for small N , sample-average MPC can significantly improve performance relative to certainty-equivalent MPC.

I. INTRODUCTION

Model predictive control (MPC) can be viewed as a heuristic method for approximately solving optimal control problems in the presence of uncertainty. At each MPC iteration, an optimal control subproblem is solved over a truncated, receding horizon to generate a planned control trajectory. The first control action in the trajectory is implemented, the system evolves, and the process repeats.

MPC has seen widespread adoption in the chemical process industry, and is gaining traction in others. [1] This popularity is likely due to the method’s ability to handle systems with many inputs and outputs, hard constraints, predictable disturbances, and explicit performance objectives. Most MPC implementations assume certainty equivalence: Each subproblem is solved for one nominal trajectory of the uncertain influences on the system, and uncertainty is otherwise ignored. There are two active research tracks that aim to improve MPC performance by using more sophisticated descriptions of uncertainty. These tracks are *robust* and *stochastic* MPC.

In robust MPC, each subproblem’s uncertain parameters – *e.g.*, state estimates, disturbances, noise, and uncertain entries in system matrices – are assumed to belong to a bounded set, constraints must hold for all possible parameter realizations, and the worst-case cost is minimized. In stochastic MPC, general probability distributions on uncertain parameters are allowed (including those with unbounded support), constraints must hold with a given probability, and the expected cost is minimized. In this paper, we focus on stochastic MPC.

Stochastic MPC can be implemented efficiently in some special cases [2], but is intractable in general. This is due in part to chance constraints – *i.e.*, constraints that must hold with a given probability – which typically make the stochastic MPC subproblem’s feasible region nonconvex. Another complication is the objective function’s expectation integral, for which a closed-form expression may not be available, and whose evaluation may require computationally expensive, high-dimensional numerical integration. Stochastic MPC also requires the joint distribution of all uncertain parameters at all time steps. In practice, this distribution may be unavailable or itself uncertain.

Recently, researchers have used randomized algorithms to develop tractable approximations of chance-constrained MPC for uncertain linear systems. [3], [4], [5], [6], [7], [8], [9] The key idea behind this approach is to replace the single deterministic uncertain parameter trajectory used in certainty-equivalent MPC by N randomly sampled trajectories (referred to as *scenarios*). Each chance constraint is replaced by the N deterministic constraints generated by the scenarios. Based on theory developed by Calafiore, Campi, and Garatti in [10], [11], [12] for randomized convex programs, a lower bound on N can be obtained to guarantee, at a given confidence level, that chance constraints will be satisfied. No explicit distributional information is needed, as long as enough scenarios can be procured.

The focus of references [3] through [9] is the use of random scenarios to satisfy chance constraints. In this paper, we complement this work by exploring the use of random scenarios to reduce cost. In particular, we investigate the effect of minimizing the sample-average cost as a proxy for the expected cost. This natural idea, known as sample-average approximation (SAA), has been studied in the stochastic programming community since the early 1990s. [13]

SAA has been applied to stochastic MPC in [3], [9], [14], [15]. Of particular relevance are [3] and [9], which consider linear systems with multiplicative and additive uncertainty and general convex costs and constraints. In [3], Blackmore *et al.* propose a sample-based stochastic MPC algorithm and demonstrate its ability to satisfy chance constraints in two numerical case studies. In [9], Schildbach *et al.* obtain a sharpened bound on the required number of scenarios when only the time-averaged probability of constraint violation must be limited. The setting of this paper is similarly general to [3] and [9]. Our consistency analyses further justify the use of SAA in randomized MPC.

To our knowledge, [3] and [14] contain the only SAA consistency analyses in the stochastic MPC literature. Both

The authors are with the Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14853. {kj82, kz33}@cornell.edu

[3] and [14] show that the sample-average cost is a consistent estimator of the expected cost. In [14], Batina *et al.* derive from Hoeffding's inequality a lower bound on N to guarantee that, with high probability, the sample-average cost closely approximates the expected cost. The analyses in [3] and [14] do not, however, guarantee that the *minimizers* or *optimal value* of a sample-average MPC (SAMPC) subproblem approximate the minimizers or optimal value of the corresponding stochastic MPC subproblem. In other words, there are currently no guarantees that an SAMPC control action is in any sense optimal. Our primary goal is to fill this gap by adapting recent theory developed for more general stochastic programs by Mak, Kleywegt, Shapiro, *et al.* in [16], [17], [18].

This paper is organized as follows. In §II, we define the stochastic optimal control problem of interest, propose a SAMPC algorithm to solve it approximately, and demonstrate the tractability of SAMPC. In §III, we discuss the consistency and bias of the SAMPC estimators. We present a numerical example in §IV and conclude in §V.

