CDS 270-2: Lecture 3-2
Receding Horizon Control

Tamás Keviczky

12 April 2006

Goals:

- Describe the receding horizon control principle
- Discuss main ingredients and different approaches for stability
- Example 1: Discrete-time formulation with terminal cost and constraint
- Example 2: Continuous-time unconstrained formulation using CLF

Reading:

- D.Q. Mayne, J.B. Rawlings, C.V. Rao and P.O.M. Scokaert, “Constrained model predictive control: Stability and optimality”, *Automatica*, 2000, Vol. 36, No. 6, pp. 789-814.
- R.M. Murray et al., “Online Control Customization via Optimization-Based Control”, In *Software-Enabled Control: Information Technology for Dynamical Systems*, T. Samad and G. Balas (eds.), IEEE Press, 2001

Receding horizon control principle

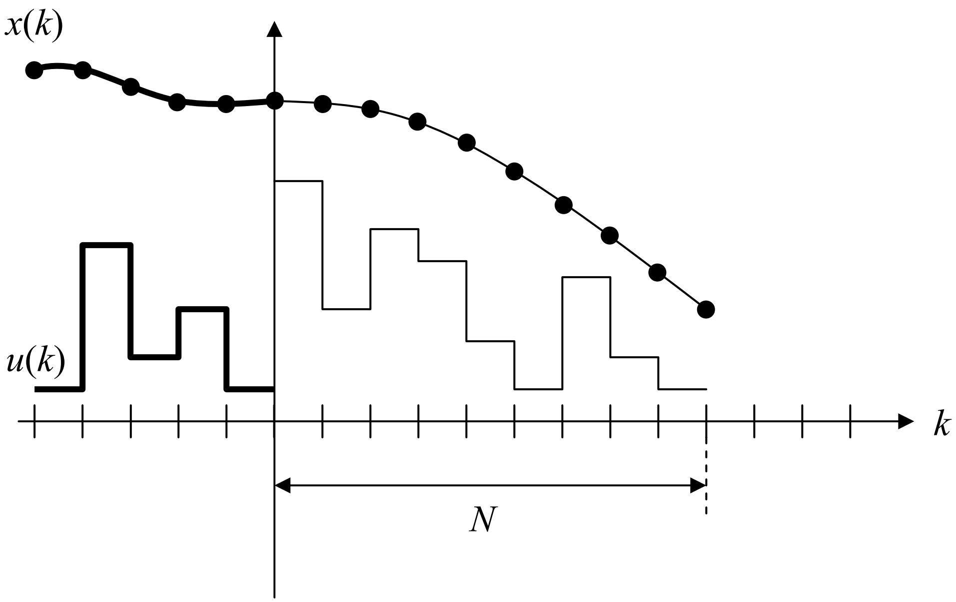

Receding Horizon Control (RHC) is a form of control, in which:

- The current control action is obtained by solving on-line, at each sampling instant, a finite horizon open-loop optimal control problem, using the current state of the plant as the initial state.

- The first part of the optimal control sequence is then applied to the plant and the procedure repeated for future sampling times.

Historical perspective

RHC definition, advantages, features

- Current control action is obtained by solving a finite horizon optimal control problem, using current state of the plant as initial state.
- Ability to cope with hard state and control constraints
- Computationally intensive (online optimization)

Became popular in process industries

- Infrequent control updates (even hours) allow time for computation
- Operating points are required to be on the boundary of admissible states and controls (for efficiency)

Application on systems with fast dynamics

- Deal with nonlinear, constrained systems
- Online optimization allows updates to cost, model parameters, constraints (flexibility)
- Dramatic increase in available computational power
- Developments in theory and methodologies provide more insight and reduce computational requirements (differential flatness, explicit solution)

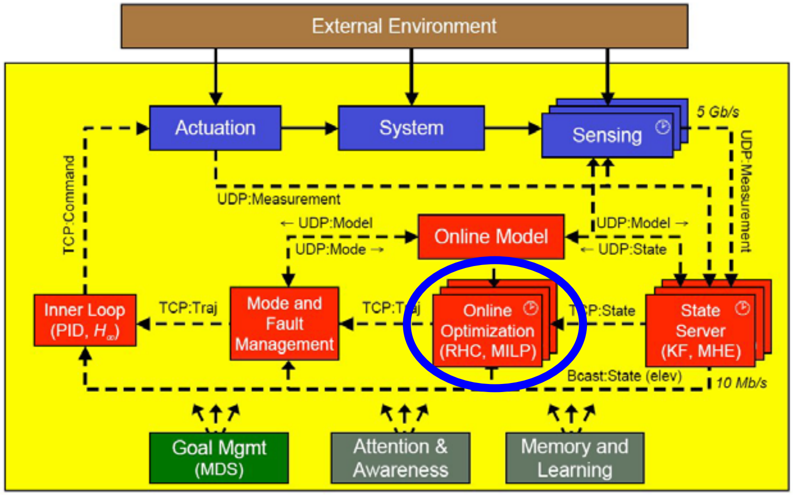

Discrete-time nonlinear constrained RHC formulation

The following finite time optimal control problem is solved in every step:

$$\mathcal{P}_N(x_t): \qquad J_N^*(x_t) \triangleq \min_{\mathbf{u}} \sum_{k=0}^{N-1} \overset{\text{Stage cost}}{\overset{\uparrow}{q(x_k,u_k)}} + \overset{\text{Terminal cost}}{\overset{\uparrow}{V(x_N)}}$$

$$\begin{aligned} \text{subject to } \; & x_{k+1} = f(x_k,u_k) && k=0,\ldots,N-1, \longrightarrow \text{Dynamics} \\ & u_k \in \mathcal{U} && k=0,\ldots,N-1, \longrightarrow \text{Input constraint} \\ & x_k \in \mathcal{X} && k=0,\ldots,N, \longrightarrow \text{State constraint} \\ & x_0 = x_t, && \\ & x_N \in \mathcal{X}_f \subset \mathcal{X}. && \longrightarrow \text{Terminal constraint} \end{aligned}$$

