Stochastic MPC for Energy Management in Smart Grids with Conditional Value at Risk as Penalty Function

Venkatasubramanian, Janani ; Rostampour, Vahab; Keviczky, Tamas

DOI
10.1109/ISGT-Europe47291.2020.9248769

Publication date
2020

Document Version
Final published version

Published in
Proceedings of the 2020 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe)

2020 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe)
Virtual, October 26-28, 2020

Stochastic MPC for Energy Management in Smart Grids with Conditional Value at Risk as Penalty Function

Janani Venkatasubramanian
Institute for Systems Theory and Automatic Control, University of Stuttgart
Stuttgart, Germany
janani.venkatasubramanian@ist.uni-stuttgart.de

Vahab Rostampour
Engineering and Technology Institute University of Groningen
Groningen, The Netherlands
v.rostampour@rug.nl

Tamás Keviczky
Delft Center for Systems and Control Delft University of Technology
Delft, The Netherlands
t.keviczky@tudelft.nl

***Abstract*—This paper considers imbalance problems arising in Energy Management in Smart Grids (SG) as discrete-time stochastic linear systems subject to chance constraints, and proposes a Model Predictive Control (MPC) approach to solve them. It is well-known that handling the closed-loop constraint feasibility of such systems is in general difficult due to the presence of a potentially unbounded uncertainty source. To overcome such a difficulty, we propose two new ideas. We first reformulate the chance constraint using the so-called *Conditional Value at Risk* (CVaR), which is known to be the tightest convex approximation for chance constraints. We then *relax* the CVaR constraint using a *penalty function* depending on a *coefficient parameter*. An optimal solution is therefore obtained by solving a single unconstrained problem which, intuitively, takes into consideration a *risk* of the system trajectories in an undesirable state. A case study using an academic example is presented to estimate the a-posteriori probability of the coefficient parameter in order to show when such a penalty function is exact by means of probabilistic constraint fulfillment.**

I. Introduction

In Smart Grid (SG) applications, the best performance that satisfies constraints is usually achieved near the boundary of the feasible set, and thus, potential constraint violation due to uncertainty, e.g., photovoltaics, wind power, is unavoidable. By adding an extra cost (penalty) on the SG system during constraint violation [1], [2], this can be represented as performance degradation. Penalty methods are pervasive in the optimization literature to approximate a constrained problem as an unconstrained problem [3], [4]. In the context of deterministic Model Predictive Control (MPC), a penalty term for state constraint violations was introduced in [5]–[7]. Another interesting work [1], proposes MPC with a penalty method for stochastic linear systems and provides a theoretical guarantee on the solvability of the optimization problem, however, this method fails in the case of unbounded uncertainty.

The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Janani Venkatasubramanian. Vahab Rostampour is supported by the incentives and algorithms for efficient, reliable, sustainable and socially acceptable energy system integration (ERSAS) research program funded by the Dutch organization for scientific research (NWO) with grant number 647-002-005.

MPC using heuristic Monte Carlo sampling for uncertain SG systems has been presented in [8]. Stochastic MPC with a probabilistic view on SG systems with uncertainty was developed in [9]. In this paper, we present a Stochastic MPC strategy with probabilistic (chance) constraint for uncertain SG systems. We then formulate a penalty function using Conditional Value at Risk (CVaR) which is widely known as the *tightest* convex approximation of the chance constraint. Using a weighted penalty function of CVaR to account for performance degradation and to take appropriate measures to allow only a certain admissible level of constraint violation, we finally provide an extensive study on the a-posteriori probability estimation of the coefficient parameter of the penalty function for an academic example SG system.

II. Preliminaries

Consider the following stochastic linear time-invariant system model:

$$\begin{cases} x_{k+1} &= Ax_k + Bu_k + Ew_k \\ z_k &= Cx_k + Du_k \end{cases}, \tag{1}$$

where $x_k \in \mathbb{R}^{n_x}$ is the state, $u_k \in \mathbb{R}^{n_u}$ is the control input and $z_k \in \mathbb{R}^{n_z}$ is the output at time k. The disturbance input $w_k \in \mathbb{R}^{n_w}$ is an exogenous disturbance with unknown current and future values with probability distribution P_w and support $\mathcal{W} \subset \mathbb{R}^{n_w}$. The system matrices A, B, E, C and D are of suitable dimensions with real elements.

