Sample-average approximation in stochastic model predictive control

IEEE Conference on Decision and Control

Kevin J. Kircher and K. Max Zhang, Cornell University

December 14, 2016

Background

Main convergence result

Example: efficiently heating an apartment

Background

Main convergence result

Example: efficiently heating an apartment

A stochastic optimal control problem

$$
\begin{array}{ll}
\underset{\mu_0,\ldots,\mu_{T-1}}{\text{minimize}} & \mathbf{E} \sum_{t \in \mathcal{T}} g_t(u_t, x_t, \delta_t) \\
\text{subject to} & x_{t+1} = A_t(\delta_t) x_t + B_t(\delta_t) u_t + w_t(\delta_t) \ \forall \ t \in \mathcal{T} \\
& u_t = \mu_t(x_0, \ldots, x_t) \in \mathcal{U}_t \ \forall \ t \in \mathcal{T}
\end{array} \tag{1}
$$

- $t \in \mathcal{T} := \{0, \ldots, T-1\}$ indexes discrete time
- states x_t are perfectly observed
- control constraint sets $\mathcal{U}_t$ are closed, nonempty, convex
- stage costs $(u, x) \mapsto g_t(u, x, \delta_t)$ are convex for all δ_t
- random vectors δ_t contain all uncertain influences on system
- samples can be generated from the (arbitrary, possibly unknown) distributions of the δ_t

What about state constraints?

- conspicuously absent: $\mathbf{P}\{x_t \in \mathcal{X}_t(\delta_t)\} \geq 1 - \alpha_t$
- but penalties on deviations of x_t from $\mathcal{X}_t(\delta_t)$ are allowed
- justification:
  - ◇ in practice, constraints are often softened to avoid infeasibility
  - ◇ in some applications, $x_t \notin \mathcal{X}_t(\delta_t)$ is not catastrophic
  - ◇ but large deviations may be significantly worse than small ones (chance constraints do not capture this)
  - ◇ scenario approximations of chance constraints are well understood (Calafiore '12, '13; Prandini '12; Matusko '12; Schildbach '12, '14)

Problem (1) is intractable in general

- optimization variables μ_t are (infinite dimensional) functions
- control horizon may be very long
- expectation integrals may be ill-defined or hard to compute/approximate

Stochastic MPC

choose MPC horizon K (define $\mathcal{K} := \{0, \ldots, K-1\}$)

1. measure x_t
2. solve

$$
\begin{array}{ll}
\underset{u_{0|t},\ldots,u_{K-1|t}}{\text{minimize}} & \mathbf{E} \sum_{k \in \mathcal{K}} g_{t+k}(u_{k|t}, x_{k|t}, \delta_{k|t}) \\
\text{subject to} & x_{k+1|t} = A_{t+k}(\delta_{k|t}) x_{k|t} + B_{t+k}(\delta_{k|t}) u_{k|t} + w_{t+k}(\delta_{k|t}) \\
& \qquad \forall\ k \in \mathcal{K} \\
& u_{k|t} \in \mathcal{U}_{t+k}\ \forall\ k \in \mathcal{K}
\end{array}
$$

   for $(u^\star_{0|t}, \ldots, u^\star_{K-1|t})$

3. set $u_t = u^\star_{0|t}$, increment t, repeat

(still intractable)

Sample-average MPC

choose MPC horizon K, sample size N (define $\mathcal{N} := \{1, \ldots, N\}$)

1. measure x_t
2. sample iid scenarios $(\delta^1_{0|t}, \ldots, \delta^1_{K-1|t}), \ldots, (\delta^N_{0|t}, \ldots, \delta^N_{K-1|t})$
3. solve

$$
\begin{array}{ll}
\underset{\substack{u_{0|t},\ldots,u_{K-1|t} \\ x^1_{1|t},\ldots,x^1_{K|t} \\ \cdots \\ x^N_{1|t},\ldots,x^N_{K|t}}}{\text{minimize}} & (1/N) \sum_{i\in\mathcal{N}} \sum_{k\in\mathcal{K}} g_{t+k}(u_{k|t}, x^i_{k|t}, \delta^i_{k|t}) \\
\text{subject to} & x^i_{k+1|t} = A_{t+k}(\delta^i_{k|t}) x^i_{k|t} + B_{t+k}(\delta^i_{k|t}) u_{k|t} + w_{t+k}(\delta^i_{k|t}) \\
& \qquad \forall\ k \in \mathcal{K},\ i \in \mathcal{N} \\
& u_{k|t} \in \mathcal{U}_{t+k}\ \forall\ k \in \mathcal{K}
\end{array}
$$