II. PROBLEM STATEMENT

A. Notation

We denote the space of real n -dimensional vectors by \mathbf{R}^n . For vectors $x \in \mathbf{R}^n$ and $y \in \mathbf{R}^m$, let (x, y) be the column vector in \mathbf{R}^{n+m} obtained by stacking x above y . We denote expectation by \mathbf{E} , probability by \mathbf{P} , a generic norm by $\|\cdot\|$, and the Euclidean norm by $\|\cdot\|_2$. We define the distance from a vector x to a set Y by $\text{dist}(x, Y) := \inf \{\|x - y\| \mid y \in Y\}$, the deviation of a set X from Y by $\text{dev}(X, Y) := \sup \{\text{dist}(x, Y) \mid x \in X\}$, and the epigraph of a function $f : \mathbf{R}^n \rightarrow \mathbf{R}$ by $\text{epi } f := \{(x, s) \in \mathbf{R}^n \times \mathbf{R} \mid f(x) \leq s\}$.

B. Stochastic optimal control problem

We consider the stochastic optimal control problem

$$\begin{aligned} & \text{minimize} && \mathbf{E} \left[\sum_{t \in \mathcal{T}} g_t(u_t, x_t, \delta_t) \right] \\ & \text{subject to} && x_{t+1} = A_t(\delta_t)x_t + B_t(\delta_t)u_t + w_t(\delta_t) \quad \forall t \in \mathcal{T} \\ & && y_t = C_t(\delta_t)x_t + D_t(\delta_t)u_t + v_t(\delta_t) \quad \forall t \in \mathcal{T} \\ & && u_t = \mu_t(y_0, \dots, y_t) \in U_t \quad \forall t \in \mathcal{T}. \end{aligned} \tag{1}$$

Here t indexes discrete time, $\mathcal{T} := \{0, \dots, T-1\}$, T is the control horizon, $x_t \in \mathbf{R}^n$ is the state of a dynamical system, $u_t \in \mathbf{R}^m$ is the control, and $y_t \in \mathbf{R}^q$ is the measurement. The expectation is taken with respect to the joint distribution of the initial state x_0 and the uncertain parameters $\delta_0, \dots, \delta_{T-1} \in \mathbf{R}^p$. The system matrices $A_t(\delta_t)$, $B_t(\delta_t)$, $C_t(\delta_t)$, and $D_t(\delta_t)$, the disturbance $w_t(\delta_t)$, and the noise $v_t(\delta_t)$ have suitable dimensions. The optimization variables are the functions $\mu_t : \mathbf{R}^{(t+1)q} \rightarrow \mathbf{R}^m$, which map measurements into controls.

Assumption 1 (Uncertainty). *For each $t \in \mathcal{T}$, the state x_t and uncertain parameter δ_t are random variables with distributions supported on $X_t \subseteq \mathbf{R}^n$ and $\Delta_t \subseteq \mathbf{R}^p$. For any MPC horizon $K \in \{1, \dots, T-t\}$, independent, identically distributed (i.i.d.) samples can be procured from*

the distribution of the uncertain parameter trajectory $\delta_t := (x_t, \delta_t, \dots, \delta_{t+K-1}) \in \mathbf{R}^{n+Kp}$, conditioned on the information $\mathcal{I}_t := (u_0, \dots, u_{t-1}, y_0, \dots, y_t) \in \mathbf{R}^{tm+(t+1)q}$. This conditional distribution does not depend on u_t, \dots, u_{t+K-1} .

We do not assume that explicit distributional information is available. Distributions may be continuous or discrete, have bounded or unbounded support, *etc.* The system matrices, disturbance, and noise depend arbitrarily on the uncertain parameters.

Assumption 2 (Costs and constraints). *For all $t \in \mathcal{T}$, the stage cost $g_t(\cdot, \cdot, \delta_t)$ is almost surely convex and its expected value is well-defined and finite. The control constraint sets U_t are closed, nonempty and convex.*

Although problem (1) includes no explicit state constraints, the stage costs can include convex 'soft constraints' such as

$$\eta(\text{dist}(x_t, X_t(\delta_t))), \tag{2}$$

where $\eta : \mathbf{R} \rightarrow \mathbf{R}$ is a convex nondecreasing function and $X_t(\delta_t)$ is a nonempty convex set for almost every $\delta_t \in \Delta_t$. More generally, problem (1) could also include chance constraints such as

$$\mathbf{P} \{x_t \in X_t(\delta_t)\} \geq 1 - \alpha,$$

where $\alpha \in (0, 1)$ is a design parameter. To isolate the effect of samples on cost, however, we do not consider chance constraints in this paper. They are readily handled in the current framework using techniques developed in references [3] through [9].

While analytically solvable in a few special cases (notably the LQG system), problem (1) is intractable in general: The optimization variables μ_0, \dots, μ_{T-1} are infinite-dimensional, the expectation integrals may be ill-defined or expensive to compute, and the imperfect state information generates challenging exploration/exploitation trade-offs. For these reasons, it is common to separate the problems of estimation and control, and to solve each by some suboptimal scheme. In this paper, we restrict our attention to approximate solution of the control problem via MPC.

