FEEDBACK MODEL PREDICTIVE CONTROL BY RANDOMIZED ALGORITHMS

Ivo Batina* **Anton A. Stoorvogel*†** **Siep Weiland‡**

Keywords: predictive control, Monte-Carlo methods, stochastic systems, input constraints

Abstract

In this paper we present a further development of an algorithm for stochastic disturbance rejection in model predictive control with input constraints based on randomized algorithms. The algorithm presented in [1] can solve the problem of stochastic disturbance rejection approximately but with high accuracy at the expense of a large computational effort. The algorithm described here uses a predefined controller structure in the optimization and it is significantly less computationally demanding but with a price of some performance loss. Via an example it is shown that the algorithm gives considerable reduction in the computational time and that performance loss is rather small compared to the algorithm in [1].

1 Introduction

Model Predictive Control (MPC) has gained a wide acceptance in industry and a lot of attention in the academic community worldwide [3]. The common feature of these control techniques is a direct use of a model for the prediction of the controlled plant behavior. At each time instant a finite horizon optimal control problem (possibly subject to some input and state constrains) is solved, taking the current state of the plant as the initial condition. Only the first control move in the computed sequence is applied to the plant and the optimization is repeated in the next time instant.

However, when the optimization ignores the effects of possible future changes in the disturbance and/or a model mismatch, closed-loop performance can be poor.

In the case of stochastic disturbances it seems that minimization of the expected cost is a natural way of dealing with disturbance rejection, analogous to the well known LQ theory. As in the deterministic case, constraints make an analytical solution to the optimization problem untractable. In the case that an analytic solution is not possible and that standard computational methods are too complex, randomized algorithms have been applied in control theory mostly in connection with robustness but recently also in connection with input constraints (see [4]).

In [1] we presented a disturbance rejection scheme for MPC based on randomized algorithms which minimizes an empirical mean of the cost function. The optimization at each step is based on closed loop optimization, therefore it takes into account the effect of disturbances. Because we do not perform any a priori parameterization of the feedback laws over the control horizon, the algorithm is computationally demanding but it gives a reliable measure of the achievable performance.

In this paper, we present a new algorithm for disturbance rejection in MPC. In the algorithm we propose here, minimization is over a class of saturated feedback controllers. A significant reduction in the computational effort is achieved by predefining a controller structure in the closed loop optimization. The result is an algorithm that is computationally less demanding than the one in [1] at the expense of some performance loss. Via an example, we show that the disturbance rejection performance of our algorithm is still significantly better when compared with an algorithm based on obtained standard MPC scheme and that is very close to the results of the algorithm in [1].

The paper is organized as follows. Overview of the available approaches to the disturbance rejection in MPC is given in section 2. Predictive control based on randomized algorithm is presented in section 3. Finally, a numerical example is presented in section 4.

2 Model predictive control and disturbances

Consider a linear time-invariant plant represented with the following state space model:

$$\begin{aligned} x(k+1) &= Ax(k) + Bu(k) + Gw(k) \\ z(k) &= C_z x(k) + D_z u(k) \end{aligned} \tag{1}$$

where $u(k) \in \mathbf{U} \subset \mathbb{R}^m$ is the control input and $\mathbf{U}$ is a compact, convex set which contains an open neighborhood of the origin, $z(k) \in \mathbb{R}^p$ is the controlled output and $x(k) \in \mathbb{R}^n$ is the state. Finally, $w(k) \in \mathbf{W} \subseteq \mathbb{R}^w$ is a disturbance. Variable k is supposed to range over the nonnegative integers $\mathbb{Z}_+$.

It assumed that there is no mismatch between the plant and the model i.e. that the model describes the plant completely. The state of the plant is measured.

We assume that the plant, which is subject to amplitude constraints on the input, satisfies condition for the global asymptotic stabilization via feedback (see [5]), i.e. that all eigenvalues of the system matrix A lie on or inside the unit circle.

Let $t \in \mathbb{Z}_+$ be a fixed time instant representing the current time.