Assumptions

- f, V and q are continuous functions
- $\mathcal{U} \subset R^m$ is a compact set, $\mathcal{X} \subset R^n$ and $\mathcal{X}_f \subset R^n$ are closed sets
- a feasible solution exists to problem $\mathcal{P}_N(x_t)$

Typical choices

- $q(x,u) = \|Qx\|_p + \|Ru\|_p$
- $V(x) = \|Px\|_p$

($\|\cdot\|_p$ is the vector p-norm, with $p = 1, 2, \ldots, \infty$.)

Discrete-time nonlinear constrained RHC formulation

RHC procedure

The minimizing control sequence is a function of the current state x_t

$$\mathbf{u}^*_{[0,\ldots,N-1],t}(x_t) = \left\{u^*_{0,t}, u^*_{1,t}, \ldots, u^*_{N-1,t}\right\}$$

The control applied to the plant at time t is the first element of this sequence

$$u_t = u^*_{0,t}$$

At the next time instant the above procedure is repeated for another N-step-ahead optimization horizon.

The receding horizon procedure implicitly defines a time-invariant control policy $\kappa : \mathcal{X} \rightarrow \mathcal{U}$ of the form

$$\kappa(x_t) = u^*_0$$

Implicit RHC

It is common in receding horizon control applications to compute on-line numerically, at time t, using the current state x_t, the optimal control move $\kappa(x_t)$. In this case, we call it an *implicit* receding horizon optimal policy.

Explicit RHC

In some cases (e.g. linear system, linear constraints, quadratic cost), we can explicitly calculate the control law $\kappa(\cdot)$. In this case, we say that we have an *explicit* receding horizon optimal policy.

Receding horizon control features

Main motivation

- Ability to handle control problems where off-line computation of a control law is difficult or impossible or otherwise undesireable (e.g. flexibility).

Few examples

- constrained systems
- unconstrained nonlinear systems (special structure of plant dynamics needed for off-line control law computation)

Contrast

- *RHC*
  online optimal control problem solution, mathematical programming

- *Standard feedback solution*
  solving Hamilton-Jacobi-Bellman diff. equation (dynamic programming), more difficult in general

Thus RHC differs from standard approaches mainly in its *implementation*.

However, due to requirements on the control solution time,
finite horizon formulation is necessary, which leads to issues…

Issues with RHC

Issues arise from implementation of *finite horizon* constrained optimal problems.
Predicted open-loop trajectories are not the same as closed-loop trajectories!
(even assuming perfect model, no disturbances)

$$\left\{u^*_{0,t},\ldots,u^*_{N-1,t}\right\} \triangleq \arg\min_{\mathbf{u}} \sum_{k=0}^{N-1}\left(x_k^{\mathrm{T}} x_k + u_k^T u_k\right) + x_N^{\mathrm{T}} x_N$$

$$\text{subject to } \; x_{k+1} = \begin{bmatrix} 2 & 1 \\ 0 & 3 \end{bmatrix} x_k + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u_k,$$

$$x_0 = x_t.$$

$\Rightarrow$

$u_t = u^*_{0,t}$

$N = 2$ unstable!

($N = 3$ stable)

Optimality $\not\Rightarrow$ Stability

Reason

Stability is an *asymptotic* notion, defined over infinite time, as opposed to optimality.
Optimality of such control problems *does not* imply stability!

How to get stability of RHC?

Ideas

- Use optimality of fixed horizon control sequence to prove stability of RHC.
- Use the value function

$$J_N^*(x)$$

as a Lyapunov function.

Problem

- Optimization problem, finite future horizon
- Stability, infinite future horizon

↓

Trick

- Add appropriate weighting on terminal state
- Account for things beyond the horizon

Lyapunov stability

We say that the equilibrium point is

- **Lyapunov stable** in $\mathcal{X}$

  if for any $\varepsilon > 0$, there exists $\delta > 0$ such that $x_0 \in \mathcal{X}$ and $\|x_0\| < \delta \;\Rightarrow\; \|x_k\| < \varepsilon$ for all $k \geq 0$.

- **Attractive** in $\mathcal{X}$

  if there exists $\delta > 0$ such that $x_0 \in \mathcal{X}$ and $\|x_0\| < \delta \;\Rightarrow\; \lim_{k \to \infty} x_k = 0$.

- **Asymptotically stable**

  if both Lyapunov stable and attractive.

Stability properties can be tested by finding a so-called *Lyapunov function* $V : \mathcal{X} \to [0, \infty)$, which satisfies certain conditions.