C. Certainty-equivalent model predictive control

In certainty-equivalent MPC, the subproblem at stage t is solved under the nominal scenario $\hat{\delta}_t := (\hat{x}_t, \hat{\delta}_{0|t}, \dots, \hat{\delta}_{K-1|t}) \in \mathbf{R}^{n+Kp}$. Here K is the MPC horizon and $\hat{\delta}_{k|t}$ is a prediction of δ_{t+k} (typically its expected value, conditioned on \mathcal{I}_t). Letting $\mathcal{K} = \{0, \dots, K-1\}$, the certainty-equivalent MPC subproblem at stage t is

$$\begin{aligned} & \text{minimize} && \sum_{k \in \mathcal{K}} g_{t+k}(u_{k|t}, x_{k|t}, \hat{\delta}_{k|t}) \\ & \text{subject to} && x_{0|t} = \hat{x}_t \\ & && x_{k+1|t} = A_{t+k}(\hat{\delta}_{k|t})x_{k|t} + B_{t+k}(\hat{\delta}_{k|t})u_{k|t} \\ & && \quad + w_{t+k}(\hat{\delta}_{k|t}) \quad \forall k \in \mathcal{K} \\ & && u_{k|t} \in U_{t+k} \quad \forall k \in \mathcal{K}. \end{aligned}$$

The nominal states $x_{k|t}$ can be eliminated using the equality constraints, reducing the optimization variables to the planned control trajectory $\mathbf{u}_t := (u_{0|t}, \dots, u_{K-1|t}) \in \mathbf{R}^{Km}$.

This gives an equivalent statement of the certainty-equivalent MPC subproblem:

$$\begin{aligned} & \text{minimize} && \psi_t(\mathbf{u}_t, \hat{\boldsymbol{\delta}}_t) \\ & \text{subject to} && \mathbf{u}_t \in \mathbf{U}_t, \end{aligned} \quad (3)$$

where

$$\mathbf{U}_t := U_t \times \cdots \times U_{t+K-1} \subseteq \mathbf{R}^{Km}$$

$$\psi_t(\mathbf{u}, \boldsymbol{\delta}) := \sum_{k \in \mathcal{K}} g_{t+k}(E_{k|t}\mathbf{u}, F_{k|t}(\boldsymbol{\delta})\mathbf{u} + f_{k|t}(\boldsymbol{\delta}), G_{k|t}\boldsymbol{\delta}).$$

Here the matrices $E_{k|t}$ and $G_{k|t}$ satisfy $E_{k|t}\mathbf{u}_t = u_{k|t}$ and $G_{k|t}\boldsymbol{\delta}_t = \delta_{t+k}$. Similarly, the functions $F_{k|t}$ and $f_{k|t}$ satisfy

$$\begin{aligned} F_{k|t}(\boldsymbol{\delta}_t)\mathbf{u}_t &= A_{k-1}^\delta \cdots A_1^\delta B_0^\delta u_{0|t} + A_{k-1}^\delta \cdots A_2^\delta B_1^\delta u_{1|t} \\ &+ \cdots + A_{k-1}^\delta B_{k-2}^\delta u_{k-2|t} + B_{k-1}^\delta u_{k-1|t} \\ f_{k|t}(\boldsymbol{\delta}_t) &= A_{k-1}^\delta \cdots A_0^\delta x_t + A_{k-1}^\delta \cdots A_1^\delta w_0^\delta \\ &+ A_{k-1}^\delta \cdots A_2^\delta w_1^\delta + \cdots + A_{k-1}^\delta w_{k-2}^\delta + w_{k-1}^\delta. \end{aligned}$$

For brevity, we have introduced the shorthand A_k^δ for $A_{t+k}(\delta_{t+k})$, B_k^δ for $B_{t+k}(\delta_{t+k})$, and w_k^δ for $w_{t+k}(\delta_{t+k})$. We denote the support of the conditional distribution of $\boldsymbol{\delta}_t$ on \mathcal{I}_t by $\boldsymbol{\Delta}_t$.

Remark 1 (Other formulations). *The objective function in the MPC subproblem (3), and its analogues in §II-D and §II-E, can be augmented with a terminal cost to improve closed-loop performance or stability. To explicitly account for recourse, the control law can be parameterized as, e.g., affine disturbance feedback (see [9], [19]).*

D. Stochastic model predictive control

In the context of the stochastic optimal control problem (1), the stochastic MPC subproblem at stage t is

$$\begin{aligned} & \text{minimize} && \phi_t(\mathbf{u}_t) := \mathbf{E} \psi_t(\mathbf{u}_t, \boldsymbol{\delta}_t) \\ & \text{subject to} && \mathbf{u}_t \in \mathbf{U}_t. \end{aligned} \quad (4)$$

As discussed in §I, the stochastic MPC algorithm may not be implementable, because problem 4 may be intractable due either to incomplete distributional information or costly expectation integrals. Nevertheless, we state the algorithm here as a benchmark against which the SAMPC algorithm in §II-E can be compared.

Algorithm 1 (Stochastic MPC). Given K , set $t = 0$. While $t < T$,

- 1) Measure y_t and update state/parameter estimates.
- 2) Compute a solution $(u_{0|t}^*, \dots, u_{K-1|t}^*)$ to problem (4).
- 3) Implement $u_{0|t}^*$, increment t , and go to step 1.