*Department of Mathematics and Computer Science, Eindhoven Univ. of Technology, P.O. Box 513, 5600 MB Eindhoven, E-mail: {i.batina,a.a.stoorvogel}@tue.nl, Tel: +31.40.247{4277,2378}, Fax: +31.40.2442489.

†Department of Information Technology and Systems, Delft Univ. of Technology, P.O. Box 5031, 2600 GA Delft, The Netherlands, E-Mail: A.A.Stoorvogel@its.tudelft.nl, Tel: +31.15.2781912, Fax: +31.15.2787209.

‡Department of Electrical Engineering, Eindhoven Univ. of Technology, P.O. Box 513, 5600 MB Eindhoven, E-Mail: S.Weiland@tue.nl, Tel: +31.40.2473577, Fax: +31.40.2434582.

Consider the control horizon $I_t := \{t + k | k \in T\}$, $T := [0, N]$ with length $N > 0$. The control horizon is defined as time dependent interval, it recedes with the time.

The main purpose of the control horizon is to provide a receding time slot over which an optimization, based on the predicted behavior of the plant, is performed. A prediction is obtained by performing the recursion defined by the model equations (1) for $N + 1$ steps with the measured state of the plant, denoted as x_t, as an initial condition.

In the standard MPC setting one defines a control and a disturbance sequence over the control horizon. The (1) is time invariant and the current time can be set to 0 without loss of generality. We denote a control sequence as $u : T \to \mathbf{U}$ and a disturbance sequence as $w : T \to \mathbf{W}$. With x_t as an initial condition i.e. $x(0) := x_t$, the state and the controlled output of the model (1) subject to the control u and the disturbance w are denoted as $x : T \to \mathbb{R}^n$ and $z : T \to \mathbb{R}^p$ respectively.

The cost function commonly considered in MPC is then given by:

$$J^{op}(x_t, u, w) := \sum_{k \in T} \|z(k)\|^2 + x(N+1)' Q x(N+1) \quad (2)$$

where $(\cdot)'$ denotes a matrix transpose and $Q = Q' \geq 0$ denotes an end-point penalty. The end-point penalty is described by a symmetric matrix $Q \in \mathbb{R}^{n \times n}$. It is known that if the end-point penalty is sufficiently large, then the closed loop system will be stable. For further details about the choice of Q we refer to [7].

The standard MPC optimization with the model (1) amounts to:

$$V(x_t, w) := \inf_u \{J^{op}(x_t, u, w)\}. \quad (3)$$

The optimization (3) is only well posed if the initial state and the disturbance sequence w are given. In the literature two approaches have been suggested to deal with with the problem of an unknown disturbance.

- Assume that the disturbance is either zero or a known constant over the optimization interval.
- Use a worst case approach, [6] where we minimize over u and maximize over w. In other words consider the following optimization problem:

  $$\min_u \max_w J^{op}(x_t) \quad (4)$$

  where w is restricted to take values in some bounded set $\mathbf{W}$.

The first approach is in a sense too optimistic and ignores the effect the disturbance can have on the system. The second approach is too pessimistic and therefore the results are often very conservative. A third approach is to optimize:

$$\inf_u \mathbb{E} J^{op}(x_t)$$

where $\mathbb{E}$ denotes the expectation and the disturbance w is assumed to be a white noise stochastic process w taking values at each time t in the set $\mathbf{W}$ with some probability measure. Here u is optimized in open loop over the control horizon. It is then not difficult to show that the resulting input u is the same when we consider the stochastic disturbance w or the deterministic and constant disturbance $\mathbb{E}w$ (i.e. we replace w by its expectation). Hence this approach is basically the same as assuming the disturbance to be constant over the control horizon.

In the paper [1] we took a fourth approach. We choose u at some instant k, $k \in T$ in the control horizon as a function of $x(k)$ and, possibly, earlier states. Equivalently, we can choose $u(k)$ as a function of earlier disturbances.

Formally, we define a feedback control law at k, $k \in T$ as a mapping $\pi_k : \mathbb{R}^n \to \mathbb{U}$. On the control horizon we have a sequence of mappings $\pi := \{\pi_k\}_{k=0}^N$.

