Model Predictive Control

Lecture: Robust MPC 2

Colin Jones

Laboratoire d'Automatique, EPFL

Goals of Robust Constrained Control

Uncertain constrained linear system
$x^+ = Ax + Bu + w \qquad (x, u) \in \mathbb{X}, \mathbb{U} \qquad w \in \mathbb{W}$

Design control law $u = \kappa(x)$ such that the system:

1. Satifies constraints : $\{x_i\} \subset \mathbb{X}$, $\{u_i\} \subset \mathbb{U}$ for all disturbance realizations
2. Is stable: Converges to a neighbourhood of the origin
3. Optimizes (expected/worst-case) "performance"
4. Maximizes the set $\{x_0 \mid \text{Conditions 1-3 are met}\}$

Challenge: Cannot predict where the state of the system will evolve

We can only compute a set of trajectories that the system *may* follow

Idea: Design a control law that will satisfy constraints and stabilize the system *for all possible disturbances*

Uncertain State Evolution

Given the current state x_0, the model $x^+ = Ax + Bu + w$ and the set $\mathbb{W}$, where can the state be i steps in the future?

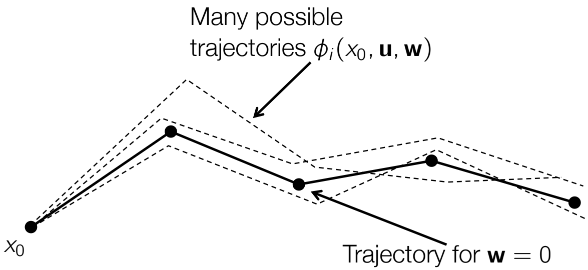

Define $\phi_i(x_0, \mathbf{u}, \mathbf{w})$ as the state that the system will be in at time i if the state at time zero is x_0, we apply the input $\mathbf{u} := \{u_0, \ldots, u_{N-1}\}$ and we observe the disturbance $\mathbf{w} := \{w_0, \ldots, w_{N-1}\}$.

Robust Constraint Satisfaction

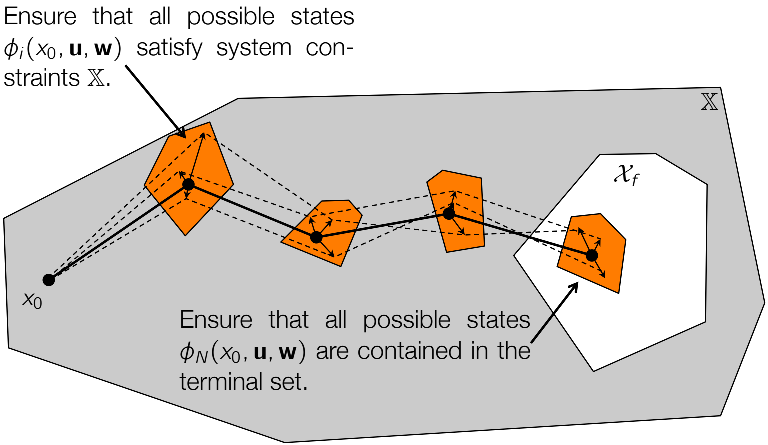

The idea: Compute a set of tighter constraints such that if **the nominal system** meets these constraints, then the uncertain system will too.
We then do MPC **on the nominal system.**

Robust Constraint Satisfaction

Goal: Ensure that constraints are satisfied for the MPC sequence.

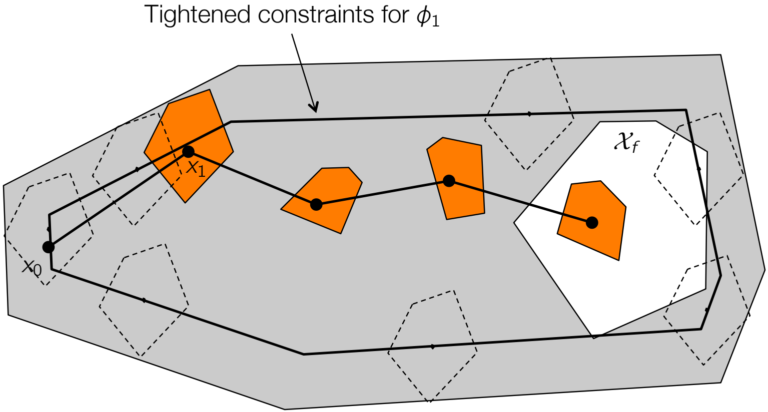

Require: $x_i \in \mathbb{X} \ominus \begin{bmatrix} I & A^0 & \dots & A^{i-1} \end{bmatrix} \mathbb{W}^i$ and

Nominal x_i satisfies tighter constraints → Uncertain state does too

Putting it Together

Robust Open-Loop MPC

$$\begin{aligned} \min_{\mathbf{u}} \quad & \sum_{i=0}^{N-1} l(x_i, u_i) + V_f(x_N) \\ \text{s.t.} \quad & x_{i+1} = Ax_i + Bu_i \\ & x_i \in \mathbb{X} \ominus \mathcal{A}_i \mathbb{W}^i \\ & u_i \in \mathbb{U} \\ & x_N \in \tilde{\mathcal{X}}_f \end{aligned}$$

where $\mathcal{A}_i := \begin{bmatrix} A^0 & A^1 & \dots & A^i \end{bmatrix}$ and $\tilde{\mathcal{X}}_f$ is a robust invariant set for the system $x^+ = (A + BK)x$ for some stabilizing K.

We do **nominal MPC**, but with tighter constraints on the states and inputs.

We can be sure that if the nominal system satisfies the tighter constraints, then the uncertain system will satisfy the real constraints.

$\Rightarrow$ Downside is that $\mathcal{A}^i \mathbb{W}^i$ can be very large

Outline

1. Closed-Loop Predictions

2. Tube-MPC

3. Nominal MPC with noise

MPC as a Game

Two players: Controller vs Disturbance

$$x^+ = f(x, u) + w$$

1. Controller chooses his move u
2. Disturbance decides on his move w **after seeing the controller's move**

MPC as a Game

Two players: Controller vs Disturbance

$$x^+ = f(x, u) + w$$

1. Controller chooses his move u
2. Disturbance decides on his move w **after seeing the controller's move**

What are we assuming when making robust predictions?

1. Controller chooses a **sequence** of N moves in the future $\{u_0, \ldots, u_{N-1}\}$
2. Disturbance chooses N moves **knowing all N moves of the controller**

We are assuming that the controller will do the same thing in the future no matter what the disturbance does!

Can we do better?

Closed-Loop Predictions

What should the future prediction look like?

1. Controller decides his first move u_0

2. Disturbance chooses his first move w_0

3. Controller decides his second move $u_1(x_1)$ **as a function of the first disturbance** w_0 **(recall** $x_1 = Ax_0 + Bu_0 + w_0$**)**

4. Disturbance chooses his second move w_1 as a function of u_1

5. Controller decides his second move $u_2(x_2)$ **as a function of the first two disturbances** w_0**,** w_1

6. ...

