


Model Predictive Control Toolbox™

Getting Started Guide

R2013a

*Alberto Bemporad
Manfred Morari
N. Lawrence Ricker*



MATLAB®

How to Contact MathWorks



www.mathworks.com Web
comp.soft-sys.matlab Newsgroup
www.mathworks.com/contact_TS.html Technical Support



suggest@mathworks.com Product enhancement suggestions
bugs@mathworks.com Bug reports
doc@mathworks.com Documentation error reports
service@mathworks.com Order status, license renewals, passcodes
info@mathworks.com Sales, pricing, and general information



508-647-7000 (Phone)



508-647-7001 (Fax)



The MathWorks, Inc.
3 Apple Hill Drive
Natick, MA 01760-2098

For contact information about worldwide offices, see the MathWorks Web site.

Model Predictive Control Toolbox™ Getting Started Guide

© COPYRIGHT 2005–2013 by The MathWorks, Inc.

The software described in this document is furnished under a license agreement. The software may be used or copied only under the terms of the license agreement. No part of this manual may be photocopied or reproduced in any form without prior written consent from The MathWorks, Inc.

FEDERAL ACQUISITION: This provision applies to all acquisitions of the Program and Documentation by, for, or through the federal government of the United States. By accepting delivery of the Program or Documentation, the government hereby agrees that this software or documentation qualifies as commercial computer software or commercial computer software documentation as such terms are used or defined in FAR 12.212, DFARS Part 227.72, and DFARS 252.227-7014. Accordingly, the terms and conditions of this Agreement and only those rights specified in this Agreement, shall pertain to and govern the use, modification, reproduction, release, performance, display, and disclosure of the Program and Documentation by the federal government (or other entity acquiring for or through the federal government) and shall supersede any conflicting contractual terms or conditions. If this License fails to meet the government's needs or is inconsistent in any respect with federal procurement law, the government agrees to return the Program and Documentation, unused, to The MathWorks, Inc.

Trademarks

MATLAB and Simulink are registered trademarks of The MathWorks, Inc. See www.mathworks.com/trademarks for a list of additional trademarks. Other product or brand names may be trademarks or registered trademarks of their respective holders.

Patents

MathWorks products are protected by one or more U.S. patents. Please see www.mathworks.com/patents for more information.

Revision History

October 2004	First printing	New for Version 2.1 (Release 14SP1)
March 2005	Online only	Revised for Version 2.2 (Release 14SP2)
September 2005	Online only	Revised for Version 2.2.1 (Release 14SP3)
March 2006	Online only	Revised for Version 2.2.2 (Release 2006a)
September 2006	Online only	Revised for Version 2.2.3 (Release 2006b)
March 2007	Online only	Revised for Version 2.2.4 (Release 2007a)
September 2007	Online only	Revised for Version 2.3 (Release 2007b)
March 2008	Online only	Revised for Version 2.3.1 (Release 2008a)
October 2008	Online only	Revised for Version 3.0 (Release 2008b)
March 2009	Online only	Revised for Version 3.1 (Release 2009a)
September 2009	Online only	Revised for Version 3.1.1 (Release 2009b)
March 2010	Online only	Revised for Version 3.2 (Release 2010a)
September 2010	Online only	Revised for Version 3.2.1 (Release 2010b)
April 2011	Online only	Revised for Version 3.3 (Release 2011a)
September 2011	Online only	Revised for Version 4.0 (Release 2011b)
March 2012	Online only	Revised for Version 4.1 (Release 2012a)
September 2012	Online only	Revised for Version 4.1.1 (Release 2012b)
March 2013	Online only	Revised for Version 4.1.2 (Release 2013a)

Introduction

1

Product Description	1-2
Key Features	1-2
Using the Documentation	1-3
Related Products	1-4
Acknowledgments	1-5
Bibliography	1-6

Building Models

2

Plant Model	2-2
Plant Inputs and Outputs	2-3
Inputs	2-3
Outputs	2-3
Construct Linear, Time Invariant (LTI) Models	2-4
Transfer Function Models	2-4
Zero/Pole/Gain Models	2-5
State-Space Models	2-5
LTI Object Properties	2-7
LTI Model Characteristics	2-11
Specify Multi-Input Multi-Output (MIMO) Plants	2-12
CSTR Model	2-15

