On the Convergence of Stochastic MPC to Terminal Modes of Operation

Diego Muñoz-Carpintero
Universidad de Chile

Mark Cannon
University of Oxford

26 June 2019

Universidad de Chile

Introduction to Stochastic MPC

System	$x_{k+1} = Ax_k + Bu_k + w_k$
Constraints	$\mathbb{P}\{(x,u) \in \mathcal{Y}\} \geq p$
	$\mathcal{Y} = \{(x,u) : Fx + Gu \leq 1\}$
Stochastic disturbance	$w \in \mathcal{W}$

Stochastic MPC

At $k = 0, 1, \ldots$:

- obtain x_k and optimize $\{u_k(x_k), \ldots, u_{k+N-1}(x_{k+N-1})\}$:

$$\begin{aligned} \min_{u_k(\cdot),\ldots,u_{k+N-1}(\cdot)} \quad & \mathbb{E}\Big\{F(x_{k+N}) + \sum_{j=0}^{N-1} \ell(x_{k+j}, u_{k+j})\Big\} \\ \text{s.t.} \quad & x_{k+j+1} = Ax_{k+j} + Bu_{k+j} + w_{k+j} \\ & \mathbb{P}\{(x_{k+j}, u_{k+j}) \in \mathcal{Y}\} \geq p \\ & x_{k+N} \in \tilde{\mathbb{X}}_f \end{aligned}$$

- apply first element of optimal sequence: $u_k = u_k^*(x_k)$

Stability and convergence results for Stochastic MPC

1. Negative Drift Conditions

There exist measurable functions $V : \mathbb{R}^{n_x} \to [0, \infty)$, $\Psi : \mathbb{R}^{n_x} \to [0, \infty)$ and a bounded and measurable set $\mathcal{Z} \subset \mathbb{R}^{n_x}$, such that

$$\mathbb{E}\{V(x_1) \,|\, x_0 = x\} - V(x) \leq -\Psi(x) \quad \forall x \notin \mathcal{Z}$$

- This implies boundedness of $\mathbb{E}\{V(x_k) \,|\, x_0 = x\}_{k \in \mathbb{N}}$
- Stochastic MPC typically ensures a drift condition, e.g.:

  $$\mathbb{E}\{V(x_1) \,|\, x_0 = x\} - V(x) \leq -(1 - \lambda)V(x) \quad \forall x \notin \mathcal{Z}$$

  for $\lambda \in (0, 1)$, for some $\mathcal{Z}$

  But this doesn't give non-conservative ultimate bounds on x or a probabilistic description of the terminal regime

Stability and convergence results for Stochastic MPC

2. Convergence of average performance

$$\lim_{k\to\infty} \frac{1}{k} \sum_{j=1}^{k} \mathbb{E}\{x_j^\top Q x_j + u_j^\top R u_j\} \le L_{ss}$$

- An alternative notion of convergence of Stochastic MPC via asymptotic average performance
- Results in conservative bounds except in special cases (e.g. if u_k converges to certainty equivalent optimal feedback)

Stability and convergence results for Stochastic MPC

3. Input to state stability (ISS)

The origin of $x_{k+1} = f(x_k, w_k)$ is ISS with region of attraction $\mathbf{X} \subseteq \mathbb{R}^{n_x}$ if $x_k \in \mathbf{X}$ for all k, all $x_0 \in \mathbf{X}$ and all $w \in \mathcal{W}$, and

$$\|x_k\| \leq \beta(\|x_0\|, k) + \gamma(\sup_{t<k}\{\|w_t\|\})$$

where β is a $\mathcal{KL}$-function and γ is a $\mathcal{K}$-function

Lemma: ISS [Jiang & Wang, 2001]