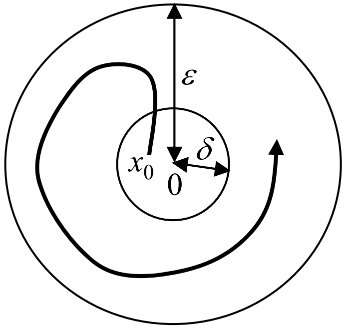

Lyapunov stable

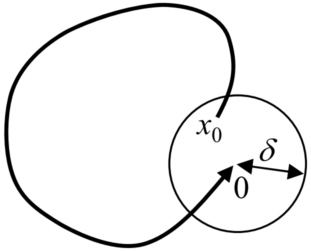

Attractive

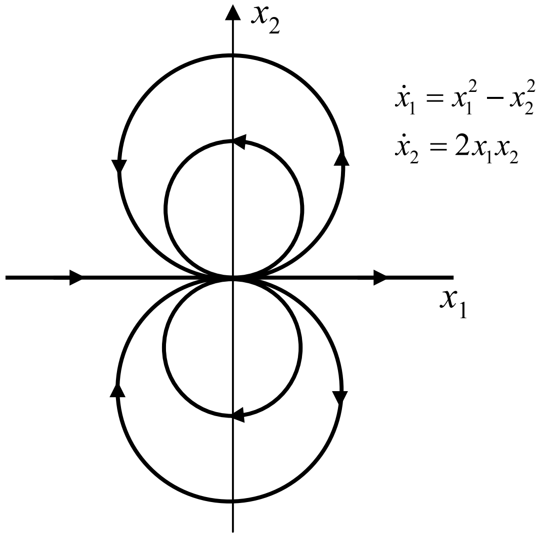

Attractive but not stable

Theorem for attractivity

Theorem (Attractivity in $\mathcal{X}$)

Let $\mathcal{X}$ be a nonempty set in R^n. Let $f : R^n \to R^n$ be such that $f(0) = 0$ and $f(\mathcal{X}) \subset \mathcal{X}$.
Assume there exists a Lyapunov function $V : \mathcal{X} \to [0, \infty)$, with the following properties:

i. $V(\cdot)$ decreases along the trajectories of $x_{k+1} = f(x_k)$ that start in $\mathcal{X}$ in the following way:

$$V(f(x)) - V(x) \leq -\gamma(\|x\|) \text{ for all } x \in \mathcal{X},$$

where $\gamma : [0, \infty) \to [0, \infty)$ is a continuous function and $\gamma(t) > 0$ for all $t > 0$.

ii. For every unbounded sequence $\{x_k\} \subset \mathcal{X}$ there is some l such that

$$\limsup_{k\to\infty} V(x_k) > V(x_l).$$

Then the origin is *globally attractive* in $\mathcal{X}$:
For all $x_0 \in \mathcal{X}$, $\lim_{k\to\infty} x_k = 0$ where the sequence is generated by $x_{k+1} = f(x_k)$.

Property (i) guarantees that $\lim_{k\to\infty} V(x_k) \geq 0$ exists.

Property (ii) prevents situations such as →

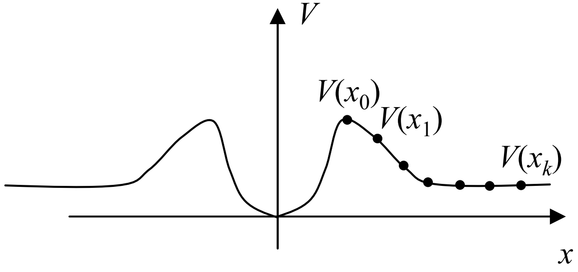

Theorem for Lyapunov stability

Theorem (Lyapunov Stability)

- Let $\mathcal{X}$ be a set in R^n that contains an open neighborhood of the origin $\mathcal{X}_\eta(0) = \{x \in R^n : \|x\| < \eta\}$.

- Let $f : R^n \rightarrow R^n$ be such that $f(0) = 0$ and $f(\mathcal{X}) \subset \mathcal{X}$.

- Assume that there exists a Lyapunov function $V : \mathcal{X} \rightarrow [0, \infty)$, $V(0) = 0$, satisfying the following properties:

  i. $V(\cdot)$ is continuous on $\mathcal{X}_\eta(0)$,

  ii. If $\{x_k\} \subset \mathcal{X}$ is such that $\lim_{k\rightarrow\infty} V(x_k) = 0$ then $\lim_{k\rightarrow\infty} x_k = 0$,

  iii. $V(f(x)) - V(x) \leq 0$ for all $x \in \mathcal{X}_\eta(0)$.

Then the origin is a *Lyapunov stable equilibrium* point for $x_{k+1} = f(x_k)$ in $\mathcal{X}$.

Recipe for stability of RHC

- Define a *terminal control law* and an associated *terminal cost function* that captures the impact of using the terminal control law over infinite time.
- Usually, the chosen terminal control laws are relatively simple and only feasible in a restricted (local) region.
- Must be able to steer the system into restricted terminal region over finite time window.
- Terminal region must be *invariant* under the terminal control law

Typical ingredients to get sufficient conditions for stability:

The Terminal Triple $(\mathcal{X}_f, \kappa_f, V)$

i. a *terminal constraint set* $\mathcal{X}_f$ in the state space which is invariant under the terminal control law,

ii. a feasible *terminal control law* κ_f that keeps the state in the terminal constraint set,

iii. a *terminal cost function* V, which usually corresponds to the objective function value generated by the use of the terminal control law over infinite time.

Definitions

Set of Feasible Initial States

The set $\mathcal{X}_0$ of feasible initial states is the set of initial states $x \in \mathcal{X}$ for which there exist feasible state and control sequences for the fixed horizon optimal control problem $\mathcal{P}_N(x)$.

Positively Invariant Set

The set $\mathcal{S} \subset R^n$ is positively invariant for the system $x_{k+1} = f(x_k, u_k)$ under the control $u_k = \kappa(x_k)$ if $f(x, \kappa(x)) \in \mathcal{S}$ for all $x \in \mathcal{S}$.

