SZTAKI

A **Sampling-and-Discarding** Approach to Stochastic Model Predictive Control for **Renewable Energy Systems**

Balázs Cs. Csáji, Krisztián B. Kis & András Kovács

SZTAKI: Institute for Computer Science and Control, Budapest, Hungary

21st IFAC World Congress, July 11-17, 2020

Overview

- We apply the scenario approach to stochastic model predictive control for a renewable energy system (public lighting microgrid).
- The scenario approach provides an elegant compromise between robust and chance-constrained (convex) optimization paradigms.
- First, the controllable and the (quasi-periodic) uncontrollable parts are decomposed; the latter is modeled by a Box-Jenkins system.
- For the controllable part, a linear state space model is applied with affine controllers for which various parametrizations are considered.
- Several numerical experiments are presented on the microgid.
- The effects of controller parametrizations, orders, reoptimization frequencies and discarding unfavorable trajectories are studied.
- The results indicate that even a low order, time-independent controller with a slow reoptimization frequency can be efficient.

Decomposable Markov Model

- We consider the following (possibly nonlinear) Markovian model:

$$x_t = f(x_{t-1}, u_t, \varepsilon_t),$$

  where $x_t \in \mathbb{R}^{d_x}$ is the state, $u_t \in \mathbb{R}^{d_u}$ is the control input, and $\varepsilon_t \in \mathbb{R}_t^{d_\varepsilon}$ is the driving noise term, at (discrete) time $t = 1, 2, \ldots$

- In some situations $\{x_t\}$ is only partially controllable and it can be decomposed into two parts (possibly after a state transformation):

$$x'_t = f'(x_{t-1}, u_t, \varepsilon_t), \qquad \text{and} \qquad x''_t = f''(x''_{t-1}, \varepsilon_t),$$

  where $\{x''_t\}$ is unaffected by the chosen inputs, f' and f'' are the controlled and uncontrolled parts of the dynamics, respectively.

- We will apply time-varying, state-feedback controllers,

$$u_t = \pi_t(x_{t-1}),$$

  where function π_t is called the (Markov) control policy at time t.

Linear Dynamics and Feedback Controllers

- An arch-typical choice for f is to use linear dynamics, i.e.,

$$x_t = f(x_{t-1}, u_t, \varepsilon_t) = A x_{t-1} + B u_t + \varepsilon_t,$$

where A and B are (constant) matrices of appropriate size.

- A common approach is to use a linear controller, that is

$$u_t = \pi_t(x_{t-1}) = K_t x_{t-1},$$

which however, leads to nonconvex MPC optimization problems.

- On the other hand, the policy can be reparametrized as

$$u_t = \pi_t(x_{t-1}) = \varphi_t + \sum_{i=1}^{t-1} \Phi_{t,i} \varepsilon_i,$$

to ensure that we have a convex MPC optimization problem.

- Note that, of course, $\varepsilon_t = x_t - A x_{t-1} - B u_t$, thus, available.

Controller Parametrizations

– A (time-varying) full affine controller is parametrized as

$$u_t = \pi_t(x_{t-1}) = \varphi_t + \sum_{i=1}^{t-1} \Phi_{t,i}\varepsilon_i.$$

– A time-varying affine controller with past order p is defined as

$$u_t = \pi_t(x_{t-1}) = \varphi_t + \sum_{\substack{i=1\vee\\(t-1-p)}}^{t-1} \Phi_{t,i}\varepsilon_i,$$

where $a \vee b = \max(a, b)$, i.e., max p past values are considered.

– Finally, a time-independent affine controller with past order p is

$$u_t = \pi(x_{t-1}) = \varphi + \sum_{\substack{i=1\vee\\(t-1-p)}}^{t-1} \Phi_i\varepsilon_i,$$

Model Predictive Control

- The scheme for an optimization step of MPC for horizon n is

$$\begin{aligned}\underset{\pi\in\Pi}{\text{minimize}} \quad & J_n^{\pi}(x_0) = \sum_{k=0}^{n-1} \ell_k(x_k, u_{k+1}) \\ \text{subject to} \quad & x_0 = x_{t_0}^{*} \\ & u_k = \pi_k(x_{k-1}) \\ & x_k = f(x_{k-1}, u_k) \\ & u \in \mathcal{U},\ x \in \mathcal{X} \\ & k = 1, \dots, n\end{aligned}$$

  where $x \doteq (x_0, \dots, x_n)^{\mathrm{T}}$, $u \doteq (u_0, \dots, u_n)^{\mathrm{T}}$, $\pi \doteq (\pi_1, \dots, \pi_n)$ are sequences of states, inputs and policies, respectively.
- $\mathcal{X}$, $\mathcal{U}$ are constraint sets for the allowed (sequences of) states and inputs; and $\{\ell_k\}$ are (given) $\mathbb{R}$-valued immediate-cost functions.