The origin of $x_{k+1} = f(x_k, w_k)$ is ISS with region of attraction $\mathbf{X} \subseteq \mathbb{R}^{n_x}$ if $\mathbf{X}$ contains the origin in its interior and is robustly invariant, and a continuous function $V : \mathbf{X} \to \mathbb{R}_+$ (called an ISS-Lyapunov function) exists satisfying, for all $x \in \mathbf{X}$ and $w \in \mathcal{W}$,

$$\alpha_1(\|x\|) \leq V(x) \leq \alpha_2(\|x\|)$$

$$V\big(f(x, w)\big) - V(x) \leq -\alpha_3(\|x\|) + \sigma(\|w\|)$$

where $\alpha_1, \alpha_2, \alpha_3$ are $\mathcal{K}_\infty$-functions and σ is a $\mathcal{K}$-function

Stability and convergence results for Stochastic MPC

3. Input to state stability (ISS)

The origin of $x_{k+1} = f(x_k, w_k)$ is ISS with region of attraction $\mathbf{X} \subseteq \mathbb{R}^{n_x}$ if $x_k \in \mathbf{X}$ for all k, all $x_0 \in \mathbf{X}$ and all $w \in \mathcal{W}$, and

$$\|x_k\| \leq \beta(\|x_0\|, k) + \gamma(\sup_{t<k}\{\|w_t\|\})$$

where β is a $\mathcal{KL}$-function and γ is a $\mathcal{K}$-function

Lemma: ISS [Jiang & Wang, 2001]

The origin of $x_{k+1} = f(x_k, w_k)$ is ISS with region of attraction $\mathbf{X} \subseteq \mathbb{R}^{n_x}$ if $\mathbf{X}$ contains the origin in its interior and is robustly invariant, and a continuous function $V : \mathbf{X} \to \mathbb{R}_+$ (called an ISS-Lyapunov function) exists satisfying, for all $x \in \mathbf{X}$ and $w \in \mathcal{W}$,

$$\alpha_1(\|x\|) \leq V(x) \leq \alpha_2(\|x\|)$$

$$V\big(f(x, w)\big) - V(x) \leq -\alpha_3(\|x\|) + \sigma(\|w\|)$$

where $\alpha_1, \alpha_2, \alpha_3$ are $\mathcal{K}_\infty$-functions and σ is a $\mathcal{K}$-function

Stability and convergence results for Stochastic MPC

ISS implies:

- the origin is asymptotically stable for $x_{k+1} = f(x_k, 0)$
- all state trajectories are bounded since $\mathcal{W}$ is bounded
- all trajectories converge to the origin as $k \to \infty$ if $w_k \to 0$

. . . but it doesn't provide

- non-conservative ultimate bounds on x
- a probabilistic description of the terminal regime

Stability and convergence results for Stochastic MPC

Observation: many Stochastic MPC analyses give qualitative stability/convergence results but do not characterize asymptotic behaviour exactly

Goal: general conditions characterizing exact asymptotic behavior under Stochastic MPC

Tools:
(i) results on convergence of Markov chains
(ii) ISS properties of controlled systems

Convergence for ISS systems

Definition (Markov chain)

Consider a measurable space $(\mathbf{X}, \mathcal{B}(\mathbf{X}))$ *and a stochastic process* $\mathbf{x} := \{x_k \in \mathbf{X}\}_{k\in\mathbb{N}}$ *defined on* $(\Omega, \mathcal{F})$, *where* $\mathcal{F}$ *is a* σ*-algebra on* $\Omega := \Pi_{i=0}^{\infty}\mathbf{X}_i$, *and* $\mathbf{X}_i = \mathbf{X}$ *for all* i. *Then* $\mathbf{x}$ *is a time-homogenous Markov chain with transition probability function*

$$P(x, \mathcal{A}) := \mathbb{P}\{x_{k+1} \in \mathcal{A} : x_k = x\}$$

if the distribution of $\mathbf{x}$ *satisfies the Markov property*

$$\mathbb{P}\{x_{k+1} \in \mathcal{A} : x_j = \bar{x}_j,\ j \in \mathbb{N}_k\} = P(\bar{x}_k, \mathcal{A})$$

Definition (Invariant measure)

For the Markov chain $\mathbf{x}$ *an invariant probability measure is a stationary distribution, i.e. a probability measure* π *satisfying*

$$\pi(\mathcal{A}) = \int \pi(dx)P(x, \mathcal{A}), \quad \forall \mathcal{A} \in \mathcal{B}(\mathbf{X})$$