Closed-Loop Predictions

We want to optimize over a **sequence of functions** $\{u_0, \mu_1(\cdot), \ldots, \mu_{N-1}(\cdot)\}$, where $\mu_i(x_i) : \mathbb{R}^n \to \mathbb{R}^m$ is called a **control policy**, and maps the state at time i to an input at time i.

Notes:

- This is the same as making μ a function of the disturbances to time i, since the state is a function of the disturbances up to that point
- The first input u_0 is a function of the current state, which is known. Therefore it is not a function, but a single value.

The problem: We can't optimize over arbitrary functions!

Closed-Loop MPC

A solution: Assume some structure on the functions μ_i

Pre-stabilization $\mu_i(x) = Kx + v_i$

- Fixed K, such that $A + BK$ is stable
- Simple, often conservative

Linear feedback $\mu_i(x) = K_i x + v_i$

- Optimize over K_i and v_i
- Non-convex. Extremely difficult to solve...

Disturbance feedback $\mu_i(x) = \sum_{j=0}^{i-1} M_{ij} w_j + v_i$

- Optimize over M_{ij} and v_i
- Equivalent to linear feedback, but convex!
- Can be very effective, but computationally intense.

Tube-MPC $\mu_i(x) = v_i + K(x - \bar{x}_i)$

- Fixed K, such that $A + BK$ is stable
- Optimize over $\bar{x}_i$ and v_i
- Simple, and can be effective

We will cover tube-MPC in this lecture.

Outline

1. Closed-Loop Predictions

2. Tube-MPC

3. Nominal MPC with noise

Tube-MPC

$$x^+ = Ax + Bu + w \qquad (x,u) \in \mathbb{X} \times \mathbb{U} \qquad w \in \mathbb{W}$$

The idea: Seperate the available control authority into two parts

1. A portion that steers the noise-free system to the origin $z^+ = Az + Bv$
2. A portion that compensates for deviations from this system
   $e^+ = (A + BK)e + w$

We fix the linear feedback controller K offline, and optimize over the nominal trajectory $\{v_0, \ldots, v_{N-1}\}$, which results in a convex problem.

[0]Further reading: D.Q. Mayne, M.M. Seron and S.V. Rakovic, Robust model predictive control of constrained linear systems with bounded disturbances, Automatica, Volume 41, Issue 2, February 2005

System Decomposition

Define a 'nominal', noise-free system:

$$z_{i+1} = Az_i + Bv_i$$

Define a 'tracking' controller, to keep the real trajectory close to the nominal

$$u_i = K(x_i - z_i) + v_i$$

for some linear controller K, which stabilizes the nominal system.

Define the error $e_i = x_i - z_i$, which gives the error dynamics:

$$\begin{aligned}
e_{i+1} &= x_{i+1} - z_{i+1} \\
&= Ax_i + Bu_i + w_i - Az_i - Bv_i \\
&= Ax_i + BK(x_i - z_i) + Bv_i + w_i - Az_i - Bv_i \\
&= (A + BK)(x_i - z_i) + w_i \\
&= (A + BK)e_i + w_i
\end{aligned}$$

Error Dynamics

Bound maximum error, or how far the ‘real’ trajectory is from the nominal

$$e_{i+1} = (A + BK)e_i + w_i \qquad\qquad w_i \in \mathbb{W}$$

Dynamics $A + BK$ are stable, and the set $\mathbb{W}$ is bounded, so there is some set $\mathcal{E}$ that e will stay inside for all time.

We want the smallest such set (the ‘minimal invariant set’)

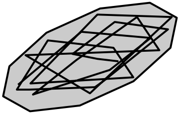

We will cover how to compute this set later

Tube-MPC : The Idea

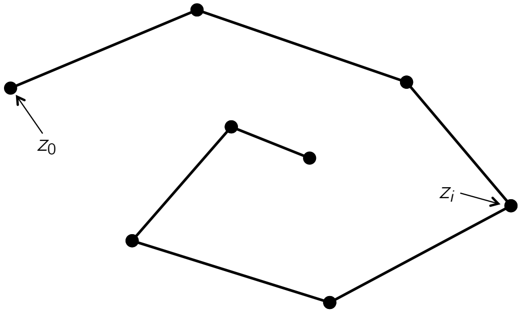

We want to ignore the noise and plan the **nominal trajectory**

Tube-MPC : The Idea

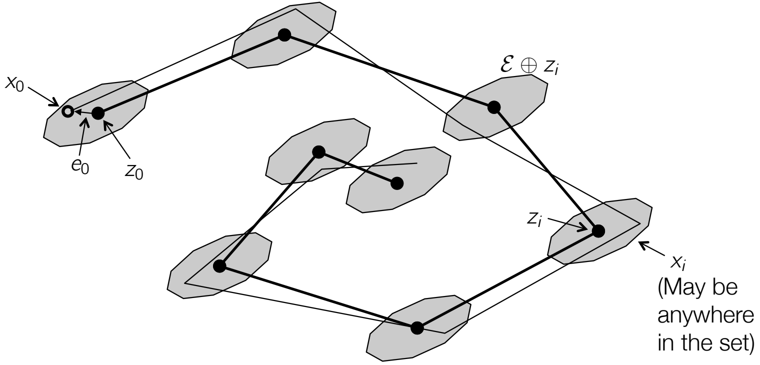

We know that the real trajectory stays 'nearby' the nominal one: $x_i \in z_i \oplus \mathcal{E}$ **because we plan to apply the controller** $u_i = K(x_i - z_i) + v_i$ **in the future** (we won't actually do this, but it's a valid sub-optimal plan)

Tube-MPC : The Idea

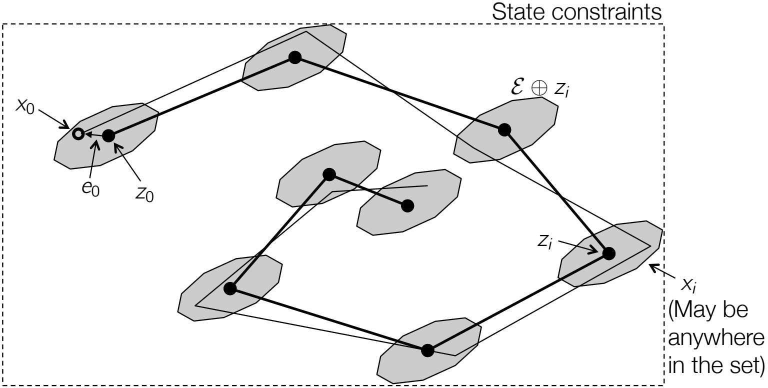

We must ensure that all possible state trajectories satisfy the constraints
This is now equivalent to ensuring that $z_i \oplus \mathcal{E} \subset \mathbb{X}$
(Satisfying input constraints is now more complex - more later)

Tube-MPC

What do we need to make this work?

- Compute the set $\mathcal{E}$ that the error will remain inside
- Modify constraints on nominal trajectory $\{z_i\}$ so that $z_i \oplus \mathcal{E} \subset \mathbb{X}$ and $v_i \in \mathbb{U} \ominus K\mathcal{E}$
- Formulate as convex optimization problem

... and then prove that

- Constraints are robustly satisfied
- The closed-loop system is robustly stable

Tube-MPC

What do we need to make this work?