E. Sample-average model predictive control

In cases where the stochastic MPC subproblem (4) is intractable, the SAA method minimizes the sample-average cost under the i.i.d. scenarios $\boldsymbol{\delta}_t^i := (x_{0|t}^i, \delta_{0|t}^i, \dots, \delta_{K-1|t}^i)$, $i = 1, \dots, N$, drawn from the conditional distribution of $\boldsymbol{\delta}_t$ on \mathcal{I}_t . In practice, these scenarios can be generated by resampling historical data or calling a pseudorandom number generator. At stage t , this gives rise to the subproblem

$$\begin{aligned} & \text{minimize} && \hat{\phi}_t^N(\mathbf{u}_t) := (1/N) \sum_{i=1}^N \psi_t(\mathbf{u}_t, \boldsymbol{\delta}_t^i) \\ & \text{subject to} && \mathbf{u}_t \in \mathbf{U}_t. \end{aligned} \quad (5)$$

The following proposition establishes the tractability of SAMPC for uncertain linear systems.

Proposition 1 (Feasibility and convexity). *Suppose Assumption 2 holds. Then the SAMPC subproblem (5) is feasible and almost surely convex.*

Proof. Feasibility follows from the nonemptiness of each U_t . Convexity of \mathbf{U}_t follows from that of each U_t , and convexity of $\hat{\phi}_t^N$ can be seen from the definition of ψ_t as follows. From Assumption 2, each g_{t+k} is almost surely convex in its first two arguments, which are affine functions of \mathbf{u}_t . Composition with an affine function preserves convexity, so $g_{t+k}(E_{k|t}\mathbf{u}_t, F_{k|t}(\boldsymbol{\delta}_t)\mathbf{u}_t + f_{k|t}(\boldsymbol{\delta}_t), G_{k|t}\boldsymbol{\delta}_t)$ is almost surely convex in \mathbf{u}_t for each $k \in \mathcal{K}$. Nonnegative weighted sums also preserve convexity, so each $\psi_t(\mathbf{u}_t, \boldsymbol{\delta}_t^i)$, and therefore $\hat{\phi}_t^N(\mathbf{u}_t)$, is almost surely convex in \mathbf{u}_t . ■

Remark 2. *The SAMPC subproblem (5) retains the problem class of the certainty-equivalent MPC subproblem (3). For example, if problem (3) is a linear program, then so is problem (5). Furthermore, the two problems have the same number of optimization variables and constraints.*

We now propose a SAMPC algorithm for approximately solving the stochastic optimal control problem (1).

Algorithm 2 (SAMPC). Given K and N , set $t = 0$. While $t < T$,

- 1) Measure y_t and update state/parameter estimates.
- 2) Sample the scenarios $\boldsymbol{\delta}_t^1, \dots, \boldsymbol{\delta}_t^N$ from the conditional distribution of $\boldsymbol{\delta}_t$ on \mathcal{I}_t .
- 3) Compute a solution $(\hat{u}_{0|t}^N, \dots, \hat{u}_{K-1|t}^N)$ to problem (5).
- 4) Implement $\hat{u}_{0|t}^N$, increment t , and go to step 1.

Remark 3. *Algorithm 2 requires a subroutine that can solve the feasible convex program (5). Under mild additional assumptions on the structure of $\hat{\phi}_t^N$ and \mathbf{U}_t , this can be accomplished in polynomial time by commercially-available solvers. Algorithm 2 can be modified to use a different sample size, MPC horizon, or optimization subroutine at each time step to exploit the time-varying problem structure.*

Analyzing the performance of Algorithm 2 on the stochastic optimal control problem (1) is difficult due to the heuristics involved (*i.e.*, separation of estimation and control and repeated optimization over a truncated, receding horizon). We can, however, make rigorous claims about the convergence of SAMPC to stochastic MPC. This is the focus of §III.

III. STATISTICAL PROPERTIES OF SAMPC ESTIMATORS

For fixed $\mathbf{u}_t \in \mathbf{U}_t$, the objective function $\hat{\phi}_t^N(\mathbf{u}_t)$ of problem (5) depends on the scenarios $\boldsymbol{\delta}_t^1, \dots, \boldsymbol{\delta}_t^N$, and is therefore random. Since $\boldsymbol{\delta}_t^1, \dots, \boldsymbol{\delta}_t^N$ are sampled from the distribution of $\boldsymbol{\delta}_t$, for any $\mathbf{u}_t \in \mathbf{U}_t$ we have

$$\mathbf{E} \hat{\phi}_t^N(\mathbf{u}_t) = \frac{1}{N} \sum_{i=1}^N \mathbf{E} \psi_t(\mathbf{u}_t, \boldsymbol{\delta}_t^i) = \mathbf{E} \psi_t(\mathbf{u}_t, \boldsymbol{\delta}_t) = \phi_t(\mathbf{u}_t).$$

In other words, $\hat{\phi}_t^N(\mathbf{u}_t)$ is an unbiased estimator of $\phi_t(\mathbf{u}_t)$. Under mild regularity conditions, the Strong Law of Large Numbers gives that for all $\mathbf{u}_t \in \mathbf{U}_t$, $\hat{\phi}_t^N(\mathbf{u}_t)$ is also a consistent estimator of $\phi_t(\mathbf{u}_t)$:

$$\hat{\phi}_t^N(\mathbf{u}_t) \rightarrow \phi_t(\mathbf{u}_t) \text{ almost surely as } N \rightarrow \infty. \quad (6)$$

These facts have been noted in [3] and [14].

