Presented as late breaking results contribution
21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

Indirect and Direct Feedback in Stochastic Model Predictive Control ⋆

Lukas Hewing and Melanie N. Zeilinger

Institute for Dynamic Systems and Control, ETH Zurich, Zurich, Switzerland (e-mail: lhewing|mzeilinger@ethz.ch).

Abstract: In contrast to the unconstrained case, there is no closed-form solution to constrained optimal control problems of linear systems under additive stochastic noise. Stochastic model predictive control (SMPC) is an approximate solution strategy for such problems, in which a simplified problem is repeatedly solved over a reduced prediction horizon. In this contribution, we compare two forms of feedback in SMPC formulations in terms of their closed-loop performance, as well as their conservatism regarding constraint satisfaction. First, we consider a direct feedback formulation, which corresponds to the typical implementation of SMPC schemes. This formulation aims to satisfy constraints with respect to the predicted state distribution conditioned on the current measured state at each time step during the receding horizon control. The second, denoted indirect feedback, introduces feedback through the cost only, and instead considers constraints by introducing a suitable virtual or nominal state. This results in a linear evolution of a closed-loop error state which can be used for constraint tightening, providing closed-loop constraint satisfaction. In numerical examples, we demonstrate that this can significantly improve performance, as well as reduce conservatism in closed-loop and that it recovers the unconstrained optimal solution given by LQR control when it is feasible also for the constrained optimal control problem.

Keywords: stochastic model predictive control, chance constraints, predictive control

1. INTRODUCTION

Stochastic model predictive control (SMPC) can be viewed as an approximate solution strategy for constrained stochastic optimal control problems (SOCP), in which an often simplified problem is repeatedly solved in receding horizon fashion. This view raises questions regarding the formulation of constraints in the SMPC optimization problem since one typically desires satisfaction of constraints with respect to the closed-loop system, i.e. satisfaction of the original constraints of the SOCP. Expressing this in the formulation of the receding horizon SMPC problem, however, proves challenging and established SMPC approaches typically aim to derive a sequence of control inputs or control laws such that the predicted uncertain state sequence conditioned on the currently measured state satisfies the chance constraints. We refer to this as a *direct feedback* formulation, see for instance Lorenzen et al. (2017); Farina et al. (2015); Hewing and Zeilinger (2018); Cannon et al. (2011); Kouvaritakis et al. (2010); Korda et al. (2011). While the prediction from the state measurement introduces feedback in the receding horizon control formulation, and is similar to nominal MPC, constraint satisfaction is often hard to relate to the original SOCP, i.e. the closed-loop control system, for instance due to feasibility issues.

For disturbances of bounded support, robust MPC techniques can be used to ensure recursive feasibility, i.e. the persistent satisfaction of constraints in prediction, from which (conservative) satisfaction of the original or closed-loop constraints follows (Lorenzen et al., 2017; Kouvaritakis et al., 2010). One way to address the conservatism in these approaches was presented in Korda et al. (2014), in which the constraint tightening is adapted over time to adjust to a desirable level of satisfaction. For disturbances of unbounded support, however, or if recursive feasibility cannot be guaranteed, guarantees with respect to the closed-loop system or original SOCP are lost (Kouvaritakis and Cannon, 2016; Farina et al., 2016); see also Hewing and Zeilinger (2018) where guarantees are recovered under additional assumptions. It can therefore be argued, that the formulation of chance constraints conditioned on the measured state is a main source of feasibility issues in direct feedback SMPC, and can lead to conservatism and decreased performance, even if recursively feasible.

This contribution investigates these issues by means of comparison with an SMPC approach recently proposed in Hewing et al. (2018), which uses an *indirect feedback* formulation to express chance constraint with respect to a closed-loop distribution, i.e. directly according to the original SOCP. We present a re-formulation of this procedure emphasizing this relationship, and facilitating further investigations. We compare the formulation to an established direct feedback variant, as well as the true optimal solution of the original SOCP and show that—in contrast to direct feedback—the indirect feedback variant recovers this true optimal solution, while direct feedback results in very conservative closed-loop trajectories, allowing essentially no violation of constraints for any noise realization.