Assumptions on the data of problem $\mathcal{P}_N(x)$

Conditions for Stability

A1 The stage cost $q(x, u)$ satisfies $q(0, 0) = 0$ and $q(x, u) \geq \gamma(\|x\|)$ for all $x \in \mathcal{X}_0$, $u \in \mathcal{U}$, where $\gamma : [0, \infty) \to [0, \infty)$ is continuous, $\gamma(t) > 0$ for all $t > 0$, and $\lim_{t\to\infty} \gamma(t) = \infty$.

A2 The terminal cost $V(x)$ satisfies $V(0) = 0$, $V(x) \geq 0$ for all $x \in \mathcal{X}_f$, and there exists a terminal control law $\kappa_f : \mathcal{X}_f \to \mathcal{U}$ such that $V(f(x, \kappa_f(x))) - V(x) \leq q(x, \kappa_f(x))$ for all $x \in \mathcal{X}_f$.

A3 The set $\mathcal{X}_f$ is positively invariant for the system dynamics under $\kappa_f(x)$: $f(x, \kappa_f(x)) \in \mathcal{X}_f$ for all $x \in \mathcal{X}_f$.

A4 The terminal control $\kappa_f(x)$ satisfies the control constraints in $\mathcal{X}_f$: $\kappa_f(x) \in \mathcal{U}$ for all $x \in \mathcal{X}_f$.

A5 The sets $\mathcal{U}$ and $\mathcal{X}_f$ contain the origin of their respective spaces.

Theorem for RHC stability

Theorem (Stability of Receding Horizon Control)

Assume that conditions **A1** to **A5** are satisfied and consider the closed loop dynamics $f(x, \kappa(x))$ under RHC. The following statements hold:

i. The set $\mathcal{X}_0$ of feasible initial states is positively invariant for the closed loop.

ii. The origin is globally attractive in $\mathcal{X}_0$ for the closed loop system.

iii. If, in addition to **A1–A5**, $0 \in \operatorname{int} \mathcal{X}_0$ and $J_N^*(\cdot)$ is continuous on some neighborhood of the origin, then the origin is asymptotically stable in $\mathcal{X}_0$ for the closed loop system.

Proof of RHC stability

Main ideas of the proof

- Use shifted optimal sequence as a feasible sequence in the next step
- Use value function as a Lyapunov function

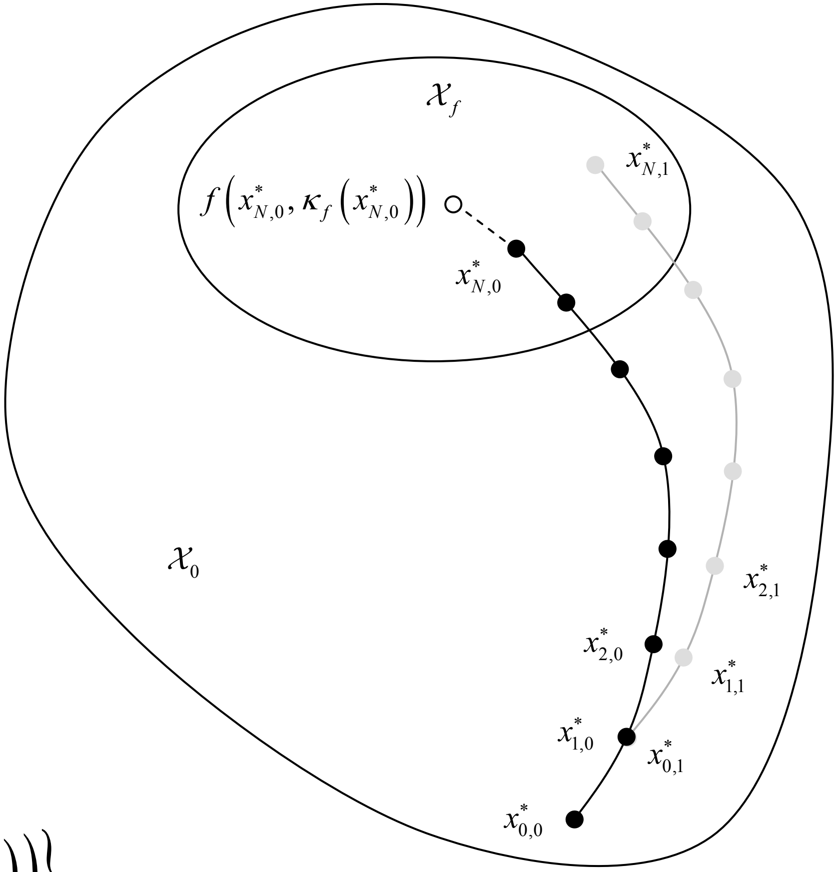

Optimal sequence at time $t = 0$

$$\mathbf{x}^*_{[0,\ldots,N],0} = \left\{x^*_{0,0}, x^*_{1,0}, \ldots, x^*_{N,0}\right\}$$

Feasible sequence at time $t = 1$

$$\mathbf{x}_{[0,\ldots,N],1} = \left\{x^*_{1,0}, \ldots, x^*_{N,0}, f\left(x^*_{N,0}, \kappa_f\left(x^*_{N,0}\right)\right)\right\}$$

Proof of RHC stability (i)