The cost that we consider is given as:

$$J(x_t, \pi, w) := \sum_{k \in T} \|C_z x(k) + D_z \pi_k(x(k))\|^2 + \|x(N+1)\|_Q^2 \quad (5)$$

where $\|x(N+1)\|_Q^2 := \langle x, Qx \rangle$ and $Q = Q' \in \mathbb{R}^{n \times n}$ is an end-point penalty.

The optimization problem to be solved is then given by:

$$V(x_t) := \inf_\pi \mathbb{E}_w J(x_t, \pi) \quad (6)$$

where $\mathbb{E}_{(\cdot)}$ denotes the conditional expectation with respect to $(\cdot)$.

The control obtained by solving (6) is implemented according to the receding control paradigm. Suppose that π^* is the sequence of feedback laws over the control horizon that minimizes $\mathbb{E}_w J(x_t, \pi)$. Only the first element of π^* is significant in the receding horizon implementation. It determines the current input for the plant as a function of the current measurement. In the next time instant, the control horizon is shifted forward and optimization problem (6) is solved for the new state measurement.

Note that (5) and (6) are time invariant. The receding horizon controller in the setting described above, is given as:

$$u(t) = \pi_0^*(x_t) \quad \forall\ t \in \mathbb{Z}_+ \quad (7)$$

An analytic solution of the optimization problem (6) is very difficult to obtain. Instead of finding an optimal input in $\mathbb{R}^z$ for some z (a finite-dimensional space) we need to find optimal inputs as functions from $\mathbb{R}^q$ to $\mathbb{R}^s$ (an infinite-dimensional space).

In the paper [1], we find an approximate solution of this problem by using randomized algorithms. Although computationally very intensive we can approximately compute the optimally achievable performance with high accuracy.

In this paper, we restrict the class of possible feedback mappings to one of the form $u(k) = \sigma(Fx(k))$, $k \in T$ i.e. a linear

feedback type with saturation. The reduction of the computational load is achieved at the expense of some performance loss. The algorithm from [1], in this paper named Algorithm 1, can be used to capture the loss of the performance and to characterize the trade-off between accuracy and computational effort.

3 Model predictive control by randomized algorithms

If one wants to compute an expectation as in the optimization criterion (6) for a given input u one can do a computation based on the distribution of the stochastic disturbance. An analytical computation of that kind for the case that we consider is extremely difficult. An alternative methodology is to compute the empirical mean instead of the conditional expectation in (6). For computation of the empirical mean we need the cost for a number of realizations of the stochastic disturbance w. The realizations are chosen randomly, according to the distribution of w. Because of that, an algorithm in which the empirical mean is used is called a randomized algorithm.

First, we formally define the empirical mean and an important related result the so-called Hoeffding's inequality.

3.1 Empirical mean

Assume a set Θ and a probability measure P on Θ are given. Let f be a scalar-valued function measurable with respect to P, defined on Θ

$$f : \Theta \to \mathbb{R} \tag{8}$$

The expectation or mean of f over Θ can be expressed in terms of the following integral:

$$\mathbb{E}(f) = \int_{\Theta} f(\theta) dP \tag{9}$$

The aim is to approximate (9) by drawing m independent, identically distributed (i.i.d) samples $\vartheta = \{\theta_1, \cdots, \theta_m\}$ from Θ according to the probability measure P. This defines the *empirical mean*:

$$\hat{\mathbb{E}}(f, \vartheta) := \frac{1}{M} \sum_{j=1}^{M} f(\theta_j)$$

We say that we will have confidence δ in this approximation if the probability that the empirical mean differs from the true expectation by more than ε is less than δ. Equivalently, we have confidence δ in our approximation if the probability that the empirical mean differs less than ε from the true expectation is larger than $1 - \delta$.

An upper bound for the confidence δ is given by Hoeffding's inequality [8]. For all $\varepsilon > 0$ we have:

$$\text{Prob}\left(|\hat{\mathbb{E}}(s, \vartheta) - \mathbb{E}(f)| > \varepsilon\right) \leq 2e^{-2m\varepsilon^2}$$

where Prob indicates the probability. In other words, the confidence δ is larger than $1 - 2e^{-2m\varepsilon^2}$.