The system (1) is subject to constraints on the state and control input. Consider constraints of the form,

$$\mathcal{X} := \{x \in \mathcal{R}^{n_x} : g(x,w) \leq 0\},$$
$$\mathcal{U} := \{u \in \mathcal{R}^{n_u} : h(u) \leq 0\},$$

where $g : \mathbb{R}^{n_x} \times \mathbb{R}^{n_w} \to \mathbb{R}^{m_1}$ and $h : \mathbb{R}^{n_u} \to \mathbb{R}^{m_2}$. The constraints on the state x define a closed convex set $\mathcal{X} \subset \mathcal{R}^{n_x}$. The constraints on the control input u define a non-empty measurable compact convex control set $\mathcal{U} \subset \mathcal{R}^{n_u}$. Let N be a positive integer number that represents the length of prediction horizon. Retaining the state within the feasible set $\mathcal{X}$ for the entire prediction horizon requires the initial state to belong to a set of feasible initial states given as,

$$\mathcal{X}_N = \{ x_k : \exists \mathbf{u}_k \text{ s.t. } x_{i|k} \in \mathcal{X} \text{ and } u_{i|k} \in \mathcal{U} \, \forall k \geq 0 \,, \; i = 0, 1, ..., N-1 \} \,. \tag{2}$$

where $\mathbf{u}_k = \{u_{0|k}, ..., u_{N-1|k}\}$, the sequence of control inputs, is the decision vector in the optimization problem in MPC. This requirement is often too conservative and results in poor performance. In such cases, chance constraints can be viewed as a compromise on the requirement to enforce hard constraints in an uncertain system which may be very expensive or even impossible. Consider now the chance constraint on the state trajectories as,

$$\mathbb{P}(\bar{g}(x, w) \leq 0) \geq 1 - \alpha \tag{3}$$

where $\bar{g}(x, w) = \max\{g_1(x, w), ..., g_{m_1}(x, w)\}$ and is interpreted as the requirement of the probability of any predicted state x not belonging to $\mathcal{X}$ to be less than α.

The optimal control problem is defined in terms of a performance index $J_N(x_k, \mathbf{u}_k, \mathbf{w}_k)$ that is evaluated over the horizon of N steps and is solved at each time step, where $\mathbf{w}_k = \{w_{0|k}, ..., w_{N-1|k}\}$. The predicted cost is given by

$$J_k(x_k, \mathbf{u}_k, \mathbf{w}_k) = p(x_{N|k}) + \sum_{i=0}^{N-1} q(x_{i|k}, u_{i|k}) \tag{4}$$

where the function $q : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}_+$ gives the stage cost and the function $p : \mathbb{R}^{n_x} \to \mathbb{R}_+$ is the terminal cost. To ensure that the optimal value of the cost is well defined, we have the following assumption. The matrix pair (A, B) is stabilizable, the matrix pair (A, C) is observable, and all eigenvalues of the system matrix A lie on or inside the unit circle.

To account for the stochasticity of w_k, it is appropriate to use the expectation of the predicted cost $J_k(x_k, \mathbf{u}_k, \mathbf{w}_k)$, which is given by

$$\bar{J}_k(x_k, \mathbf{u}_k) = \mathbb{E}\left[\left(p(x_{N|k}) + \sum_{i=0}^{N-1} q(x_{i|k}, u_{i|k}) \right) \Big| x_{0|k} \right] \tag{5}$$

where $x_{0|k} = x_k$ is the initial state measured at time k. The terminal cost ensures that the system is closed-loop stable and the controller found is stabilizing.

The functions $p(\cdot)$ and $q(\cdot)$ are assumed to be convex functions.