Linearize Simulink Models	2-17
Linearization Using MATLAB Code	2-17
Linearization Using Linear Analysis Tool in Simulink Control Design	2-20
Identify Plant from Data	2-25
Ways to Design MPC Controllers for Identified Plant Models	2-27
Input Output Structure of Identified Model	2-28
Design Controller for Identified Plant	2-29
Design Controller Using Noise Model of Identified Model	2-31
Working with Impulse-Response Models	2-33
Bibliography	2-34

Designing Controllers Using the Design Tool GUI

3

Design Controller Using the Design Tool	3-2
Start the Design Tool	3-2
Load a Plant Model	3-3
Navigate Using the Tree View	3-6
Perform Linear Simulations	3-10
Customize Response Plots	3-18
Change Controller Settings	3-23
Define Soft Output Constraints	3-42
Save Your Work	3-46
Load Your Saved Work	3-48
Test Controller Robustness	3-50
Plant Model Perturbation	3-50

Simulation Tests	3-50
Design Controller for Plant with Delays	3-53
Importing the Plant Model	3-53
Specifying Controller Horizons	3-54
Design Controller for Nonsquare Plant	3-58
More Outputs Than Manipulated Variables	3-58
More Manipulated Variables Than Outputs	3-59
Nonlinear Plants	3-60
MPC Controller Block	3-60
Initiating the Controller Design	3-61
Validating the Linearized Model	3-63
Modifying the Linearized Model	3-65
Linear Simulation Tests	3-67
Nonlinear Simulation Tests	3-68
Modifying the Controller Using the Design Tool	3-70
Exiting the Design Tool	3-70

Designing Controllers Using the Command Line

4

Design Controller at the Command Line	4-2
Create a Controller Object	4-2
View and Alter Controller Properties	4-3
Perform Linear Simulations	4-6
Save Calculated Results	4-7
Simulate Controller with Nonlinear Plant	4-9
Nonlinear CSTR Application	4-9
Example Code for Successive Linearization	4-10
CSTR Results and Discussion	4-12
Control Based On Multiple Plant Models	4-15
A Two-Model Plant	4-15
Designing the Two Controllers	4-17
Simulating Controller Performance	4-18

Compute Steady-State Gain	4-24
Extract Controller	4-25
Bibliography	4-27

Designing and Testing Controllers in Simulink

5

Design Controller in Simulink	5-2
Test an Existing Controller	5-17
Schedule Controllers at Multiple Operating Points ...	5-18
A Two-Model Plant	5-18
Animation of the Multi-Model Example	5-19
Designing the Two Controllers	5-21
Simulating Controller Performance	5-21

Index

Introduction

- “Product Description” on page 1-2
- “Using the Documentation” on page 1-3
- “Acknowledgments” on page 1-5
- “Bibliography” on page 1-6

Product Description

Design and simulate model predictive controllers

Model Predictive Control Toolbox™ provides tools for systematically analyzing, designing, and tuning model predictive controllers. You can design and simulate model predictive controllers using functions in MATLAB® or blocks in Simulink®. You can set and modify the predictive model, control and prediction horizons, input and output constraints, and weights. The toolbox enables you to diagnose issues that could lead to run-time failures and provides advice on changing weights and constraints to improve performance and robustness. By running different scenarios in linear and nonlinear simulations, you can evaluate controller performance. You can adjust controller performance as it runs by tuning weights and varying constraints. For rapid prototyping and embedded system design, the toolbox supports C-code generation.