To compare the SAMPC and stochastic MPC algorithms, however, it is necessary to study not only the objective function $\hat{\phi}_t^N$ of problem (5), but also the optimal value \hat{v}_t^N and set $\hat{\mathbf{U}}_t^N$ of minimizers. These are also random, and can be viewed as estimators of the optimal value v_t^* and set \mathbf{U}_t^* of minimizers of the stochastic MPC subproblem (4).

The following proposition, proved by Mak *et al.* in Theorems 1 and 2 of [16], shows that \hat{v}_t^N is a negatively biased estimator of v_t^* for any N , and that the bias tends monotonically to zero. This is somewhat surprising, given that $\hat{\phi}_t^N$ is an unbiased estimator of ϕ_t .

Proposition 2 (Bias, Mak [16]). *For any integer $N > 0$, $\mathbf{E} \hat{v}_t^N \leq \mathbf{E} \hat{v}_t^{N+1} \leq v_t^*$.*

The following result, proved by Shapiro in [18] using the theory of epi-convergence, establishes the consistency of the estimators generated by the SAMPC subproblem.

Proposition 3 (Subproblem consistency, Shapiro [18]). *Suppose Assumptions 1–2 and the following hold:*

- (i) *epi $\psi_t(\cdot, \delta_t)$ is closed for all $\delta_t \in \Delta_t$,*
- (ii) *epi ϕ_t is closed,*
- (iii) *\mathbf{U}_t^* is nonempty and bounded, and*
- (iv) *the pointwise Law of Large Numbers (6) holds.*

Then $\hat{v}_t^N \rightarrow v_t^$ and $\text{dev}(\hat{\mathbf{U}}_t^N, \mathbf{U}_t^*) \rightarrow 0$ almost surely as $N \rightarrow \infty$.*

If $\mathbf{E}[\psi_t(\mathbf{u}_t, \delta_t)^2]$ is finite, then by the Central Limit Theorem, for any fixed $\mathbf{u}_t \in \mathbf{U}_t$ the random variable $\sqrt{N}(\hat{\phi}_t^N(\mathbf{u}_t) - \phi_t(\mathbf{u}_t))$ converges in distribution to a zero-mean, Gaussian random variable with variance equal to $\text{Var}(\psi_t(\mathbf{u}_t, \delta_t))$. Under significantly stronger regularity conditions, similar results can be established for the SAMPC optimal value and minimizers (see §4.2 of [13]). The asymptotic distributions can be used to construct hypothesis tests about SAMPC solution quality, and give insight into the sample size required for a given accuracy.

Under additional assumptions, we can prove the consistency of the entire SAMPC control trajectory, viewed as an estimator of the stochastic MPC control trajectory.

Assumption 3 (Sub-Gaussianity). *There exists a $\sigma > 0$ such that for all $\mathbf{u}_t, \mathbf{v}_t \in \mathbf{U}_t$ and $\tau \in \mathbf{R}$, the random variable*

$$Y(\mathbf{u}_t, \mathbf{v}_t, \delta_t) := \psi_t(\mathbf{u}_t, \delta_t) - \psi_t(\mathbf{v}_t, \delta_t) - \mathbf{E}[\psi_t(\mathbf{u}_t, \delta_t) - \psi_t(\mathbf{v}_t, \delta_t)]$$

satisfies $\mathbf{E} \exp(\tau Y(\mathbf{u}_t, \mathbf{v}_t, \delta_t)) \leq \exp(\sigma^2 \tau^2 / 2)$. We say that $Y(\mathbf{u}_t, \mathbf{v}_t, \delta_t)$ is sub-Gaussian with parameter σ .

Loosely speaking, Assumption 3 guarantees that the problem variability is not too large. If each $Y(\mathbf{u}_t, \mathbf{v}_t, \delta_t)$ is

Gaussian, for example, then Assumption 3 holds with $\sigma = \sup \{\text{Var}(Y(\mathbf{u}_t, \mathbf{v}_t, \delta_t)) \mid \mathbf{u}_t, \mathbf{v}_t \in \mathbf{U}_t\}$. If $|Y(\mathbf{u}_t, \mathbf{v}_t, \delta_t)| \leq \rho$ for all $\mathbf{u}_t, \mathbf{v}_t \in \mathbf{U}_t$ and some $\rho \in (0, \infty)$, then Assumption 3 holds with $\sigma = \rho$. If no distributional information is available, σ can be estimated by sampling.