   for $(u^\star_{0|t}, \ldots, u^\star_{K-1|t}, x^{1\star}_{1|t}, \ldots, x^{1\star}_{K|t}, \ldots, x^{N\star}_{1|t}, \ldots, x^{N\star}_{K|t})$
4. set $u_t = u^\star_{0|t}$, increment t, repeat

Related work

- SAMPC has been proposed by Batina '01, Nolde '08, Blackmore '10, Schildbach '14, Verrilli '16 (MoA22.4)
- justification typically involves (pointwise) convergence of SAMPC objective function to stochastic MPC objective function as $N \to \infty$
- but do SAMPC costs and controls necessarily converge to those of stochastic MPC?
- if so, how quickly do they converge?

Our contributions:

- translation of statistical results from the stochastic programming literature into an MPC context (see paper)
- an exponential convergence theorem for the full control trajectory

Background

Main convergence result

Example: efficiently heating an apartment

Subproblem reformulations

- define $\mathbf{u}_t := (u_{0|t}, \dots, u_{K-1|t})$, $\boldsymbol{\delta}_t := (\delta_{0|t}, \dots, \delta_{K-1|t})$, $\boldsymbol{\mathcal{U}}_t := \mathcal{U}_t \times \cdots \times \mathcal{U}_{t+K-1}$
- iteratively apply dynamics to eliminate state variables
- appropriately define 'lifted' objective function ψ_t

then at stage t, the stochastic MPC subproblem can be written as

$$\begin{array}{ll} \underset{\mathbf{u}_t}{\text{minimize}} & \phi_t(\mathbf{u}_t) := \mathbf{E}\, \psi_t(\mathbf{u}_t, \boldsymbol{\delta}_t) \\ \text{subject to} & \mathbf{u}_t \in \boldsymbol{\mathcal{U}}_t \end{array}$$

and the SAMPC subproblem as

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

Theorem (SAMPC trajectory convergence)

suppose

1. ψ_t is 'not too variable' (in a certain precise sense)
2. each $\boldsymbol{\mathcal{U}}_t$ is polyhedral, each $\psi_t(\cdot, \boldsymbol{\delta}_t)$ is piecewise affine for every $\boldsymbol{\delta}_t$, each $\boldsymbol{\delta}_t$ has finite support
3. each stochastic MPC subproblem has a unique minimizer

then

- the probability that the SAMPC and stochastic MPC control trajectories differ decays exponentially quickly as N increases
- that probability reaches zero for some finite N

(proof relies heavily on a result by Shapiro '09)

Background

Main convergence result

Example: efficiently heating an apartment

Problem statement

minimize

- (linear) energy cost + (piecewise affine) discomfort cost

subject to

- SISO temperature dynamics (identified from nonlinear model)
- heater nonnegativity and capacity constraints

under uncertainty from

- occupant-specified temperature reference
- multiplicative and additive model error
- weather

(convergence theorem conditions are almost, but not quite, met)

Occupant behavior

- occupant wants
  - at least 21 °C when home and awake
  - at least 16 °C when away or asleep
- wake, depart, return and sleep times are random