(i) Positive invariance of $\mathcal{X}_0$

Let $x_t \in \mathcal{X}_0$. At step t, using the current state x_t , the receding horizon algorithm solves the optimization problem $\mathcal{P}_N(x_t)$ to obtain the optimal control and state sequences

$$\mathbf{u}^*_{[0,\ldots,N-1],t}\left(x_t\right) \triangleq \left\{u^*_{0,t}, u^*_{1,t}, \ldots, u^*_{N-1,t}\right\}$$

$$\mathbf{x}^*_{[0,\ldots,N],t}\left(x_t\right) \triangleq \left\{x^*_{0,t}, x^*_{1,t}, \ldots, x^*_{N-1,t}, x^*_{N,t}\right\}$$

Then the actual control applied to the system at time t is the first element of $\mathbf{u}^*(x_t)$,

$$u_t = \kappa(x_t) = u_0{}^*$$

Let $x_{k+1} = f(x_k, \kappa(x_k)) = f(x_k, u_0{}^*)$ be the successor state.

A feasible (not necessarily optimal) control and state sequence for $\mathcal{P}_N(x_{t+1})$ at the next step $k+1$ are

$$\mathbf{u}_{[0,\ldots,N-1],1} = \left\{u^*_{1,0}, \ldots, u^*_{N-1,0}, \kappa_f\left(x^*_{N,0}\right)\right\}$$

$$\mathbf{x}_{[0,\ldots,N],1} = \left\{x^*_{1,0}, \ldots, x^*_{N-1,0}, x^*_{N,0}, f\left(x^*_{N,0}, \kappa_f\left(x^*_{N,0}\right)\right)\right\}$$

Clearly, the first $N-1$ elements of $\mathbf{u}_{[0,\ldots,N-1],1}$ are in $\mathcal{U}$.

Also, by **A4**, the last element of $\mathbf{u}_{[0,\ldots,N-1],1}$ is in $\mathcal{U}$ since $x^*{}_N \in \mathcal{X}_f$.

Finally, by **A3**, the terminal state $f(x^*{}_N, \kappa_f(x^*{}_N))$ also lies in $\mathcal{X}_f$.

The existence of the above feasible sequences for $x_{k+1} = f(x_k, \kappa(x_k))$ shows that $x_{k+1} \in \mathcal{X}_0$. Thus $\mathcal{X}_0$ is positively invariant for the closed loop system $x_{k+1} = f(x_k, \kappa(x_k))$.

Proof of RHC stability (ii)

(ii) Attractivity

The origin is an equilibrium point for the closed loop system $x_{k+1} = f(x_k, \kappa(x_k))$. We will use the value function $J_N^*(\cdot)$ as a Lyapunov function.

First we show that $J_N^*(\cdot)$ satisfies property (i) in Attractivity Theorem (decreases along closed loop trajectories).

Let $x_k \in \mathcal{X}_0$. The increment of the Lyapunov function, when using the RHC control law and moving from x_k to $x_{k+1} = f(x_k, \kappa(x_k))$, satisfies

$$\begin{aligned} J_N^*(x_{k+1}) - J_N^*(x_k) &= J_N\left(\mathbf{x}^*_{[0,\ldots,N],k+1}, \mathbf{u}^*_{[0,\ldots,N-1],k+1}\right) - J_N\left(\mathbf{x}^*_{[0,\ldots,N],k}, \mathbf{u}^*_{[0,\ldots,N-1],k}\right) \\ &\leq J_N\left(\mathbf{x}_{[0,\ldots,N],k+1}, \mathbf{u}_{[0,\ldots,N-1],k+1}\right) - J_N\left(\mathbf{x}^*_{[0,\ldots,N],k}, \mathbf{u}^*_{[0,\ldots,N-1],k}\right) \end{aligned}$$

since, by optimality we know that

$$J_N\left(\mathbf{x}^*_{[0,\ldots,N],k+1}, \mathbf{u}^*_{[0,\ldots,N-1],k+1}\right) \leq J_N\left(\mathbf{x}_{[0,\ldots,N],k+1}, \mathbf{u}_{[0,\ldots,N-1],k+1}\right)$$

Proof of RHC stability (ii)

Substituting the shifted feasible sequences in the objective function, we obtain

$$J_N^*\left(x_{k+1}\right)-J_N^*\left(x_k\right)\leq J_N\left(\mathbf{x}_{[0,\ldots,N],k+1},\mathbf{u}_{[0,\ldots,N-1],k+1}\right)-J_N\left(\mathbf{x}^*_{[0,\ldots,N],k},\mathbf{u}^*_{[0,\ldots,N-1],k}\right)$$

$$=-q\left(x_k,\kappa\left(x_k\right)\right)+q\left(x_N^*,\kappa_f\left(x_N^*\right)\right)$$

$$+V\left(f\left(x_N^*,\kappa_f\left(x_N^*\right)\right)\right)-V\left(x_N^*\right)$$

From **A2**, and since $x^*_N \in \mathcal{X}_f$, the sum of the last three terms on the right hand side above is less than or equal to zero. Thus,

$$J_N^*\left(x_{k+1}\right)-J_N^*\left(x_k\right)\leq -q\left(x_k,\kappa\left(x_k\right)\right)\leq -\gamma\left(\left\|x_k\right\|\right)$$

for all $x_k \in \mathcal{X}_0$. This shows that $J_N^*(\cdot)$ satisfies property (i) of Attractivity Theorem.