Key Features

- Design and simulation of model predictive controllers in MATLAB and Simulink
- Customization of constraints and weights with advisory tools for improved performance and robustness
- Control of plants over a range of operating conditions using multiple model predictive controllers with bumpless control transfer
- Run-time adjustment of controller performance through constraint and weight changes
- Specialized model predictive control quadratic programming (QP) solver optimized for speed, efficiency, and robustness
- Support for C-code generation with Simulink Coder™

Using the Documentation

If you have limited experience with MATLAB or Model Predictive Control Toolbox software, read this guide first. It shows how to:

- Define your plant using Control System Toolbox™ modeling tools (LTI transfer function and state space models).
- Derive a linear plant model from a nonlinear Simulink representation.
- Design Model Predictive Control for your plant using `mpctool`, the graphical user interface (GUI), or special commands.
- Simulate Model Predictive Control performance using `mpctool`, Simulink, or commands.

If you have experience with an earlier Model Predictive Control Toolbox release, we advise you to read this document to familiarize yourself with the many new features and the new command syntax (the earlier syntax is still available, but the underlying code is no longer supported).

If you need more details, see the online documentation. To access it from the MATLAB prompt, type:

```
doc
```

When the help dialog box appears, select **Model Predictive Control Toolbox** in the **Contents** pane. This displays a roadmap with links to the available documentation components. Briefly, these are:

- Getting Started. The online version of this document.
- MPC Problem Setup. Mathematical details of the Model Predictive Control Toolbox algorithm and user specifications required for controller design.
- MPC Simulink Library. Describes the Model Predictive Controller Block and its use within Simulink.
- Case-Study Examples. Example toolbox applications.
- The Design Tool. Reference manual for `mpctool`, the GUI.

- **Functions.** Reference manual describing each Model Predictive Control Toolbox function (used for controller design and simulation in MATLAB commands and scripts).
- **Blocks.** Reference manual describing the Model Predictive Control Toolbox blocks (used for controller design and simulation in Simulink).
- **Object Reference.** Details of the controller object that represents a complete controller design.
- **Release Notes.** Summarizes major features of this release, known limitations, etc.

Related Products

MathWorks provides other products that complement and enhance the Model Predictive Control Toolbox functionality. For more information, see www.mathworks.com.

Acknowledgments

MathWorks would like to acknowledge the following contributors to Model Predictive Control Toolbox.

Alberto Bemporad

Professor of Control Systems, IMT Institute for Advanced Studies Lucca, Italy. Research interests include model predictive control, hybrid systems, optimization algorithms, and applications to automotive, aerospace, and energy systems. Fellow of the IEEE[®]. Author of the Model Predictive Control Simulink library and commands.

Manfred Morari

Professor at the Automatic Control Laboratory and former Head of Department of Information Technology and Electrical Engineering, ETH Zurich, Switzerland. Research interests include model predictive control, hybrid systems, and robust control. Fellow of the IEEE, AIChE, and IFAC. Co-author of the first version of the toolbox.

N. Lawrence Ricker

Professor of Chemical Engineering, University of Washington, Seattle, USA. Research interests include model predictive control and process optimization. Author of the quadratic programming solver and graphical user interface.

Bibliography

[1] Allgower, F., and A. Zheng, *Nonlinear Model Predictive Control*, Springer-Verlag, 2000.

[2] Camacho, E. F., and C. Bordons, *Model Predictive Control*, Springer-Verlag, 1999.

[3] Kouvaritakis, B., and M. Cannon, *Non-Linear Predictive Control: Theory & Practice*, IEE Publishing, 2001.

[4] Maciejowski, J. M., *Predictive Control with Constraints*, Pearson Education POD, 2002.

[5] Prett, D., and C. Garcia, *Fundamental Process Control*, Butterworths, 1988.


[6] Rossiter, J. A., *Model-Based Predictive Control: A Practical Approach*, CRC Press, 2003.

Building Models

- “Plant Model” on page 2-2
- “Plant Inputs and Outputs” on page 2-3
- “Construct Linear, Time Invariant (LTI) Models” on page 2-4
- “Specify Multi-Input Multi-Output (MIMO) Plants” on page 2-12
- “CSTR Model” on page 2-15
- “Linearize Simulink Models” on page 2-17
- “Identify Plant from Data” on page 2-25
- “Ways to Design MPC Controllers for Identified Plant Models” on page 2-27
- “Design Controller for Identified Plant” on page 2-29
- “Design Controller Using Noise Model of Identified Model” on page 2-31
- “Working with Impulse-Response Models” on page 2-33
- “Bibliography” on page 2-34

Plant Model

The *plant* is the system (process or device) you intend to control. The following figure shows a schematic example.