⋆ This work was supported by the Swiss National Science Foundation under grant no. PP00P2_157601 / 1.

Presented as late breaking results contribution
21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

2. PROBLEM STATEMENT

To illustrate the effect of the different feedback formulations, we consider a particularly simple constrained stochastic optimal control problem of linear time-invariant systems under additive Gaussian noise

$$x(k+1) = Ax(k) + Bu(k) + w(k) \tag{1}$$
$$w(k) \sim \mathcal{N}(0, \Sigma_w), \text{ i.i.d.}$$

which are subject to a single half-space chance constraint

$$\Pr(h^\mathsf{T} x(k) \leq 1) > p, \tag{2}$$

which needs to be satisfied at each time step individually with a minimum specified probability p. While feasibility is a central issue in stochastic MPC, we focus here on performance and possible conservatism, and assume no input constraints and that $h^\mathsf{T} B \neq 0$, such that the constraint can be satisfied at all times from any measured state $x(k)$. The discussed MPC approaches are, however, applicable to a wider range of use cases, see for instance Hewing and Zeilinger (2018); Hewing et al. (2018), similarly also under hard input constraints (Hewing and Zeilinger, 2020).

We consider a quadratic cost and large, but finite horizon $\bar{N}$, resulting in the following constrained SOCP which we aim to approximately solve using MPC techniques:

$$\bar{J}^* = \min_{\Pi} \ \mathbb{E}\left(\|x(\bar{N})\|_P^2 + \sum_{k=0}^{\bar{N}-1} \|x(k)\|_Q^2 + \|u(k)\|_R^2 \right) \tag{3a}$$
$$\text{s.t.} \quad x(k+1) = Ax(k) + Bu(k) + w(k), \tag{3b}$$
$$u(k) = \pi_k(x(0), \ldots, x(k)), \tag{3c}$$
$$w(k) \sim \mathcal{N}(0, \Sigma_w), \text{ i.i.d.}, \tag{3d}$$
$$\Pr(h^\mathsf{T} x(k) \leq 1) > p, \tag{3e}$$
$$x(0) = x_{\text{init}}, \tag{3f}$$

where we optimize over a sequence of control laws $\Pi = \{\pi_0, \ldots, \pi_{\bar{N}-1}\}$ making use of the available information up to that time step in form of state measurements.

Assumption 1. The terminal weight P is given by the solution to the associated Riccati equation

$$P = A^\mathsf{T} P A - A^\mathsf{T} P B (R + B^\mathsf{T} P B)^{-1} B^\mathsf{T} P A + Q.$$

Under this assumption, (3a) can also be interpreted as a reformulation of an infinite horizon problem, given that the unconstrained solution is feasible after time step $\bar{N}$. In the unconstrained case, i.e. without constraint (3e), the optimal solution to (3) is analytically available (Bertsekas, 2017) and given by the infinite horizon LQR controller $\pi_{\text{LQR}}(x) = K_{\text{LQR}} x$ with

$$K_{\text{LQR}} = (R + B^\mathsf{T} P B)^{-1} B^\mathsf{T} P A, \tag{4}$$

which we use for comparison in Section 4. In the constrained case, however, there is no general closed-form solution to (3) and approximate solution schemes are used, for instance stochastic MPC.