In addition, from Assumptions **A1** and **A2**, $J_N^*(\cdot)$ satisfies

$$J_N^*\left(x_k\right)\geq q\left(x_k,u_0^*\right)\geq \gamma\left(\left\|x_k\right\|\right)$$

for all $x_k \in \mathcal{X}_0$. Hence, from the assumption on γ, $J_N^*(x) \to \infty$ when $\|x\| \to \infty$, and therefore $J_N^*(\cdot)$ satisfies property (ii) in Attractivity Theorem.

It then follows from Attractivity Theorem that the origin is globally attractive in $\mathcal{X}_0$ for the closed loop system.

Proof of RHC stability (iii)

(iii) Asymptotic stability

Notice first that $J_N^*(0) = 0$ (since the optimal sequences corresponding to $x = 0$ have all their elements equal to zero).

Next, note from the lower bound on the previous slide and the properties of γ that $J_N^*(\cdot)$ satisfies property (ii) in Lyapunov Stability Theorem with $\mathcal{X} = \mathcal{X}_0$.

If, in addition, the origin is in the interior of $\mathcal{X}_0$ and $J_N^*(\cdot)$ is continuous on a neighborhood of the origin, then Lyapunov Stability Theorem shows that the origin is a stable equilibrium point for the closed loop system, and hence, combined with attractivity in $\mathcal{X}_0$, it is asymptotically stable in $\mathcal{X}_0$.

Example for linear systems

For linear systems with convex constraints and quadratic objective function, it is easy to find ingredients to guarantee stability of RHC.

The optimization problem has the following form in this case:

$$
\begin{aligned}
J_N^*(x_t) \triangleq\ & \min_{\mathbf{u}} \frac{1}{2}\sum_{k=0}^{N-1}\left(x_k^{\mathrm{T}} Q x_k + u_k^T R u_k\right) + V(x_N) \\
\text{subject to}\ & x_{k+1} = Ax_k + Bu_k \qquad k = 0, \ldots, N-1, \\
& u_k \in \mathcal{U} \qquad k = 0, \ldots, N-1, \\
& x_k \in \mathcal{X} \qquad k = 0, \ldots, N, \\
& x_0 = x_t, \\
& x_N \in \mathcal{X}_f \subset \mathcal{X}.
\end{aligned}
$$

Example for linear systems

Choices for the Terminal Triple

- The terminal cost is usually $V(x) = ½\, x^{\mathrm{T}}Px$, where P is the positive definite solution of the algebraic Riccati equation (ARE)

$$P = A^{\mathrm{T}}PA + Q - K^{\mathrm{T}}(R + B^{\mathrm{T}}PB)K,$$

- The terminal controller $\kappa_f(x)$ is chosen to be the optimal linear controller $\kappa_f(x) = -Kx$, where

$$K = (R + B^{\mathrm{T}}PB)^{-1}B^{\mathrm{T}}PA,$$

- The terminal constraint set is usually taken as the maximal output admissible set of the closed loop using K

$$\mathcal{X}_f = \mathcal{O}_\infty$$

These choices ensure exponential stability of the closed loop system in any arbitrarily large compact subset contained in the interior of $\mathcal{X}_0$.

Choice of terminal set for linear systems

The terminal set $\mathcal{X}_f$ is usually taken to be the maximal output admissible set $\mathcal{O}_\infty$ for the closed loop system using the local controller $\kappa_f(x) = -Kx$, defined as

$$\mathcal{O}_\infty \triangleq \left\{ x \,\middle|\, K\left(A - BK\right)^k x \in \mathcal{U} \text{ and } \left(A - BK\right)^k x \in \mathcal{X} \text{ for } k = 0, 1, \ldots, \infty \right\}$$

$\mathcal{O}_\infty$ is the maximal positively invariant set for the system $x_{k+1} = (A - BK)\, x_k$ in which constraints are satisfied.

With the above choice for the terminal triple $(\mathcal{X}_f, \kappa_f, V)$, conditions **A1–A5** of the main stability theorem are readily established:

- Conditions **A1** and **A5** are satisfied from the assumptions on problem $\mathcal{P}_N(x)$.
- **A3** holds since the set $\mathcal{X}_f = \mathcal{O}_\infty$ is positively invariant for the system $x_{k+1} = (A - BK)\, x_k$.
- **A4** holds since the terminal control satisfies the control constraints in $\mathcal{X}_f$.

Property A2 of the value function for linear systems

- Condition **A2** is also satisfied since using solution of the ARE, direct calculation yields that $V(x) = ½\, x^{\mathrm{T}}Px$ satisfies

$$
\begin{aligned}
V\left(f\left(x,\kappa_f(x)\right)\right)-V(x) &= \frac{1}{2}x^{\mathrm{T}}(A-BK)^{\mathrm{T}}P(A-BK)x-\frac{1}{2}x^{\mathrm{T}}Px \\
&= \frac{1}{2}x^{\mathrm{T}}\left(A^{\mathrm{T}}PA-P-K^{\mathrm{T}}B^{\mathrm{T}}PA-A^{\mathrm{T}}PBK+K^{\mathrm{T}}B^{\mathrm{T}}PBK\right)x \\
&= \frac{1}{2}x^{\mathrm{T}}\left(A^{\mathrm{T}}PA-P-K^{\mathrm{T}}\left(R+B^{\mathrm{T}}PB\right)^{-1}K-K^{\mathrm{T}}RK\right)x \\
&= -\frac{1}{2}x^{\mathrm{T}}Qx-\frac{1}{2}x^{\mathrm{T}}K^{\mathrm{T}}RKx \\
&= -q\left(x,\kappa_f(x)\right)
\end{aligned}
$$

We have verified conditions **A1–A5** of the main stability theorem, which establishes global attractivity of the origin.