- **Compute the set $\mathcal{E}$ that the error will remain inside**
- Modify constraints on nominal trajectory $\{z_i\}$ so that $z_i \oplus \mathcal{E} \subset \mathbb{X}$ and $v_i \in \mathbb{U} \ominus K\mathcal{E}$
- Formulate as convex optimization problem

... and then prove that

- Constraints are robustly satisfied
- The closed-loop system is robustly stable

Recall: Robust Invariant Set

Robust constraint satisfaction, for an **autonomous** system $x^+ = f(x, w)$, or **closed-loop** system $x^+ = f(x, \kappa(x), w)$ for a **given** controller κ.

Robust Positive Invariant set
A set $\mathcal{O}^{\mathbb{W}}$ is said to be a robust positive invariant set for the autonomous system $x_{i+1} = f(x_i, w)$ if $$x \in \mathcal{O}^{\mathbb{W}} \;\Rightarrow\; f(x, w) \in \mathcal{O}^{\mathbb{W}} \;, \quad \text{for all } w \in \mathbb{W}$$

Previously we wanted the **maximum robust invariant set**, or the largest set in which our terminal control law works.

We now want the **minimum robust invariant set**, or the smallest set that the state will remain inside despite the noise.

Uncertain State Evolution

Consider the system $x^+ = Ax + w$ and assume that $x_0 = 0$.

Where can the state evolve to? (i.e., how close can we stay to the origin?)

$$
\begin{aligned}
x_1 &= w_0 \\
x_2 &= Ax_1 + w_1 = Aw_0 + w_1 \\
&\vdots \\
x_i &= \sum_{k=0}^{i-1} A^k w_k
\end{aligned}
$$

Assume that $w_i \in \mathbb{W}$ for all i. What is the set F_i that contains all possible states x_i?

$$F_i = \bigoplus_{k=0}^{i-1} A^k \mathbb{W} \ , \quad F_0 := \{0\}$$

where $P \oplus Q := \{x + y \mid x \in P,\ y \in Q\}$

Minimum Robust Invariant Set

As sum goes to infinity, we arrive at the **minimum robust invariant set** mRPI

$$F_\infty = \bigoplus_{k=0}^{\infty} A^k \mathbb{W} \,, \qquad F_0 := \{0\}$$

If there exists an n such that $F_n = F_{n+1}$, then $F_n = F_\infty$

Minimal Invariant Set

Input: A
Output: F_∞
  $\Omega_0 \leftarrow \{0\}$
  loop
    $\Omega_{i+1} \leftarrow \Omega_i \oplus A^i \mathbb{W}$
    if $\Omega_{i+1} = \Omega_i$ **then**
      return $F_\infty = \Omega_i$
    end if
  end loop

- A finite n does not always exist, but a 'large' n is a good approximation
- If n is not finite, there are other methods of computing small invariant sets, which will be slightly larger than F_∞

Computing Minkowski Sums for Polyhedral Data

Given $P := \{x \mid Tx \le t\}$ and $Q := \{x \mid Rx \le r\}$, the Minkowski sum is:

$$
\begin{aligned}
P \oplus Q &:= \{x + y \mid x \in P,\ y \in Q\} \\
&= \{z \mid \exists x, y\ z = x + y,\ Tx \le t,\ Ry \le r\} \\
&= \{z \mid \exists y\ Tz - Ty \le t,\ Ry \le r\} \\
&= \left\{ z \;\middle|\; \exists y\ \begin{bmatrix} T & -T \\ 0 & R \end{bmatrix} \begin{pmatrix} z \\ y \end{pmatrix} \le \begin{pmatrix} t \\ r \end{pmatrix} \right\}
\end{aligned}
$$

This is a **projection** of a polyhedron from (z, y) onto z.

Minkowski Sums in MPT

Recall: We covered computation of projection in Lecture 4.

```
P = polytope(T,t);
Q = polytope(R,r);
Z = zeros(size(R,1),size(T,2));
P_plus_Q = projection(polytope([T -T; Z R], [t;r]), 1:size(T,2));
plot([P Q P_plus_Q]);
```

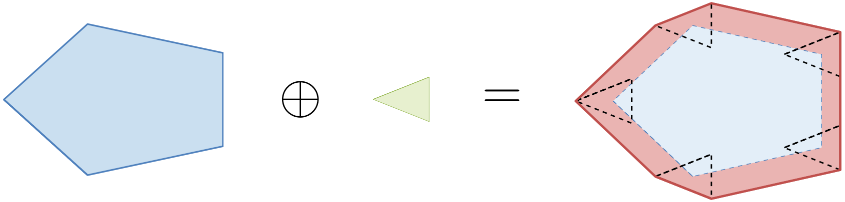

Example

System dynamics

$$x^+ = \left(\begin{bmatrix}1 & 1\\0 & 1\end{bmatrix} + \begin{bmatrix}1\\0.5\end{bmatrix}K\right)x + w \quad \mathbb{W} := \{w \mid |w_1| \leq 0.01,\ |w_2| \leq 0.1\}$$

where K is the LQR controller for $Q = I$, $R = 10$.

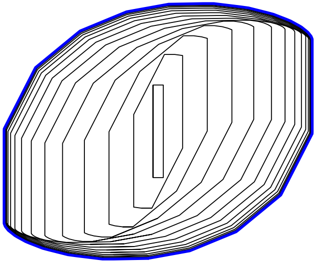

Sets $A^i\mathbb{W}$ converging to minimal robust invariant set F_∞ in the limit

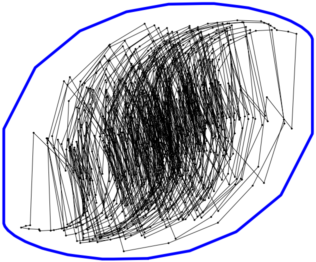

The state trajectory will stay in the set F_∞ for all time

Tube-MPC

What do we need to make this work?