Chance-Constrained Model Predictive Control

– A standard stochastic MPC formulation with chance constraints is

$$
\begin{aligned}
\underset{\pi \in \Pi}{\text{minimize}} \quad & \mathbb{E}\big[\, J_n^{\pi}(x_0) \,\big] = \mathbb{E}\left[\sum_{k=0}^{n-1} \ell_k(x_k, u_{k+1}) \right] \\
\text{subject to} \quad & x_0 = x_{t_0}^* \\
& u_k = \pi_k(x_{k-1}) \\
& x_k = f(x_{k-1}, u_k, \varepsilon_k) \\
& \mathbb{P}_{\varepsilon}\{u \in \mathcal{U},\ x \in \mathcal{X}\} \geq 1 - \delta \\
& k = 1, \dots, n
\end{aligned}
$$

where $\varepsilon \doteq (\varepsilon_1, \dots, \varepsilon_n)$ is a random sequence of uncertainties, and (constant) δ is the allowed probability of constraint violation.

– In order to make this problem tractable, J, f, $\mathcal{X}$ and $\mathcal{U}$ are typically chosen to be convex; and $\mathbb{P}_{\varepsilon}$ is often assumed to be known.

Robust Model Predictive Control

- Let us introduce uncertainty-dependent constraint sets:

$$\mathcal{Z}(\varepsilon) \doteq \{ (h;\theta) \in \mathbb{R}^{m+1} : J_n^{\pi,\varepsilon}(x_0^\varepsilon) \leq h,\; x_0^\varepsilon = x_{t_0}^*, \\ u_k^\varepsilon = \pi_k(\, x_{k-1}^\varepsilon \mid \theta\,),\; x_k^\varepsilon = f(\, x_{k-1}^\varepsilon, u_k^\varepsilon, \varepsilon_k\,), \\ x^\varepsilon \in \mathcal{X},\; u^\varepsilon \in \mathcal{U},\; k = 1, \dots, n \},$$

  where $\varepsilon \in \mathbb{R}^n$ is a noise sequence, and $\pi_k(u) \doteq \pi_k(u \mid \theta)$ is a notation to emphasize that policies $\{\pi_k\}$ are parametrized by θ.

- Then, the optimization step of robust MPC can be written as:

$$\begin{aligned} \underset{h\in\mathbb{R},\,\theta\in\mathbb{R}^m}{\text{minimize}} \quad & h \\ \text{subject to} \quad & (h;\theta) \in \mathcal{Z}(\varepsilon), \quad \text{for all } \varepsilon \in E, \end{aligned}$$

  where E is the (not necessarily convex) set of all uncertainties.

- If E is infinite, this problem can only be solved in special cases.

Scenario-Based Model Predictive Control

- The scenario MPC provides a trade-off between chance-constrained and robust MPC, if we can sample (e.g., simulate) the disturbances.
- Let $\varepsilon^{(1)}, \dots, \varepsilon^{(N)}$ be N i.i.d. "scenarios"; then the (random) scenario problem is an approximation of the worst-case one:

$$\begin{aligned} \underset{h \in \mathbb{R},\, \theta \in \mathbb{R}^m}{\text{minimize}} \quad & h \\ \text{subject to} \quad & (h; \theta) \in \mathcal{Z}(\varepsilon^{(i)}), \quad \text{for } i = 1, \dots, n, \end{aligned}$$

  which is a finite convex problem, assuming $\{\mathcal{Z}(\varepsilon^{(i)})\}$ are convex.
- The (random) optimal solution of this problem is $z_N^\star \doteq (h^\star; \theta^\star)$.
- The violation probability of a fixed $z \in \mathbb{R}^{m+1}$ is defined by

$$V(z) \doteq \mathbb{P}\{\, \varepsilon \in E : z \notin \mathcal{Z}(\varepsilon) \,\}.$$

- We want to estimate $V(z_N^\star)$. Note that it is a random variable.