Given the performance index and constraints, we are now able to formulate the following chance constrained optimization as the stochastic control problem,

$$\begin{aligned} J_k^*(x_k) = \min_{\mathbf{u}_k} \quad & \bar{J}_k(x_k, \mathbf{u}_k) \,, \\ \text{s.t.} \quad & x_{i+1|k} = A x_{i|k} + B u_{i|k} + E w_{i|k} \,, \\ & \mathbb{P}(\bar{g}(x, w) \leq 0) \geq 1 - \alpha \\ & u_{i|k} \in \mathcal{U} \,, \; i = 0, \cdots, N-1 \,, \\ & x_{0|k} = x_k \,. \end{aligned} \tag{6}$$

$J_k^*(x_k)$ is the optimal value function. The optimal control sequence $\mathbf{u}_k^*$ is applied to the system following the *receding horizon principle* [10]. However, there are a few significant difficulties with this optimization to handle chance constraints [11]. Monte-Carlo simulations and its related methods lead to heavy computational costs, and the risk of an initial infeasible state and thereby an unsolvable problem. This motivates the perspective of approximating chance constraints using CVaR and implementing them as a penalty function, leading to the main contribution of this paper.

III. Measuring Conditional Value at Risk

Consider the Value at Risk (VaR) of the function $\bar{g}(x, w)$ defined in chance constraint (3) as follows:

$$\text{VaR}_{1-\alpha}(\bar{g}(x, w)) := \min_{\eta \in \mathbb{R}} \{\eta : \mathbb{P}(\bar{g}(x, w) \leq \eta) \geq 1 - \alpha\} \tag{7}$$

While VaR represents the worst-case loss with a probability [11], [12], CVaR represents the expected loss if the worst-case threshold (VaR) is crossed. The CVaR is defined as the conditional expectation of $\bar{g}(x, w)$ exceeding VaR using the following relation:

$$\text{CVaR}_{1-\alpha}(\bar{g}(x, w)) := \mathbb{E}[\bar{g}(x, w) | \bar{g}(x, w) \geq \text{VaR}_{1-\alpha}(\bar{g}(x, w))],$$

which can also be formulated as,

$$\text{CVaR}_{1-\alpha}(\bar{g}(x, w)) = \min_{\eta \in \mathbb{R}} \left(\eta + \frac{1}{\alpha} \mathbb{E}[(\bar{g}(x, w) - \eta)_+] \right),$$

where $(\nu)_+$ is nonzero when $\nu > 0$. The formulated chance constraint in (3) can be then replaced by the CVaR constraint as,

$$\text{CVaR}_{1-\alpha}(\bar{g}(x, w)) \leq 0 \tag{8}$$

As required in CVaR, an exact evaluation of the expected value of a random function $\bar{g}(x, w)$ is expensive and hence a sampled-average approximation is used [13].

The sample-average approximation of CVaR is now given by

$$\text{CVaR}_{1-\alpha}(\bar{g}(x, w)) = \min_{\eta \in \mathbb{R}} \left(\eta + \frac{1}{\alpha N_s} \sum_{j=1}^{N_s} ((\bar{g}(x, w^j) - \eta)_+) \right),$$

where N_s represents a number of independent and identically distributed (i.i.d.) samples $w^1, ..., w^{N_s}$, called 'scenarios'. The CVaR constraint is now reduced to a combination of multiple affine constraints.

IV. Penalizing CVaR

The expected cost including the penalty function for the CVaR constraint violation is of the following form

$$J_{p,k}(x_k, \mathbf{u}_k) = \bar{J}_k(x_k, \mathbf{u}_k) + \lambda \sum_{i=0}^{N-1} p_{\mathcal{X}}(x_{i|k}) \tag{9}$$

where $p_{\mathcal{X}} : \mathbb{R}^{n_x} \to \mathbb{R}$ is the penalty function and λ is the coefficient parameter of the penalty function. The penalty on a feasible state which does not violate constraints is set to zero, i.e., $p_{\mathcal{X}}(x) = 0$, $x \in \mathcal{X}$ and for an infeasible state that violates constraints, $p_{\mathcal{X}}(x) > 0$, $x \notin \mathcal{X}$. However, the penalty on the infeasible state cannot be arbitrarily high as a choice of infinite penalty will make the optimization problem unsolvable as the cost is infinite at the infeasible states [1].