- Compute the set $\mathcal{E}$ that the error will remain inside
- **Modify constraints on nominal trajectory** $\{z_i\}$ **so that** $z_i \oplus \mathcal{E} \subset \mathbb{X}$ **and** $v_i \in \mathbb{U} \ominus K\mathcal{E}$
- Formulate as convex optimization problem

...and then prove that

- Constraints are robustly satisfied
- The closed-loop system is robustly stable

Noisy System Trajectory

Given the nominal trajectory z_i, what can the noisy system trajectory do?

$$x_i = z_i + e_i$$

Don't know what error will be at time i, but it will be in the set $\mathcal{E}$

Therefore, x_i can only be up to $\mathcal{E}$ far from z_i

$$x_i \in z_i \oplus \mathcal{E} = \{z_i + e \mid e \in \mathcal{E}\}$$

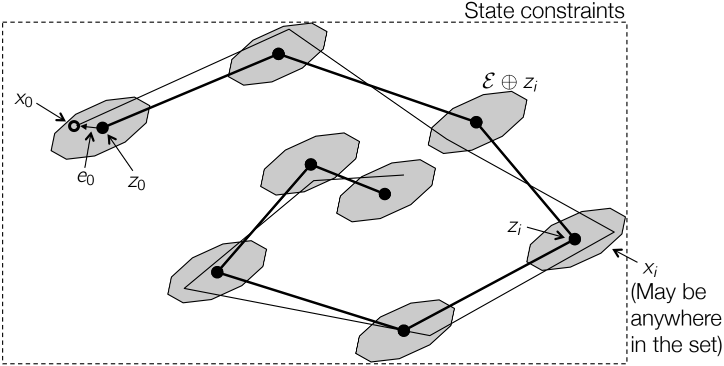

Pontryagin Difference

Pontryagin Difference

Let A and B be subsets of $\mathbb{R}^n$. The Pontryagin Difference is

$$A \ominus B := \{x \mid x + e \in A \ \forall e \in B\}$$

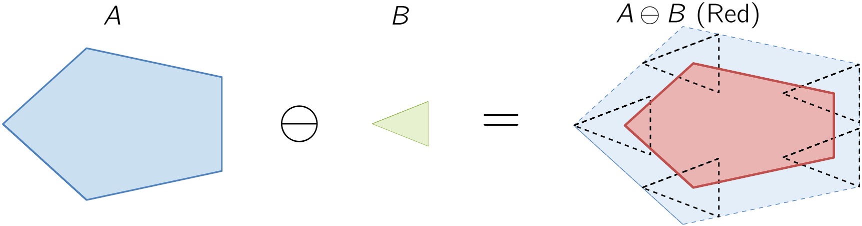

Lemma

$$x \in A \ominus B \Rightarrow x + e \in A \ \forall e \in B$$

We covered how to compute the Pontryagin Difference last week

Constraint Tightening

Goal: $(x_i, u_i) \in \mathbb{X} \times \mathbb{U}$ for all $\{w_0, \dots, w_{i-1}\} \in \mathbb{W}^i$

We want to work with the nominal system $z^+ = Az + Bv$ but ensure that the noisy system $x^+ = Ax + Bu + w$ satisfies the constraints.

Sufficient condition:

$$z_i \oplus \mathcal{E} \subseteq \mathbb{X} \qquad \Leftarrow \qquad z_i \in \mathbb{X} \ominus \mathcal{E}$$

The set $\mathcal{E}$ is known offline - we can compute the constraints $\mathbb{X} \ominus \mathcal{E}$ offline!

A similar condition holds for the inputs:

$$u_i \in K\mathcal{E} \oplus v_i \subset \mathbb{U} \qquad \Leftarrow \qquad v_i \in \mathbb{U} \ominus K\mathcal{E}$$

Tube-MPC

What do we need to make this work?

- Compute the set $\mathcal{E}$ that the error will remain inside
- Modify constraints on nominal trajectory $\{z_i\}$ so that $z_i \oplus \mathcal{E} \subset \mathbb{X}$ and $v_i \in \mathbb{U} \ominus K\mathcal{E}$
- **Formulate as convex optimization problem**

... and then prove that

- Constraints are robustly satisfied
- The closed-loop system is robustly stable

Tube-MPC Problem Formulation

Tube-MPC

Feasible set:
$$\mathcal{Z}(x_0) := \left\{ \mathbf{z}, \mathbf{v} \;\middle|\; \begin{array}{ll} z_{i+1} = Az_i + Bv_i & i \in [0,\ N-1] \\ z_i \in \mathbb{X} \ominus \mathcal{E} & i \in [0,\ N-1] \\ v_i \in \mathbb{U} \ominus K\mathcal{E} & i \in [0,\ N-1] \\ z_N \in \mathcal{X}_f & \\ x_0 \in z_0 \oplus \mathcal{E} & \end{array} \right\}$$

Cost function:
$$V(\mathbf{z}, \mathbf{v}) := \sum_{i=0}^{N-1} l(z_i, v_i) + V_f(z_N)$$

Optimization problem:
$$(\mathbf{v}^\star(x_0), \mathbf{z}^\star(x_0)) = \mathrm{argmin}_{\mathbf{v},\mathbf{z}} \{ V(\mathbf{z}, \mathbf{v}) \mid (\mathbf{z}, \mathbf{v}) \in \mathcal{Z}(x_0) \}$$

Control law:
$$\mu_{\text{tube}}(x) := K(x - z_0^\star(x)) + v_0^\star(x)$$

Main points:

- Optimizing the nominal system, with tightened state an input constraints
- First tube center is optimization variable $\rightarrow$ has to be within $\mathcal{E}$ of x_0
- The cost is with respect to the tube centers
- The terminal set is with respect to the tightened constraints

Tube-MPC

What do we need to make this work?

- Compute the set $\mathcal{E}$ that the error will remain inside
- Modify constraints on nominal trajectory $\{z_i\}$ so that $z_i \oplus \mathcal{E} \subset \mathbb{X}$ and $v_i \in \mathbb{U} \ominus K\mathcal{E}$
- Formulate as convex optimization problem

…and then prove that

- **Constraints are robustly satisfied**
- The closed-loop system is robustly stable

Tube-MPC Assumptions

Much the same as for nominal MPC:

1. The stage cost is a positive definite function, i.e. it is strictly positive and only zero at the origin

2. The terminal set is invariant **for the nominal system** under the local control law $\kappa_f(z)$:

   $$z^+ = Az + B\kappa_f(z) \in \mathcal{X}_f \quad \text{for all } z \in \mathcal{X}_f$$

   All **tightened state and input constraints** are satisfied in $\mathcal{X}_f$:

   $$\mathcal{X}_f \subseteq \mathbb{X} \ominus \mathcal{E},\ \kappa_f(z) \in \mathbb{U} \ominus \mathcal{E} \text{ for all } z \in \mathcal{X}_f$$

3. Terminal cost is a continuous Lyapunov function in the terminal set $\mathcal{X}_f$:

   $$V_f(Az + B\kappa_f(z)) - V_f(z) \leq -l(z, \kappa_f(z)) \text{ for all } z \in \mathcal{X}_f$$

Robust Invariance

Thm: Robust Invariance of Tube-MPC

The set $\mathcal{Z} := \{x \mid \mathcal{Z}(x) \neq \emptyset\}$ is a robust invariant set of the system $x^+ = Ax + B\mu_{\text{tube}}(x) + w$ subject to the constraints $(x, u) \in \mathbb{X} \times \mathbb{U}$.

Let $(\left\{v_0^\star, \ldots, v_{N-1}^\star\right\}, \left\{z_0^\star, \ldots, z_N^\star\right\})$ be the optimal solution for time x_0.

At the next point in time, the state is:

$$x_1 = Ax_0 + BK(x_0 - z_0^\star) + Bv_0^\star + w \quad \text{for some } w \in \mathbb{W}$$

i.e., the state x_1 may have many possible values. We need to show that there exists a feasible solution for **all of them**.