3. APPROXIMATE SOLUTION BY STOCHASTIC MODEL PREDICTIVE CONTROL

The approximations in deriving an SMPC approximation to problem (3) are twofold: first, we consider an optimization over a reduced horizon $N \ll \bar{N}$ and use the controller in receding horizon, i.e. implementing the first computed control input and resolving the optimization problem at the following time step. Note that the approximate solution to the SOCP (3) therefore consists of the receding horizon control law, rather than the optimization problem solved in the prediction of the MPC. We initialize the predicted state sequence at each currently measured state

$$x_{i+1} = Ax_i + Bu_i + w_i \tag{5a}$$
$$x_0 = x(k), \tag{5b}$$

with $w_i \sim \mathcal{N}(0, \Sigma_w)$ i.i.d., such that x_i represents the i-step ahead prediction, conditioned on $x(k)$. To emphasize the difference, we use the index i for predictive quantities, and parentheses k for the closed-loop.

Second, we decompose the system state into mean and deviation $x = \bar{x} + d$ in the prediction of the MPC optimization problem, and restrict the considered controller class by only optimizing over the mean input $\bar{u}$ of

$$u = Kd + \bar{u}$$

with predefined stabilizing gain K, which we also refer to as the tube controller. This results in mean and deviation prediction dynamics

$$\bar{x}_{i+1} = A\bar{x}_i + B\bar{u}_i, \tag{6a}$$
$$d_{i+1} = (A + BK)d_i + w_i, \tag{6b}$$
$$\text{var}(d_i) = \sum_{j=0}^{i} (A + BK)^j \Sigma_w ((A + BK)^\mathsf{T})^j. \tag{6c}$$

Note that the deviation d is zero mean and independent of the chosen mean control actions $\bar{u}$.

3.1 Direct Feedback Formulation

In a direct feedback formulation, we aim to enforce chance constraints (2) for the prediction dynamics conditioned on the measured state, i.e. for x_i in (5a). Since

$$h^\mathsf{T} \bar{x}_i + \phi^{-1}(p)\sqrt{h^\mathsf{T} \text{var}(d_i) h} \leq 1$$
$$\Leftrightarrow \Pr(h^\mathsf{T} x_i \leq 1) > p,$$

with ϕ^{-1} the quantile function of the standard Gaussian distribution, we can reformulate the chance constraint deterministically in terms of $\bar{x}$. Using the fact that one can equivalently optimize the cost at the mean for a quadratic objective under this controller class, this results in the direct feedback SMPC optimization problem

$$J_{\text{d}}^* = \min_{\bar{U}} \quad \|\bar{x}_N\|_P^2 + \sum_{k=0}^{N-1} \|\bar{x}_i\|_Q^2 + \|\bar{u}_i\|_Q^2 \tag{7a}$$
$$\text{s.t.} \quad \bar{x}_{i+1} = A\bar{x}_i + B\bar{u}_i, \tag{7b}$$
$$h^\mathsf{T} \bar{x}_i \leq 1 - \phi^{-1}(p)\sqrt{h^\mathsf{T} \text{var}(d_i) h}, \tag{7c}$$
$$\bar{x}_0 = x(k) \tag{7d}$$

where the applied input to system (1) is

$$u(k) = \bar{u}_0^*.$$

Note that this therefore corresponds to nominal MPC on tightened constraints, and that $\text{var}(d_i)$ can be precomputed along (6c) for the constraint tightening.

Note also that constraint (7c) reformulates the chance constraints for the predictive state distributions x_i, i.e. it provides *open-loop* constraint satisfaction. Given this simple setup without input constraints and only a single half-space state constraint, we can then guarantee recursive

Presented as late breaking results contribution
21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

feasible of (7) and (conservative) closed-loop constraint satisfaction follows immediately, since

$$\Pr(h^\mathsf{T} x(k+1) \leq 1) \geq \int \underbrace{\Pr(h^\mathsf{T} x(k+1) \leq 1 \,|\, x(k))}_{\geq p \text{ by (7c) and (7d)}} p_{x(k)}(x)\mathrm{d}x \geq p.$$

where $p_{x(k)}$ is the probability density of $x(k)$. Note that, in general, receding horizon controllers can lose closed-loop chance constraint guarantees due to loss of feasibility (Hewing and Zeilinger, 2018). Furthermore, enforcing these constraints in the prediction satisfies closed-loop constraints conservatively, as we investigate in Section 4. A way to address this is by *indirect feedback* SMPC formulations, presented in the following.