The simplest form of a penalty function that quantifies constraint violation is as follows:

$$p_{\mathcal{X}}(x) := \max\{0, \mathrm{CVaR}_{1-\alpha}(\bar{g}(x,w))\}\,. \tag{10}$$

Note that using the N_s number of scenarios, we also approximate (sampled average approximation) empirically the expected stage costs and terminal cost along with CVaR as follows:

$$J_{p,k}(x_k,\mathbf{u}_k,\boldsymbol{\eta}) = \frac{1}{N_s}\sum_{j=1}^{N_s}\Big(p(x^j_{N|k}) + \sum_{i=0}^{N-1} q(x^j_{i|k},u_{i|k})\Big)$$
$$+\lambda \max\Big\{0, \sum_{i=0}^{N-1}\Big[\eta_i + \frac{1}{\alpha N_s}\sum_{j=1}^{N_s}([g(x^j_{i|k},w^j_{i|k}) - \eta_i]_+)\Big]\Big\}\,,$$

where $x_{0|k} = x_k$ and $\boldsymbol{\eta} = \{\eta_1,...,\eta_N\}$. Note that if λ takes higher values, then, the approximate unconstrained problem becomes increasingly accurate and closer to the original constrained problem.

The reformulated problem using CVaR as a penalty function is now given as,

$$\begin{aligned} J^*_{p,k}(x_k) = \min_{\bar{\mathbf{u}}_k} \quad & J_{p,k}(x_k,\mathbf{u}_k,\boldsymbol{\eta})\,, \\ \text{s.t.} \quad & x_{i+1|k} = Ax_{i|k} + Bu_{i|k} + Ew_{i|k}\,, \\ & u_{i|k} \in \mathcal{U}\,,\, i = 0,\cdots,N-1\,, \\ & x_{0|k} = x_k\,,\, w \in \mathcal{W}\,. \end{aligned} \tag{11}$$

The above optimization problem is solved at each time step k and the first element of the optimal control input sequence, $u^*_{0|k}$, is applied to the system. The horizon is then shifted according to the receding horizon principle and the optimization problem is solved again.

V. Illustrative Example

In this section, we provide a case study similar to [1] to demonstrate the functionality of our proposed stochastic MPC using CVaR as a penalty function. Such a case study can represent an energy management framework for an interconnected network of buildings in the SG settings proposed in [14], where a large-scale stochastic dynamical model is developed to predict the thermal energy imbalance in a smart thermal grid in the following form:

$$\begin{cases} x_{k+1} &= \begin{pmatrix}1 & 0\\ 1 & 1\end{pmatrix} x_k + \begin{pmatrix}1\\0\end{pmatrix} u(k) + \begin{pmatrix}1\\0\end{pmatrix} w_k \\ z_k &= \begin{pmatrix}0 & 0\\ 0.7 & 0\\ 0 & 0.7\end{pmatrix} x_k + \begin{pmatrix}0.33\\0\\0\end{pmatrix} u_k \end{cases}\,. \tag{12}$$

The input is constrained as,

$$-0.5 \le u_k \le 0.5,\ u_k \in \mathbb{R},\ k \in \mathbb{N}\,. \tag{13}$$

The disturbance is assumed to be a Gaussian distributed random variable with zero mean and variance 0.2 in the following form:

$$w_k \in \mathcal{G}(0, 0.2),\ w_k \in \mathbb{R}\,. \tag{14}$$

The state is $x = [x_1\,,\, x_2]^\top$ and the constraint on the state is given as $x_2 \ge 0$. The initial system state is $x_0 = [0\,,\, 10]^\top$. In the presence of uncertainty, the task of the controller is to steer the state to the origin while satisfying the constraints on the system. The cost function used is of the form,

$$\bar{J}_k(x_k,\mathbf{u}_k) = \mathbb{E}\left[||x_{N|k}||^2_P + \sum_{i=0}^{N-1}(||x_{i|k}||^2_Q + ||u_{i|k}||^2_R)\right],$$

where the approximate matrices for the stage costs are,

$$Q = \begin{pmatrix}0.7^2 & 0\\ 0 & 0.7^2\end{pmatrix},\ R = \begin{pmatrix}0.33^2\end{pmatrix}\,, \tag{15}$$

and the weighting matrix for the terminal cost, by solving the discrete-time Riccati equation, is given as,

$$P = \begin{pmatrix}1.5534 & 0.9025\\ 0.9025 & 1.3334\end{pmatrix}\,. \tag{16}$$

We then solve the proposed formulation in (11) by selecting the prediction horizon $N = 5$ and the expected average-over-time constraint violation level $\epsilon = 0.05$.