Various constraints to ensure RHC stability

- **Infinite Output / Prediction Horizon**

  $N \rightarrow \infty$ (Rawlings and Muske, 1993)

- **End constraint**

  $x(t+N|t) = 0$ (Kwon and Pearson, 1977)

  (Keerthi and Gilbert 1988)

- **Relaxed terminal constraint**

  $x(t+N|t) \in \Omega$ (Scokaert and Rawlings, 1996)

- **Contraction Constraint constraint**

  $||x(t+1|t)|| \leq \alpha||x(t)||, \; \alpha < 1$ (Polak and Yang, 1993)

  (Bemporad, 1998)

All the proofs use the value function $V(t) = \min_{\mathbf{u}} J(\mathbf{u}, t)$ as a Lyapunov function.

Continuous time, unconstrained nonlinear RHC

Solve finite time optimization over *T* seconds and implement first δ seconds

$$J_T^*\left(x(t),u(\cdot)\right)=\min_{u(\cdot)}\int_t^{t+T} q(x(\tau),u(\tau))\,d\tau+V(x(t+T))$$

$$\dot{x}=f(x,u)$$

Incremental cost

Terminal cost

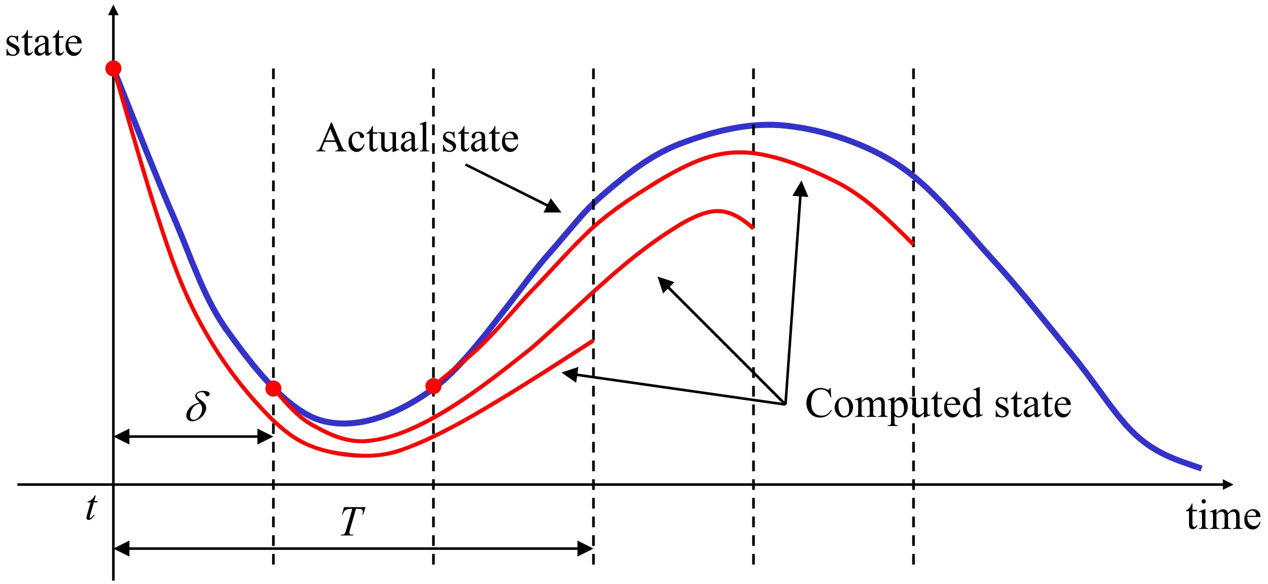

Murray, Hauser et al
SEC chapter (IEEE, 2002)

Requires that computation time be small relative to time horizons

- Initial implementation in process control, where time scales are fairly slow.
- Real-time trajectory generation enables implementation on faster systems.

Stability of Receding Horizon Control

RHC can destabilize systems if not done properly

- For properly chosen cost functions, get stability with T sufficiently large.
- For shorter horizons, V has to be chosen properly to avoid instability.

The choice of the terminal cost

- Best choice would be $V(x) = J_{\infty}^{*}(x)$ such that the optimal finite and infinite costs are the same. (Not possible, if the optimal value function were available there would be no need to solve a trajectory optimization problem.)
- The terminal cost must account for the discarded tail by ensuring that the origin can be reached from the terminal state $x(t+T)$ in an efficient manner as measured by q.

  One way to do this is to use an appropriate control Lyapunov function (CLF).

Control Lyapunov Function (CLF)

Definition

A control Lyapunov function (CLF) is a C^1, proper, positive definite function $V : R^n \to R_+$ such that

$$\inf_u \left[\dot{V}(x,u) \right] \le 0$$

where

$$\dot{V}(x,u) = \frac{dV}{dx}\frac{dx}{dt} = \frac{dV}{dx} f(x,u)$$

denotes the directional derivative in direction $f(x, u)$.

Meaning

If it is possible to make the derivative negative at every point by an appropriate choice of u then we can stabilize the system with V as a Lyapunov function for the closed loop.

It can be shown that the existence of a CLF is equivalent to the existence of an asymptotically stabilizing control law u = k(x).

Stability of Receding Horizon Control

Theorem (Jadbabaie & Hauser, 2002)

Suppose that the terminal cost $V(x)$ is a control Lyapunov function such that

$$\min_u \left(\dot{V} + q\right)\left(x, u\right) \leq 0$$

for each x in $\Omega_r = \{x : V(x) < r^2\}$, for some $r > 0$. Then, for every $T > 0$ and δ in $(0, T\,]$, the resulting receding horizon trajectories go to zero exponentially fast.