Convergence for ISS systems

Markov chain convergence results [e.g. Meyn and Tweedie, 2005]:

Let $\mathbf{x}$ be a φ-irreducible Markov chain with state space $\mathbf{X} \subseteq \mathbb{R}^{n_x}$ such that

(i) $\mathbf{x}$ is generated by $x_{k+1} = f(x_k, w_k)$, for some continuous $f : \mathbf{X} \times \mathcal{W} \to \mathbf{X}$ and a stochastic disturbance $\{w_k \in \mathcal{W}\}_{k \in \mathbb{N}}$

(ii) $\mathbf{x}$ is aperiodic

(iii) $\mathrm{supp}(\varphi)$ has non-empty interior

(iv) there is a measurable function $V : \mathbf{X} \to [0, \infty)$ such that for any $c < \infty$ the set $\mathcal{C}_V(c) \coloneqq \{y : V(y) \leq c\}$ is compact, and there is a compact set $\mathcal{C}$ satisfying for all $x_k \in \mathbf{X}$:

$$\mathbb{E}\{V(x_{k+1})\} - V(x_k) \leq -1 + b\mathbf{1}_{\mathcal{C}}(x_k)$$

Convergence for ISS systems

Markov chain convergence results [e.g. Meyn and Tweedie, 2005]:

then

Theorem (Markov chain convergence)

An invariant probability measure $\pi(\cdot)$ *exists satisfying*

$$\lim_{k\to\infty} \sup_{\mathcal{A}\in\mathcal{B}(X)} |P^k(x,\mathcal{A}) - \pi(\mathcal{A})| = 0$$

where $P^k(x,\mathcal{A}) := \mathbb{P}\{x_k \in \mathcal{A} : x_0 = x\}$, *and the Law of Large Numbers:*

$$\lim_{k\to\infty} \frac{1}{k}\sum_{j=1}^{k} h(x_j) \stackrel{a.s.}{=} \mathbb{E}_\pi\{h(x)\}$$

holds for any $h : \mathbf{X} \to \mathbb{R}$ *such that* $\mathbb{E}_\pi\{|h(x)|\} < \infty$

where $\mathbb{E}_\pi\{h(x)\} := \int \pi(dx)h(x)$

Convergence for ISS systems

We apply these results to systems of the form

$$x_{k+1} = f(x_k, w_k) \coloneqq g(x_k) + Dw_k,$$

with $x_k \in \mathbf{X}$, $w_k \in \mathcal{W}$ and $g : \mathbf{X} \to \mathbf{X}$ continuous with $g(0) = 0$

Assumption 1. (Disturbance distribution)

The disturbance sequence $\{w_k \in \mathcal{W}\}_{k \in \mathbb{N}}$ is i.i.d., with $\mathbb{E}\{w_k\} = 0$ and a non-singular probability distribution such that

$$\mathbb{P}\{\|w\| \leq \lambda\} > 0 \quad \forall \lambda > 0$$

Convergence for ISS systems

Suppose there is a linear terminal mode of operation to which we want to prove convergence

Assumption 2. (Linear terminal mode)

There exists a bounded set $\mathbf{X}_f \subseteq \mathbf{X}$ containing the origin in its interior, such that for all $x \in \mathbf{X}_f$,

(i) $f(x, w) \in \mathbf{X}_f$ for all $w \in \mathcal{W}$

(ii) $f(x, w) = Ax + Dw$ for all $x \in \mathbf{X}_f$, where A is Schur stable and (A, D) is controllable

Then the linear terminal dynamics define a transition probability function $P(x, \cdot)$ and an invariant probability measure $\pi(\cdot)$, where

- π is the probability measure of $\sum_{k=0}^{\infty} A^k D w_k$
- the support of π is the minimal invariant set $\mathbf{X}_\infty = \bigoplus_{k=0}^{\infty} A^k D \mathcal{W}$

Convergence for ISS systems

Assumption 3. (ISS)