Constraint Violation Bounds

- The following probabilistic bound can be proved for $V(z_N^\star)$:

$$\mathbb{P}\{\, V(z_N^\star) > \delta \,\} \,\leq\, \beta(\delta, d, N),$$

for all $\delta \in (0,1)$, assuming that for all possible $\{\varepsilon^{(i)}\}$, the scenario problem is feasible and has a unique solution; where $\beta(\delta, d, N)$ is

$$\beta(\delta, d, N) \,\doteq\, \sum_{i=0}^{d-1} \binom{N}{i} \delta^i (1-\delta)^{N-i},$$

with $d = \dim(z_N^\star)$, i.e., the number of decision variables.

- Given δ and β, sufficient number of scenarios can be ensured if

$$N \,\geq\, {}^1\!/\!_\delta \left(d - 1 + \log({}^1\!/\!_\beta) + \sqrt{2(d-1)\log({}^1\!/\!_\beta)} \right),$$

which sample size guarantees that $\mathbb{P}\{\, V(z_N^\star) \leq \delta \,\} \geq 1 - \beta$.

Sampling and Discarding

- We may discard k (unfavorable) scenarios and still have that if

$$\binom{k+d-1}{k} \sum_{i=0}^{k+d-1} \binom{N}{i} \delta^i (1-\delta)^{N-i} \leq \beta,$$

  then we can guarantee $V(z^\star_{N-k}) \leq \delta$ with confidence $1-\beta$.
- This result is independent of the algorithm for selecting the scenarios to be removed (but k should be chosen a priori).
- We apply the scenario approach to the optimization step of MPC which corresponds to a value-at-risk formulation of the problem.
- To select which scenarios to remove, we simulate the policy of the previous step, and remove the scenarios with the worst outcome.
- If in each step $V(z^\star_N)$ is below δ, the expected time-average of closed-loop constraint violations also remains upper bounded by δ.

Public Lighting Microgrid

- This sampling and discarding type MPC was applied to control the energy flow of the E+grid experimental public lighting microgrid.
- The system comprises 191 intelligent LED luminaries that adjust their lighting levels according to the actual traffic conditions.
- The energy is generated by roof-mounted PV panels with a total active surface area of 152.5 m^2 and peak power of 21 kWp.
- There is a battery storage with a capacity of 18.5 kWh.
- It has a bidirectional grid connection (it can buy and sell energy).
- The system must be robust against potential power outages: it should guarantee three hour island mode operation.
- The system is located in the campus of a research institute in Budapest, Hungary. Our experiments are based on real data; but the effectiveness of the controllers was evaluated by simulation.

Modeling the Energy Balance

- To generate forecasts ("scenarios"), we model the energy balance, i.e., the difference of the energy production and the consumption.
- Let $\{\varepsilon_t\}$ be the quasi-periodic energy balance (time step: 1 hour).
- Its historical averages for each hour of the day is denoted by $\{v_t\}$.
- We can model the energy balance by NARX models, that is

$$\varepsilon_t \;=\; g(\varepsilon_{t-1}, \dots, \varepsilon_{t-p}, v_t) \,+\, n_t,$$

  where n_t is the process noise at time t, and p is the order.
- Function g is realized by an SVR or an MLP (nonlinear) model.
- Another possible model is BJ (Box-Jenkins) that takes the form:

$$\varepsilon_t \;=\; F^{-1}(q)B(q)\, v_t + D^{-1}(q)C(q)\, n_t,$$

  where B, C, D, F are finite polynomials in q^{-1} (backward shift).

Experiment: Time-Series Models

- The effectiveness of these three types of models was evaluated by ten-fold cross validation on a dataset of 500 energy balance data.
- The average RMSEs with their standard deviations were compared.
- The results show that BJ models were the best for this problem.