Definition 1 (Sharp). *A point $\mathbf{u}_t^* \in \mathbf{U}_t^*$ is a sharp minimizer of the stochastic MPC subproblem (4) if there exists a constant $c > 0$ such that for all $\mathbf{u}_t \in \mathbf{U}_t$,*

$$\phi_t(\mathbf{u}_t) \geq \phi_t(\mathbf{u}_t^*) + c \|\mathbf{u}_t - \mathbf{u}_t^*\|.$$

Theorem 4 (Trajectory consistency). *Suppose Assumptions 1–2 hold and for all $t \in \mathcal{T}$, the stochastic MPC subproblem (4) has a sharp minimizer \mathbf{u}_t^* . Then for some finite N , $\hat{\mathbf{U}}_t^N = \{\mathbf{u}_t^*\}$ for all $t \in \mathcal{T}$ with probability one. If Assumption 3 also holds, then there exist constants $a, b > 0$ such that if*

$$N \geq \frac{1}{a} \ln \left(\frac{b}{\beta} \right), \quad (7)$$

then $\mathbf{P}\{\hat{\mathbf{U}}_t^N = \{\mathbf{u}_t^\} \forall t \in \mathcal{T}\} \geq (1 - \beta)^T$.*

Proof. The first result follows directly from Theorem 5.23 of [18], which also guarantees that for all $t \in \mathcal{T}$ there exist $a_t, b_t > 0$ such that $\mathbf{P}\{\hat{\mathbf{U}}_t^N = \{\mathbf{u}_t^*\}\} \geq 1 - b_t \exp(-a_t N)$. We prove the second result by induction. At $t = 0$, there exist $a_0, b_0 > 0$ such that if $N \geq (1/a_0) \ln(b_0/\beta)$, then $\mathbf{P}\{\hat{\mathbf{U}}_0^N = \{\mathbf{u}_0^*\}\} \geq 1 - \beta$. Now suppose there exist $\bar{a}_t, \bar{b}_t > 0$ such that $N \geq (1/\bar{a}_t) \ln(\bar{b}_t/\beta)$ implies $\mathbf{P}\{\hat{\mathbf{U}}_s^N = \{\mathbf{u}_s^*\} \forall s \in \{0, \dots, t-1\}\} \geq (1 - \beta)^t$. Then with probability at least $(1 - \beta)^t$, the stochastic MPC and SAMPC subproblems at stage t have the same input data, so there exist $a_t, b_t > 0$ such that $N \geq (1/\min\{\bar{a}_t, a_t\}) \ln(\max\{\bar{b}_t, b_t\}/\beta)$ implies $\mathbf{P}\{\hat{\mathbf{U}}_s^N = \{\mathbf{u}_s^*\} \forall s \in \{0, \dots, t\}\} \geq (1 - \beta)^{t+1}$. The result follows, with $a = \min_{t \in \mathcal{T}} a_t$ and $b = \max_{t \in \mathcal{T}} b_t$. ■

Theorem 4 says that if each stochastic MPC subproblem has a sharp minimizer, then for some finite N , SAMPC exactly reproduces the stochastic MPC control trajectory. Furthermore, the probability that the stochastic MPC and SAMPC trajectories differ decays exponentially quickly to zero as N increases. A unique minimizer is guaranteed to be sharp if \mathbf{U}_t is a polyhedron, $\psi_t(\cdot, \delta_t)$ is piecewise affine for every δ_t , and Δ_t is a finite set. [18]

IV. NUMERICAL EXAMPLE

In this section, we demonstrate SAMPC through the example of efficiently heating a building. The problem is to balance the competing objectives of minimizing energy cost and maintaining the indoor air temperature above an uncertain, time-varying setpoint. In this example, as in many building climate control problems, occasional temperature deviations below the setpoint are not catastrophic. They do tend to cause occupant discomfort, however, which comes at some cost to the building operator. Therefore, it is natural to penalize temperature deviations in the objective function, rather than imposing robust or chance constraints.

The building considered is a simple one-room residence. We identify a SISO ($n = m = q = 1$) linear model from data




Fig. 1. Histograms of the waking, departure, return, and sleeping times in the historical temperature setpoint data. In each simulated day, the setpoint is high when the occupant is awake and at home, and low otherwise.

generated by the MATLAB[®] `bdg` toolbox, which acts as a nonlinear, time-varying truth-model (for details, see [20], [21]). The state x_t (°C) is the indoor air temperature. The control u_t (kW) is the heat flow from an electric resistance heater with capacity $\bar{u} = 10$ kW, so $U_t = [0, \bar{u}]$ for all t . We use a time step of $\Delta t = 15$ minutes, control horizon of one day ($T = 96$), and MPC horizon of four hours ($K = 16$).

The first two of the $p = 7$ uncertain parameters in δ_t define the system matrices $A_t(\delta_t) = (\delta_t)_1$ and $B_t(\delta_t) = (\delta_t)_2$. The other system matrices $C_t(\delta_t) = 1$ and $D_t(\delta_t) = 0$ are deterministic. The third through sixth parameters define the disturbance $w_t(\delta_t) = (\delta_t)_3 T_t^\infty + (\delta_t)_4 I_t^h + (\delta_t)_5$ and the measurement noise $v_t(\delta_t) = (\delta_t)_6$. For simplicity, we assume perfect knowledge of the outdoor air temperature T_t^∞ (°C) and the total solar irradiance on a horizontal surface, I_t^h (W/m²). The final parameter $(\delta_t)_7$ is the temperature setpoint.