The system $x_{k+1} = g(x_k) + Dw_k$ has an ISS-Lyapunov function

- Clearly this implies that the origin is ISS, but it does not directly guarantee convergence to the terminal mode of operation
- The ISS property can be coupled with the stochastic nature of the disturbance sequence to prove convergence to $\mathbf{X}_f$

Convergence for ISS systems: main result

Under Assumptions 1-3, the Markov chain convergence results imply:

Theorem

The system $x_{k+1} = g(x_k) + Dw_k$ *satisfies*

$$\lim_{k\to\infty} \sup_{\mathcal{A}\in\mathcal{B}(X)} |P^k(x,\mathcal{A}) - \pi(\mathcal{A})| = 0$$

where $\pi(\cdot)$ *is the invariant probability measure associated with the terminal linear dynamics and* $P^k(x,\mathcal{A}) := \mathbb{P}\{x_k \in \mathcal{A} : x_0 = x\}$, *and*

$$\lim_{k\to\infty} \frac{1}{k} \sum_{j=1}^{k} h(x_j) \overset{a.s.}{=} \mathbb{E}_\pi\{h(x)\}$$

holds for any $h : \mathbf{X} \to \mathbb{R}$ *such that* $\mathbb{E}_\pi\{|h(x)|\} < \infty$

where $\mathbb{E}_\pi\big\{h(x)\big\} := \int \pi(dx)h(x)$

Convergence for ISS systems: main result

This implies convergence to the minimal invariant set $\mathbf{X}_\infty = \bigoplus_{k=0}^{\infty} A^k D\mathcal{W}$

Corollary

The system $x_{k+1} = g(x_k) + Dw_k$ *satisfies*

$$\lim_{k\to\infty} \mathbb{P}\{x_k \in \mathbf{X}_\infty\} = 1.$$

Convergence for Stochastic MPC

Interpretation

- the system converges to the minimal Robust Positively Invariant (mRPI) set
- average performance converges to that of the terminal linear mode

These results can be applied to many Stochastic MPC algorithms
We consider two formulations:

1. Affine in the disturbance SMPC
   P. Goulart and E. Kerrigan, *Input-to-state stability of robust receding horizon control with an expected value cost*, Automatica, 2008
2. Striped affine in the disturbance SMPC
   B. Kouvaritakis, M. Cannon, and D. Muñoz-Carpintero, *Efficient prediction strategies for disturbance compensation in stochastic MPC*, International Journal of Systems Science, 2013

Convergence for Stochastic MPC

Both strategies consider the system

$$x_{k+1} = Ax_k + Bu_k + Dw_k$$

and assume that

- ⋆ x_k is measured at time k
- ⋆ (A, B) is stabilizable
- ⋆ the disturbance sequence $\{w_k \in \mathcal{W}\}_{k\in\mathbb{N}}$ is i.i.d. with $\mathbb{E}\{w_k\} = 0$
- ⋆ the probability distribution of w_k is finitely supported in a bounded set $\mathcal{W}$ containing the origin in its interior

Here we additionally assume

- ⋆ $\mathbb{P}\{\|w\| \leq \lambda\} > 0$ for all $\lambda > 0$

Example 1: Affine in the disturbance Stochastic MPC

- State and control constraints:

$$(x_k, u_k) \in \mathbf{Z} \quad \forall k$$

$\mathbf{Z} \subset \mathbb{R}^{n_x} \times \mathbb{R}^{n_u}$ is a convex compact set with $0 \in \mathrm{int}(\mathbf{Z})$

- The predicted control sequence at time k is parameterized as

$$u_{i|k} = v_{i|k} + \sum_{j=0}^{i-1} M_{i,j} w_{j|k}, \quad i \in \mathbb{N}_{N-1}$$

where $v_{i|k}$ and $M_{i,j}$ are optimization variables at time k, and

$$u_{i|k} = K x_{i|k}, \quad i \geq N$$

Example 1: Affine in the disturbance Stochastic MPC

- MPC cost function:

$$J = \mathbb{E}\Big\{x_{N|k}^\top P x_{N|k} + \sum_{i=0}^{N-1}\big(x_{i|k}^\top Q x_{i|k} + u_{i|k}^\top R u_{i|k}\big)\Big\},$$

  with $Q \succeq 0$, $R \succ 0$, $P \succ 0$ and $(A, Q^{1/2})$ assumed detectable, where P and K satisfy the algebraic Riccati equation

$$P = Q + A^\top P A - K^\top (R + B^\top P B) K$$
$$K = -(R + B^\top P B)^{-1} B^\top P A$$