If we choose the number of samples m to satisfy the following inequality

$$m \geq \frac{1}{2\varepsilon^2} \ln \frac{2}{1 - \delta}$$

then Hoeffding's inequality tells us that the empirical mean is within ε of the true mean with confidence δ. An important comment to be made here is that the number of samples needed in order to get an estimate of the true mean with high confidence is independent of the dimension of the underlying stochastic process, i.e. the dimension of the set Θ.

3.2 Algorithm 1

At some instant s, $s \in T$ the state $x(s)$ is a stochastic variable. However, (1) is a causal system and $x(s)$ does not depend on the "future" disturbance $w_s : [s, N] \to \mathbf{W}$. That allows to express an optimal cost "to go" at each s, $s \in T$ as:

$$V_s(x(s)) = \inf_{\pi^s} \{ \mathbb{E}_{w_s} J_s(x(s), \pi^s) \}. \tag{10}$$

where for all $s \in T$ we define:

$$J_s(x(s), \pi^s, w_s) = \sum_{k=s}^{N} \|y(k)\|^2 + x(N+1)' Q x(N+1)$$

with the disturbance $w_s : [s, N] \to \mathbf{W}$ and $\pi^s := \{\pi_k\}_{k=s}^{N}$ as a sequence of feedback mappings $\pi_k : \mathbb{R}^n \to \mathbf{U}$.

The method proposed in [1] amounts to infimizing the empirical mean instead of the conditional expectation in (10). The empirical mean is computed for a number of the disturbance samples.

Suppose that we take κ samples of the disturbance $w(0)$ at $k = 0$. With that, there are κ possible states $x(1)$ for the initial condition x_t and the input $u(0)$. For each one of these possible futures we generate κ samples of the disturbance $w(1)$ which establishes κ^2 possible future states $x(2)$. On this way, by the persistent sampling of the disturbance up to $N - 1$ the number of samples of w is κ^N. The number of samples of the restricted disturbance sequence w_s is κ^{N-s}. The number of samples of w grows exponentially with the horizon.

For all $s \in [0, N]$ and for each of κ^{N-s} samples of w_s denoted as w_s^i, $i \in [1, \kappa^{N-s}]$ we write the cost function as:

$$J^i(x(s), \pi^s, w_s^i) = \sum_{k=s}^{N} \|C_z x(k) + D_z \pi_k^s(x(k))\|^2 + \\ + \|x(N+1)\|_Q^2.$$

The empirical conditional mean of the cost function in $x(s)$ given a restricted sequence of feedback mappings π^s is written as:

$$\hat{\mathbb{E}}_{w_s} J(x(s), \pi^s) := \frac{1}{\kappa^{N-s}} \sum_{i=1}^{\kappa^{N-s}} J^i(x(s), \pi^s, w_s^i) \tag{11}$$

and an empirical optimal cost "to go" as:

$$\hat{V}(x(s)) := \inf_{\pi^s} \hat{\mathbb{E}}_{w_s} J(x(s), \pi^s). \tag{12}$$

The following result states an important property of (12).

Lemma 1 *The empirical optimal cost (12) is a strictly convex function in $x(s)$ for all s.*

Proof: In [1].

The Algorithm 1 is based on a dynamic program:

$$\hat{V}_s(x(s)) := \min_{u\in\mathbf{U}} \left\{ \|C_z x(s) + D_z u\|^2 + \right. \\ \left. + \hat{\mathbb{E}}_{w(s)} \hat{V}_{s+1}(x(s+1)) \right\} \tag{13}$$

with a terminal condition:

$$\hat{V}_{N+1}(x(N+1)) := \|x(N+1)\|_Q^2$$

and an initial condition:

$$x(0) := x_t.$$

The dynamic program (13) has to be solved backwards i.e. from $k = N$ to $k = 0$.