Two sets of simulation results, one for deterministic MPC and one for stochastic MPC (11), are presented with 100 trials in each set. The design of the controller for deterministic MPC is elaborated in [15]. In both deterministic MPC and stochastic MPC, the hard state constraint is implemented using the proposed penalty function. However, in deterministic MPC the uncertain variable is fixed to its mean value, zero, over the control horizon.

Fig. 1 shows the trajectory of x_2 using deterministic MPC and the proposed stochastic MPC. When the state of the system is 'far' from the constraint boundary, both standard deterministic MPC and stochastic MPC perform similarly. However, the controller in deterministic MPC does not predict the possibility of a constraint violation when the system is close to the constraint boundary. The stochastic MPC, on the other hand, takes into account the possibility of constraint violation due to disturbance by using a number of sample disturbance values and hence, provides more realistic control input.

Fig. 2 presents the estimation of the probability density of the state x_2 using both controllers, stochastic MPC and deterministic MPC. The probability of constraint violation at all times is significantly lower while applying stochastic MPC as shown in Fig. 2. Furthermore, since the performance of stochastic MPC is determined by the penalty function, the performance must improve with an increase in the penalty parameter as the 'weight' of the 'risk' will increase in case of constraint violation.

When the penalty function is *exact* [4], and the penalty parameter chosen is greater in value than the Lagrange multiplier associated to the problem, a single unconstrained minimization

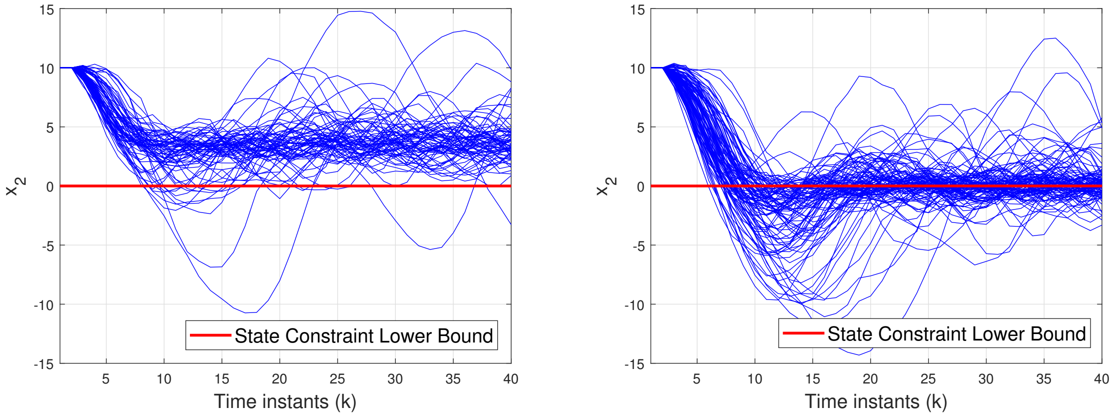

(a) Stochastic MPC (b) Deterministic MPC

Fig. 1: Trajectory of state x_2, Constraint: $x_2 \geq 0$.

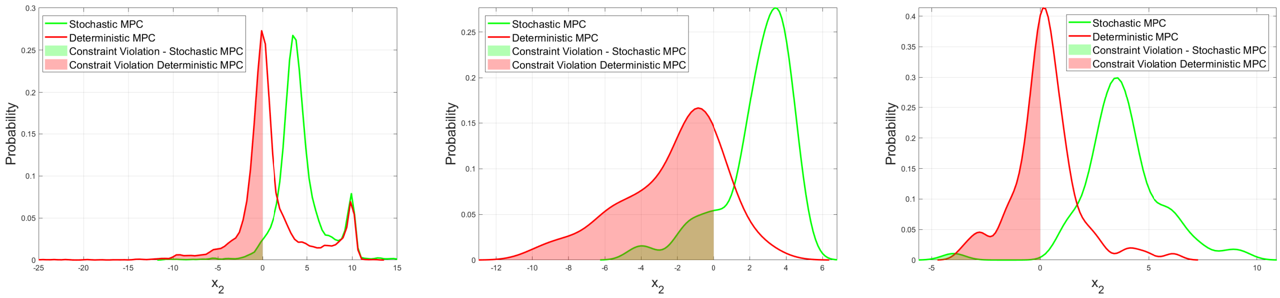

(a) (b) (c)