- A terminal constraint is included in the optimal control problem:

$$x_{N|k} \in \mathbf{X}_f,$$

  where $\mathbf{X}_f$ is a robust positively invariant set under $u = Kx$

Example 1: Affine in the disturbance Stochastic MPC

Optimal control problem solved at each instant k:

$$
\begin{aligned}
\min_{\mathbf{u}_k, \mathbf{x}_k, \theta_k} \quad & J \\
\text{subject to} \quad & x_{i+1|k} = Ax_{i|k} + Bu_{i|k} + Dw_{i|k} \\
& u_{i|k} = v_{i|k} + \sum_{j=0}^{i-1} M_{i,j} w_{j|k} \\
& (x_{i|k}, u_{i|k}) \in \mathbf{Z} \\
& x_{0|k} = x_k, \quad x_{N|k} \in \mathbf{X}_f \\
& \forall w_{i|k} \in \mathcal{W},\ \forall i \in \mathbb{N}_{N-1}
\end{aligned}
$$

where $\theta_k = \left(\{v_{i|k}\}_{i \in \mathbb{N}_{N-1}},\ \{M_{i,j}\}_{j \in \mathbb{N}_{N-1}, i \in \{1,\dots,N-1\}}\right)$

Example 1: Affine in the disturbance Stochastic MPC

- Goulart & Kerrigan (2008) prove that the origin is ISS, but no further results on convergence to the terminal mode
- Wang et al. (2008) prove convergence to the mRPI set by redefining the cost and control policy

Under the assumptions that $\left\{ \begin{array}{l} (A+BK, D) \text{ is controllable} \\ \mathbb{P}\{\|w\| \leq \lambda\} > 0 \text{ for all } \lambda > 0 \end{array} \right\}$ we have:

Theorem

For any feasible initial state, $x_0 \in \mathbf{X}$, the closed loop system satisfies

$$\lim_{k\to\infty} \mathbb{P}\{x_k \in \mathbf{X}_\infty\} = 1$$

and

$$\lim_{k\to\infty} \frac{1}{k}\sum_{j=1}^{k}(x_j^\top Q x_j + u_j^\top R u_j) \overset{a.s.}{=} \lim_{k\to\infty} \mathbb{E}\{\xi_k^\top (Q + K^\top R K)\xi_k\}$$

where $\xi_{k+1} = (A+BK)\xi_k + Dw_k$ *with* $\xi_0 = x_0$

Example 2: Striped affine in the disturbance SMPC

- States and controls are subject to probabilistic constraints

$$\mathbb{P}\{f^\top x_{k+1} + g^\top u_k \leq 1\} \geq p,$$

where $g \in \mathbb{R}^{n_x}$, $f \in \mathbb{R}^{n_u}$, $p \in (0, 1]$.

- Predicted control inputs have the structure:

$$u_{i|k} = Kx_{i|k} + c_{i|k} + \sum_{j=1}^{i-1} L_j w_{i-j|k}, \quad i \in \mathbb{N}_{N-1}$$

$$u_{i|k} = Kx_{i|k} + \sum_{j=1}^{N-1} L_j w_{i-j|k}, \quad i \geq N$$

where $c_{i|k}$ are optimization variables and $A + BK$ is Schur stable

L_j are computed offline by minimizing constraint tightening parameters bounding the effects of disturbances on constraints