By construction, the state x_1 is in the set $z_1 \oplus \mathcal{E}$ for all $\mathbb{W}$. Therefore (as in standard MPC), the sequence

$$(\left\{v_1^\star, \ldots, v_{N-1}^\star, \kappa_f(z_N^\star)\right\}, \{z_1^\star, \ldots, z_N^\star, Az_N^\star + B\kappa_f(z_N^\star)\})$$

is feasible for all x_1.

Tube-MPC

What do we need to make this work?

- Compute the set $\mathcal{E}$ that the error will remain inside
- Modify constraints on nominal trajectory $\{z_i\}$ so that $z_i \oplus \mathcal{E} \subset \mathbb{X}$ and $v_i \in \mathbb{U} \ominus K\mathcal{E}$
- Formulate as convex optimization problem

…and then prove that

- Constraints are robustly satisfied
- **The closed-loop system is robustly stable**

Robust Stability

Thm: Robust Stability of Tube-MPC

The state x of the system $x^+ = Ax + B\mu_{\text{tube}}(x) + w$ converges in the limit to the set $\mathcal{E}$.

As in standard MPC, we have the relationship:

$$
\begin{aligned}
J^\star(x_0) &= \sum_{i=0}^{N-1} l(z_i^\star, v_i^\star) + V_f(z_N^\star) \\
J^\star(x_1) &\leq \sum_{i=1}^{N} l(z_i^\star, v_i^\star) + V_f(z_{N+1}^\star) \\
&= J^\star(x_0) - \underbrace{l(z_0^\star, v_0^\star)}_{\geq 0} + \underbrace{V_f(z_{N-1}^\star) - V_f(z_N^\star) + l(z_N^\star, \kappa_f(z_N^\star))}_{\leq 0 \ (V_f \text{ is a Lyapunov function in } \mathcal{X}_f)}
\end{aligned}
$$

This shows that $\lim_{i\to\infty} J(z_0^\star(x_i)) = 0$, and therefore $\lim_{i\to\infty} z_0^\star(x_i) = 0$.

However, x_i does not tend to zero! It only stays within a robust invariant set centered at $z_0^\star(x_i)$: $\lim_{i\to\infty} \operatorname{dist}(x_i, \mathcal{E}) = 0$, where dist is any distance function.

Putting it all together: Tube MPC

To implement tube MPC:

— Offline —

1. Choose a stabilizing controller K so that $\|A + BK\| < 1$
2. Compute the minimal robust invariant set $\mathcal{E} = F_\infty$ for the system $x^+ = (A + BK)x + w$, $w \in \mathbb{W}$[1]
3. Compute the tightened constraints $\tilde{\mathbb{X}} := \mathbb{X} \ominus \mathcal{E}$, $\tilde{\mathbb{U}} := \mathbb{U} \ominus \mathcal{E}$
4. Choose terminal weight function V_f and constraint $\mathcal{X}_f$ satisfying assumptions on slide 35

— Online —

1. Measure / estimate state x
2. Solve the problem $(\mathbf{v}^\star(x), \mathbf{z}^\star(x)) = \text{argmin}_{\mathbf{v},\mathbf{z}} \{ V(\mathbf{z}, \mathbf{v}) \mid (\mathbf{z}, \mathbf{v}) \in \mathcal{Z}(x) \}$ (Slide 33)
3. Set the input to $u = K(x - z_0^\star(x)) + v_0^\star(x)$

---

[1] Note that it is often not possible to compute the minimal robust invariant set, as it may have an infinite number of facets. Therefore, we often take an invariant outer approximation.

Example

System dynamics

$$x^+ = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} x + \begin{bmatrix} 1 \\ 0.5 \end{bmatrix} u + w \qquad \mathbb{W} := \{w \mid |w_1| \leq 0.01,\ |w_2| \leq 0.1\}$$

Constraints:

$$\mathbb{X} := \{x \mid \|x\|_\infty \leq 1\} \qquad \mathbb{U} := \{u \mid \|u\| \leq 1\}$$

Stage cost is:

$$l(z, v) := {z_i}^T Q z_i + {v_i}^T R v_i$$

where

$$Q := \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \qquad R := 10$$

Offline Design - Compute Minimal Invariant Set

1. Choose a stabilizing controller K so that $\|A + BK\| < 1$
2. Compute the minimal robust invariant set $\mathcal{E} = F_\infty$ for the system $x^+ = (A + BK)x + w$, $w \in \mathbb{W}$

We take the LQR controller for $Q = I$, $R = 1$:

$$K := \begin{bmatrix} -0.5198 & -0.9400 \end{bmatrix}$$

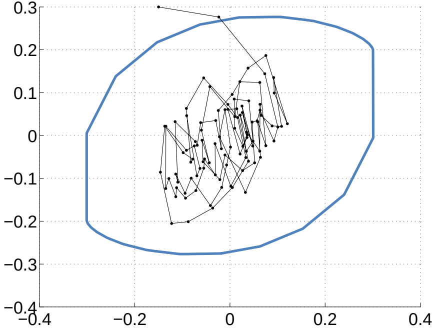

Evolution of the system $x^+ = (A + BK)x + w$ for $x_0 = \begin{bmatrix} -0.1 & 0.2 \end{bmatrix}^T$

Offline Design - Tighten State Constraints

3. Compute the tightened constraints $\tilde{\mathbb{X}} := \mathbb{X} \ominus \mathcal{E}$, $\tilde{\mathbb{U}} := \mathbb{U} \ominus K\mathcal{E}$

$$\mathbb{X} = \{x \mid \|x\|_\infty \leq 1\} = \left\{x \;\middle|\; \begin{bmatrix} I \\ -I \end{bmatrix} x \leq \begin{bmatrix} \mathbf{1} \\ \mathbf{1} \end{bmatrix} \right\}$$

If $\mathcal{E} = \{x \mid Fx \leq f\}$, then the tightened constraint sets are:

$$\mathbb{X} \ominus \mathcal{E} = \{x \mid x + e \in \mathbb{X}\ \forall e \in \mathcal{E}\} = \left\{x \;\middle|\; \begin{bmatrix} I \\ -I \end{bmatrix} x + \begin{bmatrix} I \\ -I \end{bmatrix} e \leq \begin{bmatrix} \mathbf{1} \\ \mathbf{1} \end{bmatrix}\ \forall e \in \mathcal{E} \right\}$$

$$= \left\{x \;\middle|\; \begin{bmatrix} I \\ -I \end{bmatrix} x \leq \begin{bmatrix} 1 - \max\left\{\begin{bmatrix}1 & 0\end{bmatrix} e \mid e \in \mathcal{E}\right\} \\ 1 - \max\left\{\begin{bmatrix}0 & 1\end{bmatrix} e \mid e \in \mathcal{E}\right\} \\ 1 + \max\left\{\begin{bmatrix}1 & 0\end{bmatrix} e \mid e \in \mathcal{E}\right\} \\ 1 + \max\left\{\begin{bmatrix}0 & 1\end{bmatrix} e \mid e \in \mathcal{E}\right\} \end{bmatrix} \right\}$$

The maximizations are all linear programs and can be computed offline.

The results is a polytope with smaller RHS.