Fig. 2: An estimate of the probability density of x_2. (2a) over 40 time instants, (2b) at time instant $k = 11$ (overshoot region) and (2c) at time instant $k = 40$ (steady state region)

of the cost yields the same optimal solution as the original problem. A lower bound on the Lagrange multiplier can be obtained as in [16], however this method uses the optimal value of the dual function of the cost which is not readily available. A heuristic approach is therefore applied to compute a suitable bound on the penalty parameter. Fig. 3 shows a gradual increase in the mean of x_2 with increase in the penalty parameter. Further, from data, it can be observed that the median vlaue of x_2 gradually increases indicating higher likelihood of constraint satisfaction. Fig. 4a shows the a-posteriori probability of achieving the optimal median value of x_2 with different penalty parameters, and Fig. 4b presents the trajectory of x_2 using different penalty parameters. As expected, performance increases with an increase in the penalty parameter. As the a-posteriori probability of the parameter does not increase drastically beyond a parameter value of 100, it can be inferred that the exact penalty parameter value is close to 100. A parameter beyond this value will yield similar system behaviour.

VI. Conclusions

In this paper, we account for the performance degradation, due to constraint violation, by adding a penalty into the cost function of the system at the infeasible states to give rise to a penalty method for optimization. A suitable choice of a penalty function is the CVaR function which determines the risk faced by the system at an infeasible state. Using the CVaR constraint, we accommodate the expected performance loss when the worst-case threshold for constraint violation is crossed. A coefficient parameter is used to weigh the penalty function. We then provide a heuristic study by varying such a parameter to demonstrate its effects on performance.

References

[1] I. Batina, A. A. Stoorvogel, S. Weiland *et al.*, "Optimal control of linear, stochastic systems with state and input constraints," in *Conference on Decision and Control*. IEEE, 2002, pp. 1564–1569.

[2] I. Batina, *Model predictive control for stochastic systems by randomized algorithms*. Ph.D. Dissertation, Technische Universiteit Eindhoven, 2004.

[3] D. P. Bertsekas, *Nonlinear programming*. Athena Scientific, Belmont, 1999.

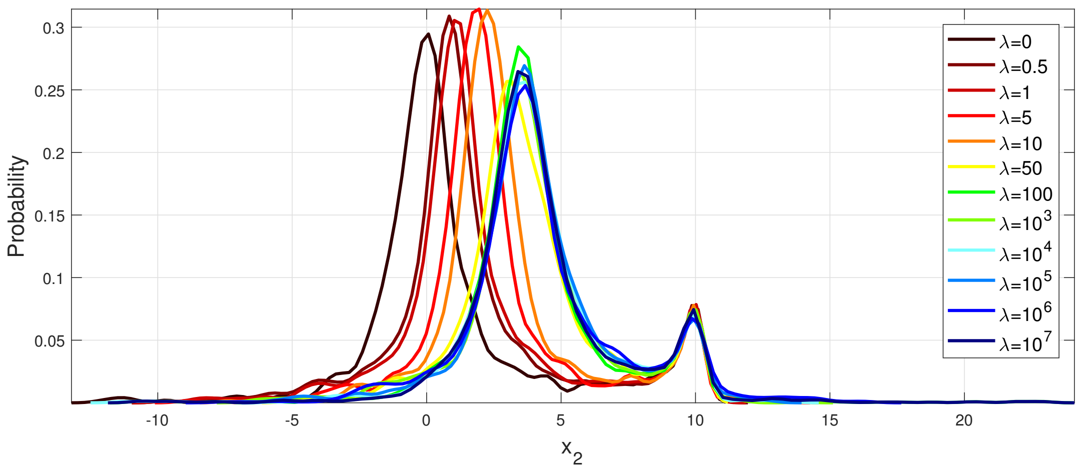

Fig. 3: Estimate of probability density of x_2 over 40 time instants over 100 trials for different values of the penalty parameter.