Example 2: Striped affine in the disturbance SMPC

- MPC cost function:

$$J = \mathbb{E}\Big\{\sum_{i=0}^{\infty}\big(x_{i|k}^\top Q x_{i|k} + u_{i|k}^\top R u_{i|k} - L_{ss}\big)\Big\}$$

where $Q, R \succ 0$, K satisfies the algebraic Riccati equation

$$P = Q + A^\top P A - K^\top (R + B^\top P B) K$$
$$K = -(R + B^\top P B)^{-1} B^\top P A$$

and

$$L_{ss} = \lim_{i\to\infty} \mathbb{E}\{x_{i|k}^\top Q x_{i|k} + u_{i|k}^\top R u_{i|k}\}$$

Example 2: Striped affine in the disturbance SMPC

Optimal control problem solved at each instant k:

$$
\begin{aligned}
\min_{\mathbf{u}_k,\mathbf{x}_k,\mathbf{c}_k} \quad & J \\
\text{subject to} \quad & x_{i+1|k} = Ax_{i|k} + Bu_{i|k} + Dw_{i|k} \\
& u_{i|k} = c_{i|k} + Kx_{i|k} + \sum_{j=1}^{i-1} L_j w_{i-j|k} \\
& \mathbb{P}\{f^\top x_{i+1|k} + g^\top u_{i|k} \leq 1\} \geq p \\
& x_{0|k} = x_k \\
& \forall w_{i|k} \in \mathcal{W},\ \forall i \in \mathbb{N}_{N+N_2-1}
\end{aligned}
$$

with N_2 is chosen large enough to ensure constraint satisfaction over an infinite prediction horizon

Example 2: Striped affine in the disturbance SMPC

- The optimal value function satisfies (Kouvaritakis et al, 2013):

$$\mathbb{E}\{V_{k+1}\} - V_k \leq -(x_k^\top Q x_k + u_k^\top R u_k) + L_{ss}$$

where $L_{ss} = l_{ss} + \mathbb{E}\{w^\top \mathcal{L}^\top \tilde{P}_c \mathcal{L} w\}$,

$$l_{ss} = \lim_{k\to\infty} \mathbb{E}\big\{\xi_k^\top (Q + K^\top R K)\xi_k\big\}$$

with $\xi_{k+1} = (A + BK)\xi_k + Dw_k$ and $\xi_0 = x_0$.

- This implies

$$\lim_{k\to\infty} \frac{1}{k} \sum_{j=1}^{k} \mathbb{E}\{x_j^\top Q x_j + u_j^\top R u_j\} \leq L_{ss}.$$

However, the state converges to the mRPI set $\mathbf{X}_\infty$ and

$$\lim_{k\to\infty} \frac{1}{k} \sum_{j=1}^{k} (x_j^\top Q x_j + u_j^\top R u_j) = l_{ss}$$

which is the asymptotic performance under $u_k = Kx_k$.

Example 2: Striped affine in the disturbance SMPC

Under the assumptions that $\left\{ \begin{array}{l} (A+BK, D) \text{ is controllable} \\ \mathbb{P}\{\|w\| \leq \lambda\} > 0 \text{ for all } \lambda > 0 \end{array} \right\}$ we have:

Theorem

For any feasible initial state, $x_0 \in \mathbf{X}$*, the closed loop system satisfies*

$$\lim_{k \to \infty} \mathbb{P}\{x_k \in \mathbf{X}_\infty\} = 1$$

and

$$\lim_{k \to \infty} \frac{1}{k} \sum_{j=1}^{k} (x_j^\top Q x_j + u_j^\top R u_j) \stackrel{a.s.}{=} \lim_{k \to \infty} \mathbb{E}\{\xi_k^\top (Q + K^\top R K) \xi_k\}$$

where $\xi_{k+1} = (A+BK)\xi_k + Dw_k$ *with* $\xi_0 = x_0$.

Concluding Remarks

- Generalized analysis of Stochastic MPC convergence:
  - Markov chain convergence results determine asymptotic behaviour of control laws that result in linear dynamics on an RPI terminal set
  - Average closed loop performance converges to that of the linear dynamics on the terminal set
  - These results are obtained using an ISS property, but the limit directly implied by the ISS Lyapunov inequality yields a worse bound
- The paper illustrates the convergence analysis by applying it to two Stochastic MPC strategies
- Future work: remove condition on controllability of (A, D)