At each stage, the controller incurs the cost

$$g_t(u_t, x_t, \delta_t) = \pi_e \Delta t u_t + \frac{\gamma}{T} ((\delta_t)_7 - x_t)_+,$$

where $\pi_e = 0.1$ \$/kWh is the electricity price. The second term in the stage cost is a special case of the penalty (2), with $X_t(\delta_t) = \{x \in \mathbf{R} \mid x \geq (\delta_t)_7\}$ and $\eta(\cdot) = (\cdot)_+ := \max\{0, \cdot\}$. The parameter γ (\$/°C) governs the trade-off between energy cost and time-averaged deviation below the setpoint. In all simulations, we set $\gamma = 750$ \$/°C, so that a time-averaged deviation of 140 Celsius-hours per year (twice the acceptable value under European building standards) incurs the same cost as running the heater at half capacity.

The system model is identified using linear regression, so the parameters $(\delta_t)_1, \dots, (\delta_t)_5$ are modeled as jointly Gaussian and white. The same applies to the measurement noise $(\delta_t)_6$. We do not, however, assume that the distribution of the sequence of setpoints $(\delta_t)_7$ is known. Rather, we suppose the controller has access to a large set of historical setpoint trajectories. For simulation purposes, we randomly generate these trajectories as follows. The setpoint is 16 °C when the resident is asleep or out, and 21 °C when they are awake and at home. Each day, the resident wakes up, leaves home, returns home, then goes to sleep. The waking, departure, return, and sleeping times vary randomly from day to day, as shown in Figure 1.

Implementing SAMPC at stage t requires randomly sam-




Fig. 2. Temperature setpoints and measurements in one Monte Carlo run. With $N = 20$, SAMPC uses 13% more energy than certainty-equivalent MPC, but is more robust to uncertainty in the timing of setpoint changes.

pling from the (unknown) conditional distribution of the upcoming setpoints $(\delta_t)_7, \dots, (\delta_{t+K-1})_7$ on \mathcal{I}_t . We accomplish this approximately by resampling from those historical setpoint trajectories that match the realizations of $(\delta_0)_7, \dots, (\delta_{t-1})_7$ within a specified tolerance. We use a similar approach in the certainty-equivalent MPC setpoint prediction.

Figure 2 shows the air temperature measurements under certainty-equivalent MPC and SAMPC with $N = 20$ in a representative Monte Carlo run. Certainty-equivalent MPC is aggressive, maintaining the air temperature near the setpoint. This leads to occasional mild deviations below the setpoint due to model error and disturbances. Additionally, in this Monte Carlo run the resident both wakes up and returns home earlier than predicted, which causes the certainty-equivalent MPC temperature to lag behind the morning and evening setpoint changes.

SAMPC modifies the certainty-equivalent MPC behavior in two ways. First, it maintains a buffer of about 0.5 °C between the air temperature and the setpoint. Second, it begins the morning and evening warm-up phases about an hour earlier than certainty-equivalent MPC, so that the building is already warm when the occupant wakes up and returns home. In this Monte Carlo run, the energy costs under certainty-equivalent MPC and SAMPC were \$7.50 and \$8.53, and the time-averaged setpoint deviations were 180 and 0 Celsius-hours per year.

A key observation is that as N increases, the expected cost of the SAMPC policy decreases exponentially, while its computational complexity increases only linearly. As illustrated in Figure 3, a sample size of $N = 20$ appears to strike a good balance between controller performance and computational complexity. The cost decrease with increasing N comes mainly from improved robustness to sudden temperature setpoint changes. This robustness comes at the cost of increased energy use, as Figure 4 shows.

V. CONCLUSION

In this paper, we proposed a randomized MPC algorithm that minimizes the sample-average cost as a proxy for the expected cost when distributional information is lacking or




Fig. 3. Average cost and computer time of the SAMPC policy over 100 Monte Carlo simulations of a 24-hour control horizon. Optimization was done in the Gurobi solver on a 2 GHz Intel Core 2 Duo processor.

expectation integrals are intractable. The algorithm naturally extends certainty-equivalent MPC, so is similarly simple to implement. The SAMPC subproblems are always feasible and retain the problem class of the corresponding certainty-equivalent MPC subproblems, as well as the number of optimization variables and constraints. Under mild regularity conditions, SAMPC converges to stochastic MPC as the sample size tends to infinity. In some cases, this convergence is exponentially fast.

An important direction for future research is the construction of sample-size bounds to guarantee that, at a given confidence level, a SAMPC minimizer is nearly optimal for the stochastic MPC subproblem. Some analytical results in this vein can be found in §5.3 of [18], but the bounds are quite conservative and not particularly constructive.

The SAMPC algorithm proposed in this paper could be modified in a number of ways. Variance reduction techniques could be used to improve SAA convergence rates. Other sample-based techniques for stochastic programming, such as stochastic approximation or retrospective approximation (see §5.2 of [13]), could be attractive alternatives to SAA.

ACKNOWLEDGMENTS

We are grateful to the Hydro Research Foundation and the U.S. Department of Energy’s Office of Renewable Energy and Energy Efficiency for funding. Thanks as well to Professors Eilyan Bitar and Shane Henderson of Cornell University for their valuable suggestions.