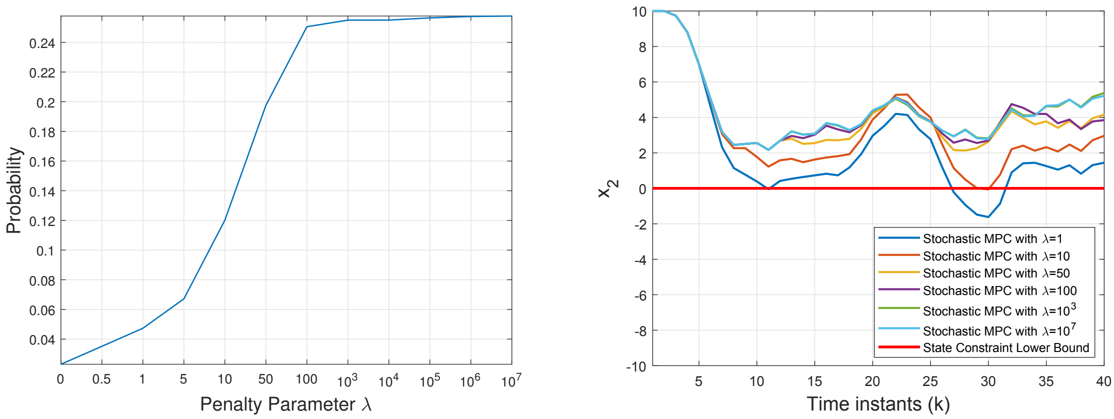

(a) A-posteriori probability of penalty parameter.

(b) Constraint: $x_2 \geq 0$ with different penalty parameter

Fig. 4: A-posteriori probability of penalty parameter together with closed-loop system trajectories of state x_2.

[4] ——, "Necessary and sufficient conditions for a penalty method to be exact," *Mathematical programming*, vol. 9, no. 1, pp. 87–99, 1975.

[5] A. Zheng and M. Morari, "Stability of model predictive control with mixed constraints," *IEEE Transactions on Automatic Control*, vol. 40, no. 10, pp. 1818–1823, 1995.

[6] P. O. Scokaert and J. B. Rawlings, "Feasibility issues in linear model predictive control," *AIChE Journal*, vol. 45, no. 8, pp. 1649–1659, 1999.

[7] E. C. Kerrigan and J. M. Maciejowski, "Soft constraints and exact penalty functions in model predictive control," in *Proc. UKACC International Conference (Control)*, 2000.

[8] S. S. Farahani, Z. Lukszo, T. Keviczky, B. De Schutter, and R. M. Murray, "Robust model predictive control for an uncertain smart thermal grid," in *2016 European Control Conference (ECC)*. IEEE, 2016, pp. 1195–1200.

[9] V. Rostampour and T. Keviczky, "Robust randomized model predictive control for energy balance in smart thermal grids," in *2016 European Control Conference (ECC)*. IEEE, 2016, pp. 1201–1208.

[10] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[11] A. Nemirovski and A. Shapiro, "Convex approximations of chance constrained programs," *SIAM Journal on Optimization*, vol. 17, no. 4, pp. 969–996, 2006.

[12] H. Sun, H. Xu, and Y. Wang, "Asymptotic analysis of sample average approximation for stochastic optimization problems with joint chance constraints via conditional value at risk and difference of convex functions," *Journal of Optimization Theory and Applications*, vol. 161, no. 1, pp. 257–284, 2014.

[13] W. Wang and S. Ahmed, "Sample average approximation of expected value constrained stochastic programs," *Operations Research Letters*, vol. 36, no. 5, pp. 515–519, 2008.

[14] V. Rostampour and T. Keviczky, "Probabilistic energy management for building climate comfort in smart thermal grids with seasonal storage systems," *IEEE Transactions on Smart Grid*, vol. 10, no. 4, pp. 3687–3697, 2018.

[15] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos, "The explicit linear quadratic regulator for constrained systems," *Automatica*, vol. 38, no. 1, pp. 3–20, 2002.

[16] A. Nedić and A. Ozdaglar, "Approximate primal solutions and rate analysis for dual subgradient methods," *SIAM Journal on Optimization*, vol. 19, no. 4, pp. 1757–1780, 2009.