Offline Design - Tighten State Constraints

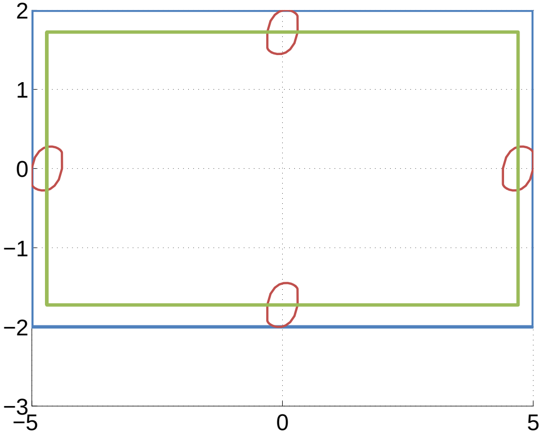

Blue : Original constraint set $\mathbb{X}$

Red : Error set $\mathcal{E}$

Green : Tightened constraints $\mathbb{X} \ominus \mathcal{E}$

Offline Design - Tighten Input Constraints

We compute $\mathbb{U} \ominus K\mathcal{E}$ in the same manner:

$$\begin{aligned}\mathbb{U} \ominus K\mathcal{E} &= \left\{ u \;\middle|\; \begin{bmatrix} 1 \\ -1 \end{bmatrix} u \le \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\} \ominus \{ Kx \mid Fx \le f \} \\ &= \left\{ u \;\middle|\; \begin{bmatrix} 1 \\ -1 \end{bmatrix} u \le \begin{bmatrix} 1 - \max\{ Kx \mid Fx \le f \} \\ 1 + \max\{ Kx \mid Fx \le f \} \end{bmatrix} \right\}\end{aligned}$$

Offline Design - Terminal Weights and Constraints

We need to find a function V_f and a set $\mathcal{X}_f$ that satisfy the conditions on slide 35:

1. The terminal set is invariant **for the nominal system** under the local control law $\kappa_f(z)$:

   $$z^+ = Az + B\kappa_f(z) \in \mathcal{X}_f \quad \text{for all } z \in \mathcal{X}_f$$

   All **tightened state and input constraints** are satisfied in $\mathcal{X}_f$:

   $$\mathcal{X}_f \subseteq \mathbb{X} \ominus \mathcal{E},\ \kappa_f(z) \in \mathbb{U} \ominus \mathcal{E} \text{ for all } z \in \mathcal{X}_f$$

2. Terminal cost is a continuous Lyapunov function in the terminal set $\mathcal{X}_f$:

   $$V_f(Az + B\kappa_f(z)) - V_f(z) \leq -l(z, \kappa_f(z)) \text{ for all } z \in \mathcal{X}_f$$

Offline Design - Terminal Constraint

We base our terminal weights and constraints on the LQR controller (many other choices possible).

Choose the terminal control law to the the LQR control law: $\kappa_f(x) = Kx$ where the weights Q and R are taken the same as for our MPC problem.

We need a set $\mathcal{X}_f$ that is invariant under this controller and contained in the tightened constraints:

$$\text{pre}(\mathcal{X}_f) \subseteq \mathcal{X}_f \quad \text{and} \quad \mathcal{X}_f \subseteq \mathbb{X} \ominus \mathcal{E} \quad \text{and} \quad K\mathcal{X}_f \subseteq \mathbb{U} \ominus K\mathcal{E}$$

We know how to compute the maximal invariant set for linear systems with polytopic constraints (Lecture: Introduction to Constrained Systems)

Offline Design - Terminal Cost

We need to find a function V_f with the property:

$$V_f(Az + B\kappa_f(z)) - V_f(z) \leq -l(z, \kappa_f(z)) \text{ for all } z \in \mathcal{X}_f$$

where we've chosen $\kappa_f(z) = Kz$ (the optimal LQR controller)

Recall the the optimal cost of the LQR control law is:

$$V^\star(z_0) = \sum_{i=0}^{\infty} {z_i}^T (Q + K^T RK) z_i = z_0^T P z_0$$

where P is the solution to a discrete-time Riccati equation.

We know that $V^\star(z)$ is a Lyapunov function for the system $z^+ = (A + BK)z$:

$$\begin{aligned} V^\star(z_1) - V^\star(z_0) &= \sum_{i=1}^{\infty} {z_i}^T (Q + K^T RK) z_i - \sum_{i=0}^{\infty} {z_i}^T (Q + K^T RK) z_i \\ &= -{z_0}^T (Q + K^T RK) z_0 = -l(z_0, \kappa_f(z_0)) \end{aligned}$$

which is exactly what we need, and therefore, we can take $V_f(z) = z^T Pz$.