REFERENCES

- [1] J. Lee, “Model Predictive Control: Review of the Three Decades of Development,” *International Journal of Control, Automation and Systems*, vol. 9, no. 3, pp. 414–424, 2011.
- [2] E. Cinquemani, M. Agarwal, D. Chatterjee, and J. Lygeros, “Convexity and convex approximations of discrete-time stochastic control problems with constraints,” *Automatica*, vol. 47, no. 9, pp. 2082–2087, 2011.
- [3] L. Blackmore, M. Ono, A. Bektassov, and B. Williams, “A probabilistic particle-control approximation of chance-constrained stochastic predictive control,” *Robotics, IEEE Transactions on*, vol. 26, no. 3, pp. 502–517, 2010.
- [4] G. Schildbach, G. Calafiore, L. Fagiano, and M. Morari, “Randomized model predictive control for stochastic linear systems,” in *American Control Conference*, 2012, pp. 417–422.

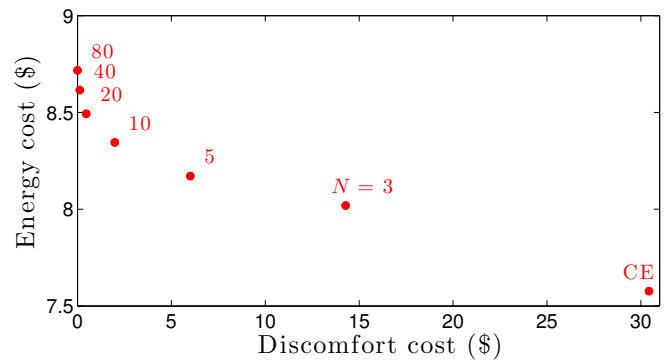


Fig. 4. Average costs of energy and discomfort over 100 Monte Carlo simulations. For all simulated $N \geq 80$, the SAMPC policy achieves zero time-averaged constraint violation and an energy cost of about \$8.75.

- [5] G. Calafiore and L. Fagiano, “Model predictive control of stochastic LPV systems via random convex programs,” in *Conference on Decision and Control*. IEEE, 2012, pp. 3233–3238.
- [6] M. Prandini, S. Garatti, and J. Lygeros, “A randomized approach to stochastic model predictive control,” in *Conference on Decision and Control*, 2012, pp. 7315–7320.
- [7] J. Matusko and F. Borrelli, “Scenario-based approach to stochastic linear predictive control,” in *Conference on Decision and Control*. IEEE, 2012, pp. 5149–5199.
- [8] G. Calafiore and L. Fagiano, “Stochastic model predictive control of LPV systems via scenario optimization,” *Automatica*, vol. 49, no. 6, pp. 1861–1866, 2013.
- [9] G. Schildbach, L. Fagiano, C. Frei, and M. Morari, “The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations,” *Automatica*, vol. 50, no. 12, pp. 3009–3018, 2014.
- [10] G. Calafiore and M. Campi, “Uncertain convex programs: Randomized solutions and confidence levels,” *Mathematical Programming*, vol. 102, no. 1, pp. 25–46, 2005.
- [11] M. Campi and S. Garatti, “The exact feasibility of randomized solutions of uncertain convex programs,” *SIAM Journal on Optimization*, vol. 19, no. 3, pp. 1211–1230, 2008.
- [12] G. Calafiore, “Random convex programs,” *SIAM Journal on Optimization*, vol. 20, no. 6, pp. 3427–3464, 2010.
- [13] S. Kim, R. Pasupathy, and S. G. Henderson, “A guide to sample average approximation,” in *Handbook of Simulation Optimization*. Springer New York, 2015, pp. 207–243.
- [14] I. Batina, A. Stoorvogel, and S. Weiland, “Stochastic disturbance rejection in model predictive control by randomized algorithms,” in *American Control Conference*, 2001, pp. 732–737.
- [15] K. Nolde, M. Uhr, and M. Morari, “Medium term scheduling of a hydro-thermal system using stochastic model predictive control,” *Automatica*, vol. 44, no. 6, pp. 1585–1594, 2008.
- [16] W. Mak, D. Morton, and R. Wood, “Monte Carlo bounding techniques for determining solution quality in stochastic programs,” *Operations Research Letters*, vol. 24, no. 1, pp. 47–56, 1999.
- [17] A. Kleywegt, A. Shapiro, and T. Homem-de Mello, “The sample average approximation method for stochastic discrete optimization,” *SIAM Journal on Optimization*, vol. 12, no. 2, pp. 479–502, 2001.
- [18] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory*, ser. MPS-SIAM Series on Optimization. Philadelphia, PA: MPS-SIAM, 2009.
- [19] D. de la Peña, A. Bemporad, and T. Alamo, “Stochastic programming applied to model predictive control,” in *Conference on Decision and Control and European Control Conference*, 2005, pp. 1361–1366.
- [20] K. J. Kircher and K. M. Zhang, “Testing building controls with the BLDG toolbox,” in *American Control Conference (ACC)*, 2016.
- [21] K. J. Kircher, W. Schaefer, and K. M. Zhang, “A simple first-principles building energy model,” *under review*, September 2016.