At some $s \in T$ the number of points for which we evaluate $\hat{V}(x(s))$ is determined by all past disturbance realization $w(\tau)$, $\tau \in [0, s-1]$. With disturbances sampled as described previously, the number of points in the state space in which we evaluate the empirical cost "to go" (13) is equal to κ^s. This yields an exponential growth on the horizon which makes this algorithm computationally very intensive but the problem can be solved with an arbitrary accuracy.

3.3 Algorithm 2

A natural remedy is to sacrifice performance for the benefit of an increased computational speed. This can be done by imposing a feedback structure on the control horizon instead of an optimization over a general feedback mapping.

Consider the structure of the form

$$u(k) = \sigma\,(Fx(k)) \quad k \in [0, N] \tag{14}$$

where σ is a saturation function that achieves that $\sigma(u) \in \mathbf{U}$ for all $u \in \mathbb{R}^m$ and F is a linear feedback control law $F : \mathbb{R}^n \to \mathbb{R}^m$.

The cost function at s, for all $s \in T$ becomes:

$$J_s^F(x(s), F, w_s) = \sum_{k=s}^{N} \|C_z x(k) + D_z \sigma\,(Fx(k))\,\|^2 + \\ + \|x(N+1)\|_Q$$

and an optimal cost:

$$\hat{V}^F(x_t)) := \inf_F \{\hat{\mathbb{E}}_{w_s} J_0^F(x_t, F)\} \tag{15}$$

Algorithm 2

Step 1 Draw κ samples for w according to the distribution of w. Set $V_0 = \infty$. Set accuracy parameter ϵ. Set $F = F_{LQ}$ where F_{LQ} is the solution of the unconstrained infinite horizon LQ problem for the system (1):

$$F_{LQ} = -(D_z^T D_z + B^T P B)^{-1} B^T P A$$

where $P = P^T \geq 0$ is the solution of:

$$P = A^T P A + C_z^T C_z - (A^T P B + C_z^T D) \\ \times (B^T P B + D_z^T D_z)^{-1} (B^T P A + D_z^T C_z)$$

Step 2 Compute

$$\hat{V}_s^F := \hat{\mathbb{E}}_{w_s} J_s^F(x(s), F)$$

for all $s \in T$.

Step 3 If $\|\hat{V}_0^F - V_0\| < \epsilon$ stop.
Otherwise: set $V_0 = \hat{V}_0^F$ and update F according to the numerical algorithm that has been chosen for the numerical minimization of (15).

The input of the plant at some time $t \in \mathbb{Z}_+$ is then computed according to (14) and in the next time instant computations in Algorithm 2 are repeated. It follows from Hoeffding's inequality that the required increase in the number of samples which are needed to preserve a certain confidence with a growing control horizon is polynomial instead of exponential. However, the algorithm will converge to a suboptimal solution of the optimization problem. In this case, the accuracy is relative to the suboptimal solution and there is a performance loss that is inherent to the procedure because we restricted our attention to only one parameterization of all possible parameterizations of feedback controllers.

We would like to stress the fact that Algorithm 1 computes the optimum with high accuracy. Therefore, we can evaluate simplifications of this scheme which reduce the computational time and show to what extend performance has been compromised.

The parameterization of control laws which we suggest in this paper is a simple one. If performance loss in this scheme turns out to be too large for a specific application, we may consider more elaborate parameterizations of control laws such as an optimization over the sequence of saturated linear feedback gains over the control horizon.

4 Numerical example

In this section we illustrate the approach of this paper by means of a numerical example. We use the Algorithm 1 as a tool to compare performance between the standard MPC scheme and the feedback MPC based on the Algorithm 2. Consider the

system of the form (1) where

$$A = \begin{bmatrix} 0.7326 & -0.0861 \\ 0.1722 & 0.9909 \end{bmatrix} \quad B = \begin{bmatrix} 0.0609 \\ 0.0064 \end{bmatrix}$$

$$G = \begin{bmatrix} 0.5 \\ 0 \end{bmatrix} \quad C = \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \quad D = \begin{bmatrix} 0.1 \\ 0 \\ 0 \end{bmatrix}$$

See [1] and [2]. For each time instant t, the stochastic disturbance $w(t)$ is assumed to be uniformly distributed on the interval $[-\alpha,\ \alpha]$ where α varies in various experiments as $\alpha = 0.5$, $\alpha = 1$, $\alpha = 1.5$, and $\alpha = 2$.