Tubes - Example

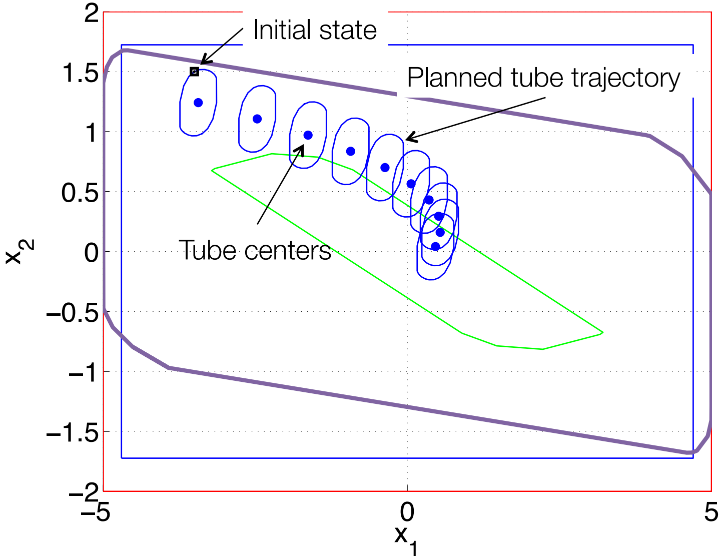

Tubes - Example

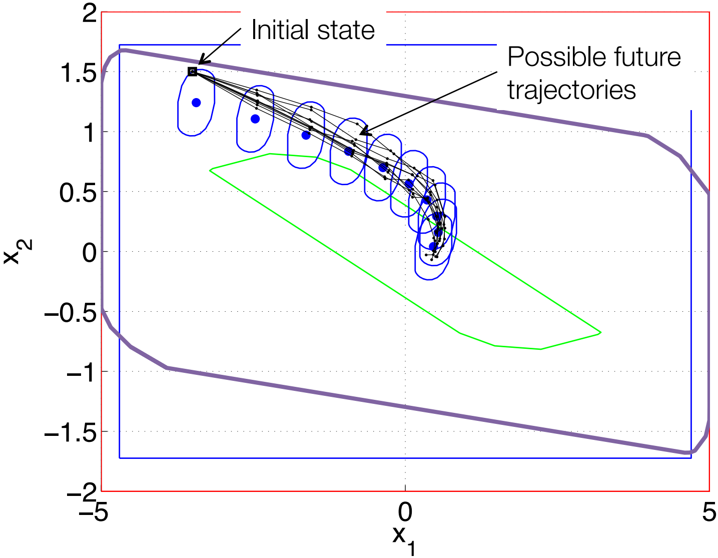

Tubes - Example

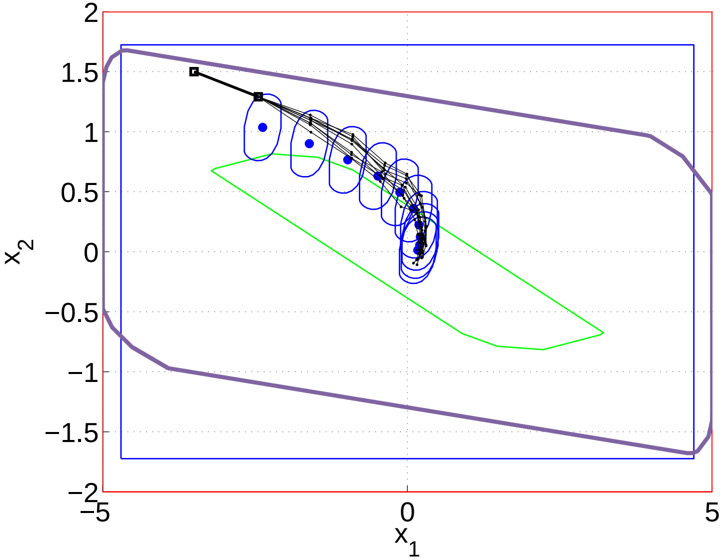

Tubes - Example

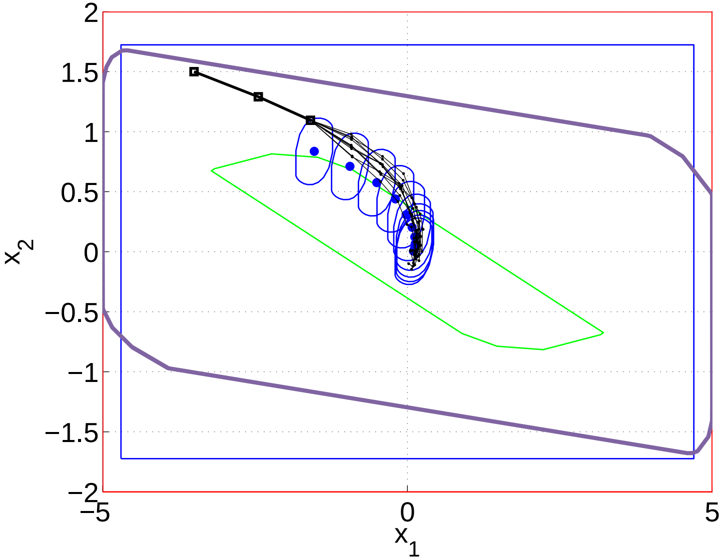

Tubes - Example

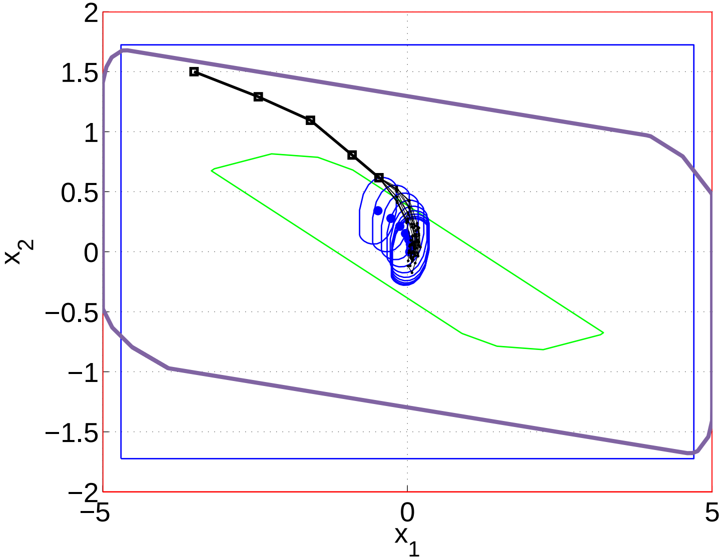

Tubes - Example

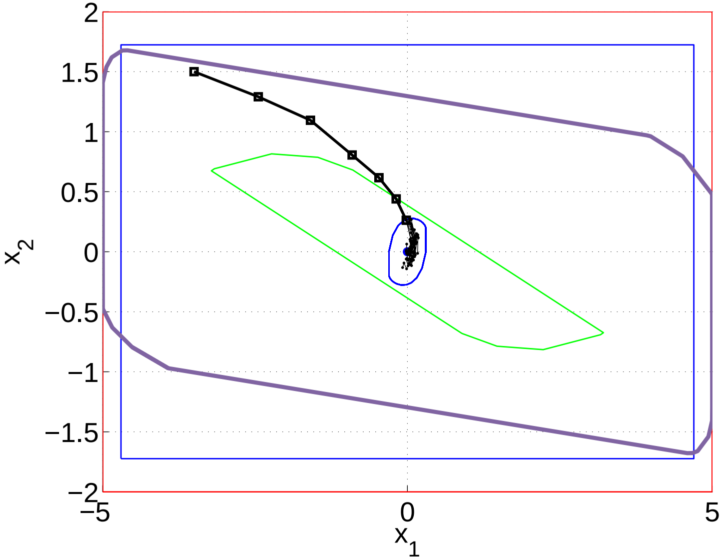

Tubes - Example

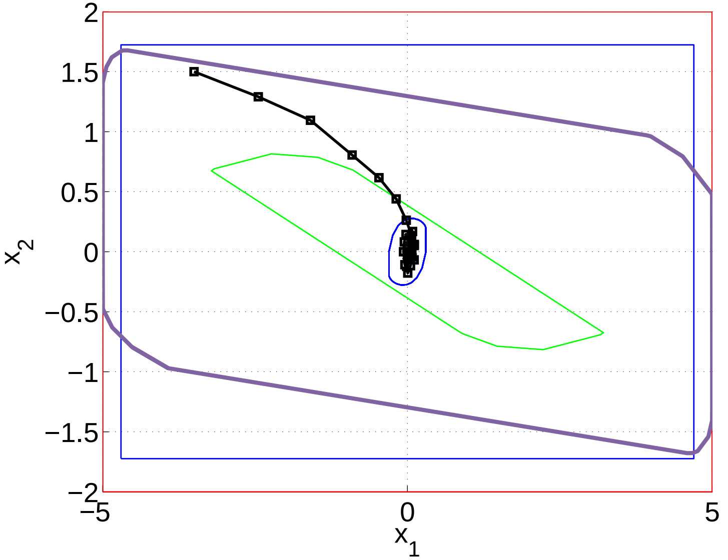

Tube MPC - Summary

Idea:

- Split input into two parts: One to steer system (v), one to compensate for the noise (Ke)

$$u = Ke + v$$

- Optimize for the nominal trajectory, ensuring that any deviations stay within constraints

Benefits:

- Less conservative than open-loop robust MPC (we're now actively compensating for noise in the prediction)
- Works for unstable systems
- Optimization problem to solve is simple

Cons:

- Sub-optimal MPC (optimal is extremely difficult)
- Reduced feasible set when compared to nominal MPC
- We need to know what $\mathbb{W}$ is (this is usually not realistic)

Outline

1. Closed-Loop Predictions

2. Tube-MPC

3. Nominal MPC with noise

Nominal MPC with Noise

We want to control the noisy system:

$$x^+ = Ax + Bu + w$$

What happens if we just ignore the noise and hope for the best?