3.2 Indirect Feedback Formulation

In an indirect feedback formulation, one does not enforce constraints on each predicted state sequence conditioned on the current state, but rather assures satisfaction directly for the closed-loop system, i.e. for $x(k)$. This is enabled by introducing a nominal or virtual state $z(k)$ and error $e(k) = x(k) - z(k)$ and the following MPC formulation

$$J^*_{\text{ind}} = \min_{V} \;\; \|\bar{x}_N\|^2_P + \sum_{k=0}^{N-1} \|\bar{x}_i\|^2_Q + \|\bar{u}_i\|^2_Q \tag{8a}$$

$$\text{s.t.} \quad \bar{x}_{i+1} = A\bar{x}_i + B\bar{u}_i, \tag{8b}$$

$$z_{i+1} = Az_i + Bv_i, \tag{8c}$$

$$\bar{e}_i = x_i - z_i, \tag{8d}$$

$$\bar{u}_i = K\bar{e}_i + v_i, \tag{8e}$$

$$h^\mathsf{T} z_i \leq 1 - \phi^{-1}(p)\sqrt{h^\mathsf{T} \mathrm{var}(e(i+k))h}, \tag{8f}$$

$$\bar{x}_0 = x(k), z_0 = z(k), \tag{8g}$$

where the applied input to system (1) is

$$u(k) = Ke(k) + v(k), \tag{9a}$$

$$v(k) = v^*_0 \tag{9b}$$

and we let our virtual system state $z(k)$ evolve according to $v(k)$, leading to

$$z(k+1) = Az(k) + Bv(k), \tag{10a}$$

$$e(k+1) = (A + BK)e(k) + w(k). \tag{10b}$$

This results in the fact that we have a linear closed form expression for the dynamics of the error in closed-loop, i.e. the deviation of the state $x(k)$ from the virtual state $z(k)$, which evolves autonomously from the chosen MPC control input. Note that similar to the *direct feedback* case, we still have $\bar{x}_i$ as the predicted mean of the state conditioned on the current measurement $x(k)$, such that feedback through state measurement $x(k)$ is introduced to the system. Since in general, the error $e(k)$ is nonzero, we furthermore have that the applied mean input in the prediction is given by

$$\bar{u}_i = K\bar{e}_i + v_i,$$

where $\bar{e}$ is the predicted mean error, conditioned on $x(k)$. As before, we can precompute the constraint tightening, which is done now considering the closed-loop error distribution, by computing $\mathrm{var}(e(i+k))$ beforehand.

In contrast to the direct feedback case, we therefore formulate constraints with respect to $e(k)$ and due to the closed-loop error dynamics (10b), enforcing constraint (8f) immediately guarantees constraint satisfaction with respect to $x(k)$, i.e. in *closed-loop* (as opposed to in prediction). To summarize, *indirect feedback* optimizes the objective with respect to the predicted distribution of the state conditioned on the current measurement $x(k)$, whereas constraints are satisfied with respect to the closed-loop distribution of the error $e(k)$, which is analytically available due to the choice of virtual state $z(k)$. This ensures recursive feasibility, even in the presence of additional (input) constraints (Hewing et al., 2018), and optimizing the cost w.r.t. x_i introduces feedback, which also affects the virtual state trajectory $z(k)$.

4. NUMERICAL COMPARISON

In the following, we investigate performance and conservatism of the presented direct and indirect feedback approach in a stochastic optimal control problem, which we deliberately chose such that the unconstrained optimal solution, i.e. the LQR solution, is feasible and therefore represents the optimal solution also to the constrained problem. We can therefore use this LQR solution as a benchmark to compare the SMPC schemes.