Remarks

- Earlier approach used terminal trajectory constraints, hard to implement in real-time.
- CLF terminal cost is difficult to find in general, but LQR-based solution at equilibrium point often works well - choose $V = x^T P\, x$ where P = Riccati solution.

Main ingredient of the proof

Denote with $x^u(\tau; x)$ the state trajectory at time τ starting from initial state x and applying a control trajectory $u(\cdot)$.

Let $(x_T^*, u_T^*)(\cdot\, ; x)$ denote an optimal trajectory of the finite horizon optimal control problem with horizon T.

Assume $x_T^*(T; x) \in \Omega_r = \{x : V(x) < r^2\}$, for some $r > 0$. Then, for each $\delta \in [0, T\,]$, the optimal cost from $x_T^*(\delta; x)$ satisfies

$$J_T^*\left(x_T^*\left(\delta; x\right)\right) \le J_T^*\left(x\right) - \int_0^\delta q\left(x_T^*\left(\tau; x\right), u_T^*\left(\tau; x\right)\right) d\tau$$

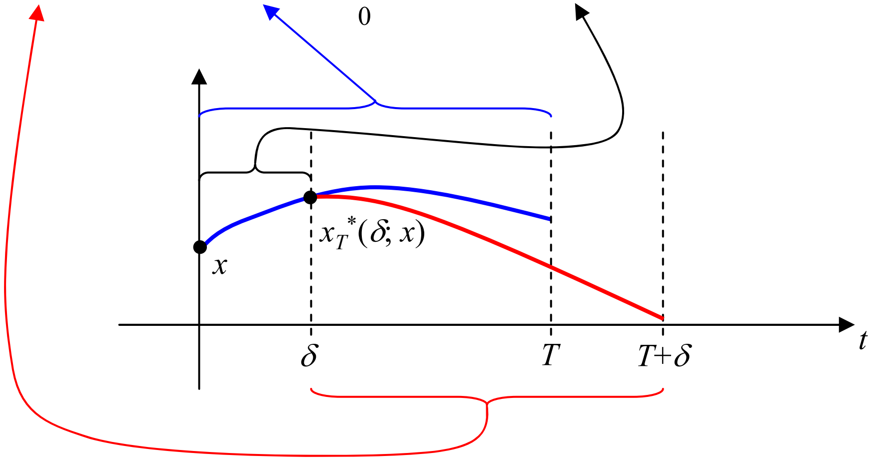

Proof sketch

Let $\left(\tilde{x}(t),\tilde{u}(t)\right)$, $t \in [0, 2T]$, be the trajectory obtained by concatenating $(x_T^{\;*}, u_T^{\;*})(t\,;x)$, $t \in [0, T]$, and $(x^k, u^k)(t - T\,; x_T^{\;*}(T\,; x))$, $t \in [T, 2T]$, which are closed-loop trajectories corresponding to a feedback law $u = k(x)$ such that $\left(\dot{V}+q\right)\left(x,k\left(x\right)\right)\le 0$.

Consider the cost of using $\tilde{u}\left(\cdot\right)$ for T seconds at the initial state $x_T^{\;*}(\delta\,; x))$, $\delta \in [0, T]$

$$\begin{aligned}
J_T\left(x_T^*\left(\delta;x\right),\tilde{u}\left(\cdot\right)\right) &= \int_{\delta}^{T+\delta} q\left(\tilde{x}\left(\tau\right),\tilde{u}\left(\tau\right)\right)d\tau+V\left(\tilde{x}\left(T+\delta\right)\right)\\
&= J_T^*\left(x\right)-\int_0^{\delta} q\left(x_T^*\left(\tau;x\right),u_T^*\left(\tau;x\right)\right)d\tau-V\left(x_T^*\left(T;x\right)\right)\\
&\quad+\int_T^{T+\delta} q\left(\tilde{x}\left(\tau\right),\tilde{u}\left(\tau\right)\right)d\tau+V\left(\tilde{x}\left(T+\delta\right)\right)\\
&\le J_T^*\left(x\right)-\int_0^{\delta} q\left(x_T^*\left(\tau;x\right),u_T^*\left(\tau;x\right)\right)d\tau
\end{aligned}$$

where we have used the facts that $q\left(\tilde{x}\left(\tau\right),\tilde{u}\left(\tau\right)\right)\le -\dot{V}\left(\tilde{x}\left(\tau\right),\tilde{u}\left(\tau\right)\right)$ for all $\tau\in\left[T,2T\right]$

and due to optimality $J_T^*\left(x_T^*\left(\delta;x\right)\right)\le J_T\left(x_T^*\left(\delta;x\right),\tilde{u}\left(\cdot\right)\right)$

Theory: RHC + CLF

Results to Date

- Characterization of exponential region of attraction based on CLF level sets
- Proof of stability with incremental improvement in cost
- Proof of stability with *general* class of terminal costs, for sufficiently long horizons

Model Predictive Framework Built on Sound Theoretical Basis

Missing pieces

- Inclusion of constraints (via CLF)
- Including robustness and disturbance rejection
- Theoretical work on receding horizon trajectory *tracking*
- Extensions to multi-vehicle, and distributed cases.
- Online adaptation of cost, dynamics, and constraints.
- Lower bounds on the horizon length in the case of a general positive-definite terminal cost.