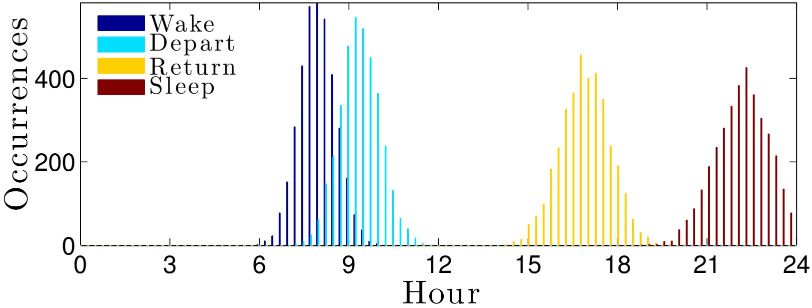

Indoor temperature in a typical Monte Carlo run

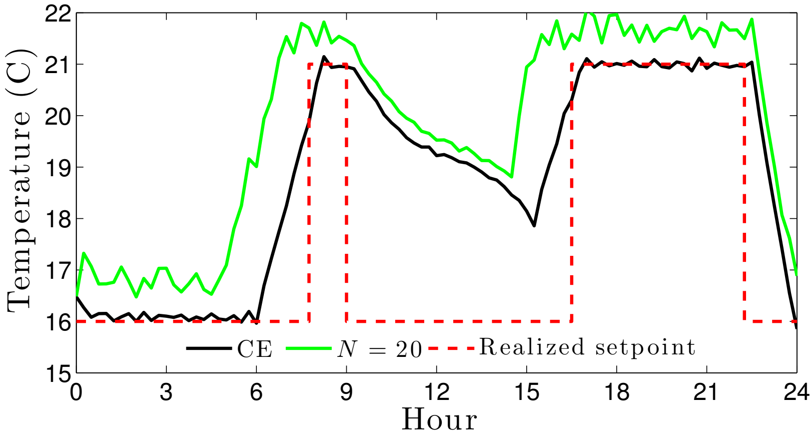

- CE means certainty-equivalent MPC

Increasing N improves comfort, uses more energy

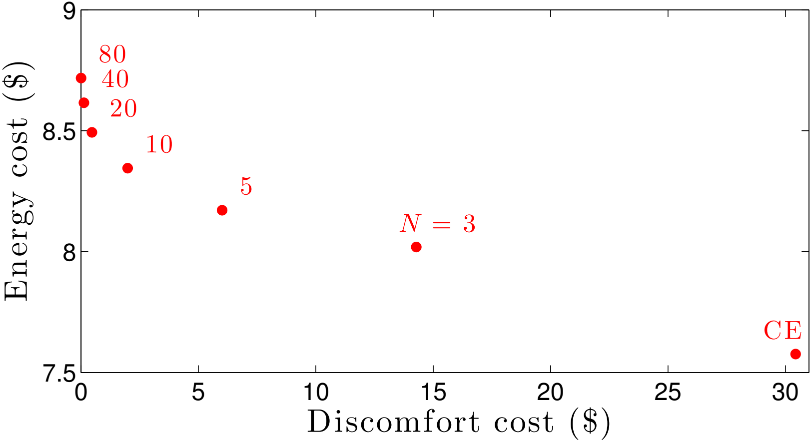

- costs are averaged over 100 Monte Carlo runs
- zero discomfort cost for all $N \geq 80$

Performance and computer time

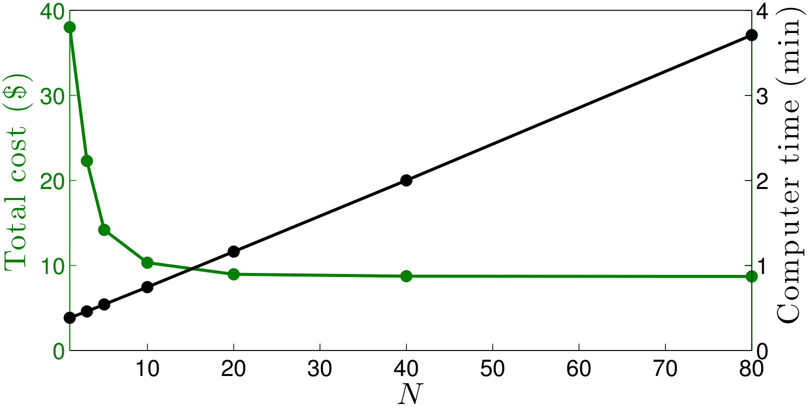

- cost, computer time are averaged over 100 Monte Carlo runs
- 15 minute time steps, 4 hour MPC horizon ($K = 16$)
- optimization in Gurobi on a 2 GHz Intel Core 2 Duo processor