4.1 Setup

We consider a specific SOCP problem of form (3) for a five dimensional integrator chain system

$$A = \begin{bmatrix} 1 & T_s & T_s^2/2! & T_s^3/3! & T_s^4/4! \\ 0 & 1 & T_s & T_s^2/2! & T_s^3/3! \\ 0 & 0 & 1 & T_s & T_s^2/2! \\ 0 & 0 & 0 & 1 & T_s \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \; B = \begin{bmatrix} T_s^5/5! \\ T_s^4/4! \\ T_s^3/3! \\ T_s^2/2! \\ T_s \end{bmatrix}$$

with $T_s = 0.1$ subject to additive i.i.d. disturbances

$$w(k) \sim \mathcal{N}(0, 1.5BB^\mathsf{T}).$$

The quadratic objective is given with cost matrices

$$Q = \mathrm{diag}([1,\, 2,\, 1,\, 1,\, 1]^\mathsf{T}),\; R = 0.1$$

and we consider a fixed initial condition

$$x(0) = [0.3,\, 0,\, 0,\, 0,\, 0]^\mathsf{T},$$

resulting in no overshoot of the nominal system under LQR control (4) in the first state. The constraint on the system is chosen as a constraint on the first state

$$h^\mathsf{T} = [1/\sqrt{h^\mathsf{T}\Sigma_\infty h},\, 0,\, 0,\, 0,\, 0],\; p = 84\%,$$

with Σ_∞ as the solution to the Lyapunov equation

$$\Sigma_\infty = (A + BK_{\text{LQR}})\Sigma_\infty(A + BK_{\text{LQR}})^\mathsf{T} + 1.5BB^\mathsf{T},$$

corresponding to the terminal variance of the system under LQR control. Due to this specific choice, and the fact that there is no overshoot for the nominal system trajectory, the chance constraint is satisfied under LQR control, representing the optimal solution to the problem and providing a benchmark to compare against. For both SMPC formulations, we make use of this LQR controller as tube controller $K = K_{\text{LQR}}$.

4.2 Numerical Results

We carried out 5000 simulations for $\bar{N} = 200$ time steps with different noise realizations for each controller type, namely *LQR*, *indirect feedback* and *direct feedback* SMPC, each with prediction horizon $N = 30$. The results are

Presented as late breaking results contribution
21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

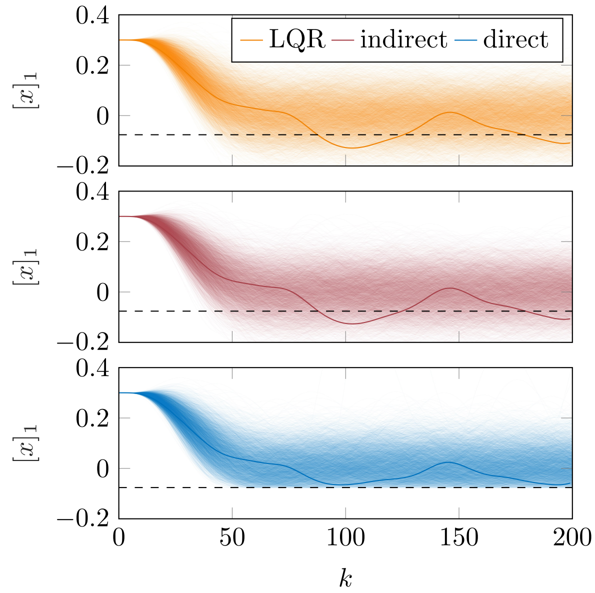

Fig. 1. Results of 5000 simulations for different control formulations showing the trajectory of the first state $[x]_1$, with one noise realization highlighted. The dashed line illustrates the chance constraint.