Setup and solve a standard MPC problem:

$$\begin{aligned} V^\star(x_0) = \min_{\mathbf{u}} \ & \sum_{i=0}^{N-1} l(x_i, u_i) + V_f(x_N) \\ \text{s.t.} \quad & x_{i+1} = Ax_i + Bu_i \\ & (x_i, u_i) \in \mathbb{X} \times \mathbb{U} \\ & x_N \in \mathcal{X}_f \end{aligned}$$

Our closed-loop system is now:

$$x^+ = Ax + Bu_0^\star(x) + w$$

Example

Consider the same example again, with the same noise, but now we just pretend it's not there in the controller.

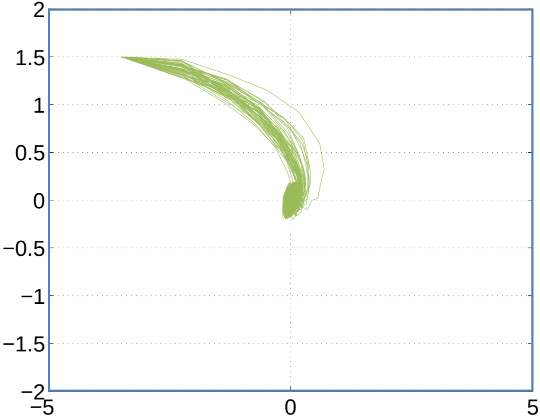

- 100 trajectories with different noise realizations
- Seems to work fine?!

Example

Consider the same example again, with the same noise, but now we just pretend it's not there in the controller.

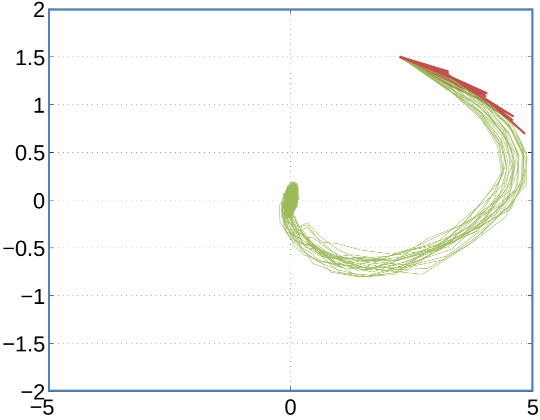

- 100 trajectories with different noise realizations
- Seems to work fine?!
- Can no longer be certain it will work!
- For some states it will work sometimes

How do we formalize this idea?

What Happens to Our Lyapunov Function?

Recall: The optimal cost $V^\star(x)$ is a Lyapunov function for the nominal system

$$V^\star(Ax + Bu^\star(x)) - V^\star(x) \le -l(x, u^\star(x))$$

However, our state at the next point in time is now

$$x^+ = Ax + Bu^\star(x) + w$$

Do we still have a Lyapunov decrease?

What Happens to Our Lyapunov Function?

Assume: Optimal cost $V^\star$ is continuous[2]

$$|V^\star(Ax + Bu^\star(x) + w) - V^\star(Ax + Bu^\star(x))|$$
$$\leq \gamma \|Ax + Bu^\star(x) + w - (Ax + Bu^\star(x))\| = \gamma\|w\|$$

Our Lyapunov decrease can be bounded as:

$$\begin{aligned}
&V^\star(Ax + Bu^\star(x) + w) - V^\star(x) \\
&= V^\star(Ax + Bu^\star(x) + w) - V^\star(x) - V^\star(Ax + Bu^\star(x)) + V^\star(Ax + Bu^\star) \\
&\leq V^\star(Ax + Bu^\star(x)) - V^\star(x) + \gamma\|w\| \\
&\leq -l(x, u^\star(x)) + \gamma\|w\|
\end{aligned}$$

- Amount of decrease grows with $\|x\|$
- Amount of increase is upper bounded by $\max\{\|w\| \mid w \in \mathbb{W}\}$

Therefore we will move towards the origin until there is a balance between the size of x and the size of w

[2] True for linear systems, convex constraints and continuous stage costs.

Input-to-State Stability

What we have shown is that our system is **Input-to-State Stable**.

Much more general theory than what is given here[3]

Asymptotic stability

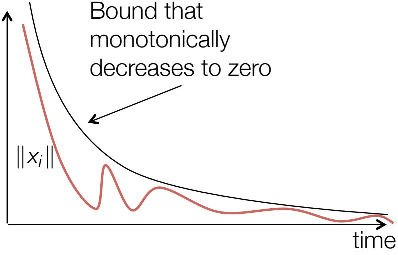

System converges to zero

ISS stability

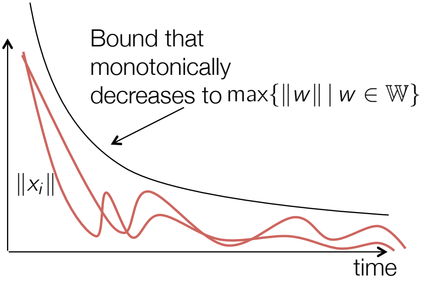

Converges to set around zero, who's size is determined by size of the noise

[3] Limon, D., Alamo, T., Raimondo, D. M., Muñoz de la Peña, D., Bravo, J. M., Ferramosca, A., and Camacho, E. F. (2009). Input-to-State Stability: A Unifying Framework for Robust Model Predictive Control. In L. Magni, D. M. Raimondo, & F. Allgöwer (Eds.), Nonlinear Model Predictive Control (Vol. 384, pp. 1-26). Berlin, Heidelberg: Springer Berlin Heidelberg. doi:10.1007/978-3-642-01094-1

Nominal MPC for Uncertain Systems - Summary

Idea

- Ignore the noise and hope it works

Benefits

- Simple
- No knowledge of the noise set $\mathbb{W}$ required - 'just works'
- Often very effective in practice (this is what most practitioners do anyway)
- Feasible set is large (we can find a solution, but it may be garbage)
- Region of attraction may be larger than other approaches

Cons

- Very difficult to determine region of attraction (set of states in which the controller works)
- Hard to tune - no obvious way to tradeoff robustness against performance

Robust MPC for Uncertain Systems - Summary

Idea

- Compensate for noise in prediction to ensure all constraints will be met

Cons

- Complex (some schemes are simple to implement, like tubes, but complex to understand)
- Must know the largest noise $\mathbb{W}$
- Often very conservative
- Feasible set may be small

Benefits

- Feasible set is invariant - we know exactly when the controller will work
- Easier to tune - knobs to tradeoff robustness against performance