Model Type		Performance by Model Order					
		1	2	3	4	5	6
MLP	RMSE	0.534	0.531	0.520	0.540	0.494	0.507
	STD	0.127	0.132	0.167	0.150	0.128	0.160
SVR	RMSE	0.563	0.539	0.519	0.527	0.513	0.527
	STD	0.092	0.135	0.151	0.162	0.135	0.148
BJ	RMSE	0.480	0.481	0.482	0.471	0.461	0.458
	STD	0.126	0.133	0.133	0.154	0.155	0.154

Experiment: Generating Trajectories

- The noise is estimated by its EDF and resampled by bootstrap.
- The figure shows generating trajectories by BJ (bootstraped noise).

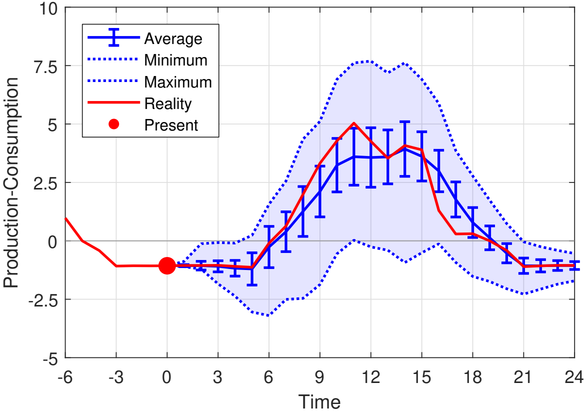

Controlling the Energy Flow

- The aim is to trade with the electricity in a cost-effective way, always guaranteeing $r = 3$ hours of island mode operation.
- The cost-to-go function of a control policy π is defined as

$$J_n^{\pi}(x_0) = \sum_{t=1}^{n} \alpha^{t-1} \big(c_t^+ u_t^+ - c_t^- u_t^- \big),$$

where $\alpha \in (0,1)$ is a discount factor; c_t^+ and c_t^- are the costs of buying and selling; u_t^+ and u_t^- are the amount of bought and sold electricity (the positive and the negative components of u_t).

- In the experiments we used: $\alpha = 0.95$, $c_t^+ = 1.0$ and $c_t^- = 0.95$.
- We consider three types of controllers: the (time-varying) full affine, the (time-varying) fixed-past and the time-independent.
- Note that the controllers are affine functions of the energy balance.

Uncertainty-Dependent Constraint Sets

- The uncertainty-dependent constraint sets are

$$
\begin{aligned}
\mathcal{Z}(\varepsilon) \,=\, \{\, (h;\theta) \in \mathbb{R}^{m+1} : J_n^{\pi,\varepsilon}(x_0) \,\leq\, h,\; x_0^{\varepsilon} = x_{t_0}^{*},\\
u_t^{\varepsilon} = \pi_t(x_{t-1}^{\varepsilon} \mid \theta),\; x_t^{\varepsilon} = x_{t-1}^{\varepsilon} + u_t^{\varepsilon} + \varepsilon_t,\\
u_t^{\varepsilon} = u_t^{+} - u_t^{-},\; u_t^{+} \geq 0,\; u_t^{-} \geq 0,\\
B \geq x_t^{\varepsilon} \geq -\varepsilon_t - \cdots - \varepsilon_{t+r},\\
R \geq -u_t^{\varepsilon} - \varepsilon_t \geq -R,\; t = 1, \dots, n \,\},
\end{aligned}
$$

where x_t is the state of charge of the battery, ε_t is the energy balance (at time t), R is the maximum charge rate of the battery, and B is its maximum capacity (R and B are given constants).

- We can simulate N i.i.d. scenarios (energy balance trajectories) and solve the resulting scenario problem, that is: a (finite) LP.

Experiment: Controller Orders

– In all experiments, the bound on the constraint violation probability was $\delta = 0.1$, and the confidence probability was $1 - \beta = 0.999$.
– Ten-fold cross validation was applied, the averages as well as the standard deviations (error bars) of the results were evaluated.