displayed in Figure 1, where the top plot shows closed-loop trajectory realizations under the optimal control law provided by the LQR controller. The resulting cost $\bar{J}^*$ from (3) can be analytically computed as

$$\bar{J}^* = 28.6,$$

together with a maximum probability of constraint violation under LQR control of 16%, due to our specific choice of constraint. Quantified results of the performed simulation runs are shown in Table 1 and show that these theoretical quantities are well matched by the average cost and the worst case empirical violation rate, computed by averaging constraint violations over all 5000 simulation runs and selecting the maximum resulting empirical violation rate over all time steps. From Figure 1 *(middle)* and the values in Table 1, it is evident that the *indirect feedback* SMPC formulation shows similar behavior to the optimal LQR solution, with the apparent small differences due to numerical noise and the finite simulation samples. This, however, is different in a *direct feedback* formulation, which represents the most common implementation of stochastic MPC schemes. Here, we observe virtually no constraint violations. In fact, we only observe 1 violation out of 5000 simulations with 200 time steps, i.e. a 0.02% maximum empirical constraint violation rate. Visually, it is similarly evident in Figure 1 that the chance constraint essentially acts as a hard constraint in the *direct feedback* formulation. Furthermore, Table 1 lists the total number of instances over all time steps and simulations in which the magnitude of the applied control input exceeds $|u(k)| \geq 4$. This number is severely elevated in the *direct* scheme, illustrating that large control actions must be applied in order arrive at a predicted distribution satisfying the constraint.

5. CONCLUSION

This note presented a discussion and performance comparison of a common implementation scheme in stochastic MPC, which we name *direct feedback*, to an alternative

Table 1. Performance Comparison

Controller	Cost	Empirical Violation	$\# \, \|u(k)\| \geq 4$
LQR	28.6	15%	10
indirect	28.7	16%	11
direct	30.0	0%	183

indirect feedback implementation. It is argued and numerically demonstrated, that a direct feedback formulation, which aims to satisfy constraints at each time step with regard to the predicted distribution given the currently measured state, leads to significant conservatism with respect to constraint satisfaction and reduced performance. The indirect feedback formulation, on the other hand, considers a closed-loop error distribution with respect to a nominal state, and enables a non-conservative treatment of constraints and the recovery of the true optimal solution.

REFERENCES

Bertsekas, D. P., 2017. Dynamic programming and optimal control. Vol. 1. Athena Scientific Belmont, MA.

Cannon, M., Kouvaritakis, B., Raković, S. V., Cheng, Q., 2011. Stochastic tubes in model predictive control with probabilistic constraints. Trans. Autom. Control 56 (1), 194–200.

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.

Farina, M., Giulioni, L., Scattolini, R., 2016. Stochastic linear Model Predictive Control with chance constraints - A review. Journal of Process Control 44, 53–67.

Hewing, L., Wabersich, K. P., Zeilinger, M. N., 2018. Recursively feasible stochastic model predictive control using indirect feedback. arXiv:1812.06860.

Hewing, L., Zeilinger, M. N., 2018. Stochastic model predictive control for linear systems using probabilistic reachable sets. In: Conf. Decision Control.

Hewing, L., Zeilinger, M. N., April 2020. Scenario-based probabilistic reachable sets for recursively feasible stochastic model predictive control. IEEE Control Systems Letters 4 (2), 450–455.

Korda, M., Gondhalekar, R., Cigler, J., Oldewurtel, F., 2011. Strongly feasible stochastic model predictive control. Conf. Decision Control and Europ. Control Conf.

Korda, M., Gondhalekar, R., Oldewurtel, F., Jones, C., 2014. Stochastic MPC framework for controlling the average constraint violation. Trans. Autom. Control 59 (7), 16. 1706–1721.

Kouvaritakis, B., Cannon, M., 2016. Model Predictive Control: Classical, Robust and Stochastic. Springer.

Kouvaritakis, B., Cannon, M., Raković, S. V., Cheng, Q., 2010. Explicit use of probabilistic distributions in linear predictive control. Automatica 46 (10), 1719–1724.

Lorenzen, M., Dabbene, F., Tempo, R., Allgöwer, F., 2017. Constraint-tightening and stability in stochastic model predictive control. Trans. Autom. Control 62 (7), 3165–3177.