The aim is to regulate the system towards the origin (disturbance rejection) while inputs are constrained to be

$$-2 \leq u(t) \leq 2$$

for all t. As an indication of the achieved level of disturbance rejection we consider the variance of the system state. The control horizon $N = 2$ and simulations are performed over a 200 sec. time interval.

We compare simulation results for the standard MPC design, based on the open loop optimization of the form (3), the design based on the Algorithm 1 and the Algorithm 2 proposed in this paper. The results are summarized in Table 1 and 2.

	α	0.5	1
Algorithm 1	var(x)	0.0242	0.1178
Standard MPC scheme	var(x)	0.0243	0.1310
Performance loss		0.5%	11.2%
Algorithm 2	var(x)	0.0243	0.1179
Performance loss		0.5%	0.1%

Table 1: Variance of system state, $\alpha = \{0.5, 1\}$

	α	1.5	2
Algorithm 1	var(x)	0.3082	0.5932
Standard MPC scheme	var(x)	0.3234	0.6020
Performance loss		4.9 %	1.4 %
Algorithm 2	var(x)	0.3095	0.5920
Performance loss		0.5 %	-0.2 %

Table 2: Variance of system state, $\alpha = \{1.5, 2\}$

The number of disturbance samples κ has been set to 20 for both algorithms. The simulations reported here have been obtained on the PC (450 MHz processor) running Matlab and Simulink. The simulations are performed for 200 time steps. Note that the speed by which computations are performed critically depends on the simulation software used. The average computational time for the Algorithm 1 is reduced by factor 8 compared to the Algorithm 2, for the system of the same order and the same number of disturbance samples. On the other hand, the performance losses of the Algorithm 2 are significantly smaller then the losses of the standard MPC scheme. When the disturbance level is small or very high all algorithms yield approximately the same performance. In the first case that is because the disturbances do not influence the system significantly and in the later case, because of the input constraints, the performance is the same as for the open-loop control.

5 Conclusion

We described an algorithm to design model predictive controllers which takes into account stochastic disturbances and constraints on input. The algorithm uses a predefined controller structure in the optimization and it is significantly less computationally demanding than algorithm in [1] but with a price of some performance loss. It is shown, by an example, that this simplification does reduce computational time needed by the expense of marginal reduction in performance. Other aspects currently under investigation are issues of state constraints and measurement feedback.

References

[1] I. Batina, A. A. Stoorvogel, and S. Weiland. Stochastic disturbance rejection in model predictive control by randomized algorithms. *In proceedings of American Control Conference*, 2001.

[2] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulus. The explicit solution of model predictive control via multiparametric quadratic programing. *In proceedings of American Control Conference*, pages 872–876, 2000.

[3] E. F. Camacho and C. Bordons. *Model Predictive Control*. Springer-Verlag, London, 1998.

[4] X. Chen and K. Zhou. Constrained robustness analysis by randomized algorithms. *IEEE Transactions on Automatic Control*, 45(6):1180–1186, 2000.

[5] A. Saberi, A. A. Stoorvogel, and P. Sannuti. *Control of linear systems with regulation and input constraints*. Springer-Verlag, London, 2000.

[6] P. O. M. Scokaert and D. Q. Mayne. Min-max feedback model predictive control for constrained linear systems. *IEEE Transactions on Automatic Control*, 43(8):1136–1142, 1998.

[7] A. A. Stoorvogel, S. Weiland, and J. H. A. Ludlage. On aysomptic stability of a receiding horizon control scheme. *In proceedings of the* 38^{th} *Conference on decision and Control*, pages 3684–3689, 1999.

[8] M. Vidyasagar. *A theory of learning and generalization*. Springer-Verlag, London, 1997.