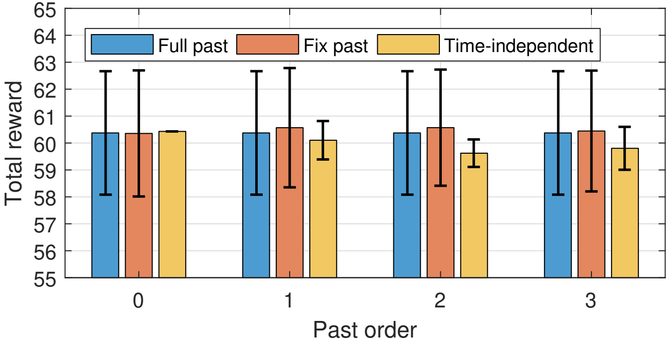

Experiment: Reoptimization Frequencies

– We compared various reoptimization frequencies applying a time-independent affine controller with (fixed) past order 1.
– Our results indicate that the efficiency only slightly decreases over time if we keep using the same parameters for more than one step.

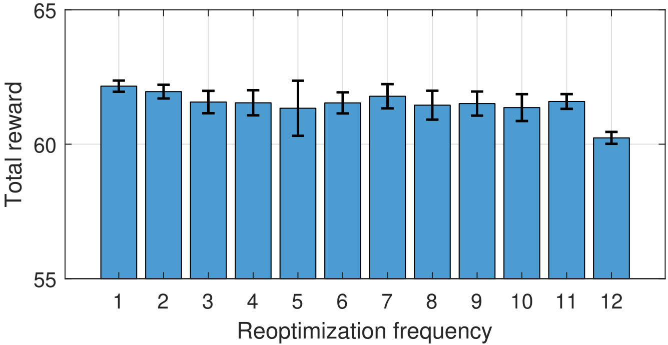

Experiment: Cumulative Rewards

– The notation (hx, ry) encodes that the optimization horizon was x and the reoptimization frequency was y (reward = profit at time t).

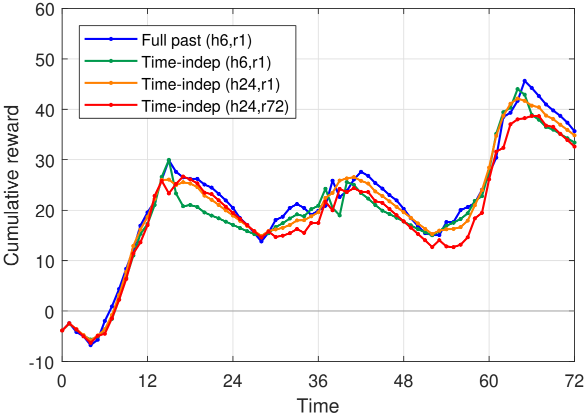

Experiment: Discarding Trajectories

– No. Disc. Traj. vs Guarantees ($N = 1000$, $d = 3$, $\beta = 0.001$).

430	151	67	43	28	21	15	8	3
50%	80%	90%	93%	95%	96%	97%	98%	99%

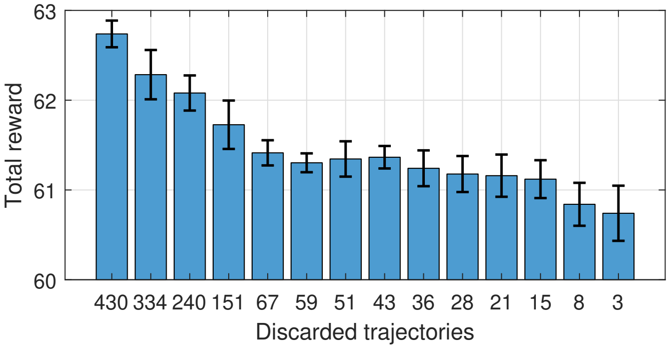

Conclusions

- We studied a sampling-and-discarding method, based on the scenario approach, to SMPC for renewable energy systems.
- It was applied to a public lighting microgrid with LED luminaries, PV panels, a battery and a bidirectional power grid connection.
- The system was decomposed into controllable and uncontrollable parts, the value-at-risk formulation of the problem was overviewed.
- Several experiments were presented, for example, about generating trajectories by bootstrap, the effects of controller parametrizations, reoptimization frequencies and discarding unfavorable scenarios.
- They demonstrated the viability of the approach, even for low order, time-independent controllers that are rarely reoptimized.
- As these controllers have much fewer parameters than a full affine one, stronger stochastic guarantees can be provided for them.

Thank you for your attention!

www.sztaki.hu/~csaji csaji@sztaki.hu