Plant with Input and Output Signals

A Model Predictive Control Toolbox design requires a *plant model*, which defines the mathematical relationship between the plant inputs and outputs. The controller uses it to predict plant behavior.

The toolbox software requires the model to be linear, time invariant (LTI). You can define such a model as follows:

- Create a transfer function, state space, or zero/pole/gain model using methods provided by the Control System Toolbox software
- Derive it from plant data using, e.g., methods provided by System Identification Toolbox™ software
- Derive it by linearizing a Simulink model

This topic illustrates each of these approaches. Control System Toolbox, Simulink, and System Identification Toolbox documentation provides additional examples and details.

Plant Inputs and Outputs

Inputs

The *plant inputs* are the independent variables affecting the plant. As shown in Plant with Input and Output Signals on page 2-2, there are three types:

Measured disturbances

The controller can't adjust them, but uses them for feedforward compensation.

Manipulated variables

The controller adjusts these in order to achieve its goals.

Unmeasured disturbances

These are independent inputs of which the controller has no direct knowledge, and for which it must compensate.

Outputs

The *plant outputs* are the dependent variables (outcomes) you wish to control or monitor. As shown in Plant with Input and Output Signals on page 2-2, there are two types:

Measured outputs

The controller uses these to estimate unmeasured quantities and as feedback on the success of its adjustments.

Unmeasured outputs

The controller estimates these based on available measurements and the plant model. The controller can also hold unmeasured outputs at setpoints or within constraint boundaries.

You must specify the input and output types when designing the controller. See "Input and Output Types" on page 2-9 for more details.

Construct Linear, Time Invariant (LTI) Models

In this section...
“Transfer Function Models” on page 2-4
“Zero/Pole/Gain Models” on page 2-5
“State-Space Models” on page 2-5
“LTI Object Properties” on page 2-7
“LTI Model Characteristics” on page 2-11

Model Predictive Control Toolbox software supports the same LTI model formats as does Control System Toolbox software. You can use whichever is most convenient for your application. It's also easy to convert from one format to another.

The following topics describe the three model formats and the commands used to construct them:

- “Transfer Function Models” on page 2-4
- “Zero/Pole/Gain Models” on page 2-5
- “State-Space Models” on page 2-5
- “LTI Object Properties” on page 2-7
- “Specify Multi-Input Multi-Output (MIMO) Plants” on page 2-12
- “LTI Model Characteristics” on page 2-11

For more details, see the Control System Toolbox documentation.

Transfer Function Models

A transfer function (TF) relates a particular input/output pair. For example, if $u(t)$ is a plant input and $y(t)$ is an output, the transfer function relating them might be:

$$\frac{Y(s)}{U(s)} = G(s) = \frac{s+2}{s^2+s+10} e^{-1.5s}$$

This TF consists of a *numerator* polynomial, $s+2$, a *denominator* polynomial, s^2+s+10 , and a delay, which is 1.5 time units here. You can define G using Control System Toolbox `tf` function:

```
Gtf1 = tf([1 2], [1 1 10], 'OutputDelay', 1.5)
```

Control System Toolbox software builds and displays it as follows:

Transfer function:

$$\exp(-1.5*s) * \frac{s + 2}{s^2 + s + 10}$$

Zero/Pole/Gain Models

Like the TF, the zero/pole/gain (ZPK) format relates an input/output pair. The difference is that the ZPK numerator and denominator polynomials are factored, as in

$$G(s) = 2.5 \frac{s + 0.45}{(s + 0.3)(s + 0.1 + 0.7i)(s + 0.1 - 0.7i)}$$

(zeros and/or poles are complex numbers in general).

You define the ZPK model by specifying the zero(s), pole(s), and gain as in

```
Gzpk1 = zpk(-0.45, [-0.3, -0.1+0.7*i, -0.1-0.7*i], 2.5)
```


State-Space Models

The state-space format is convenient if your model is a set of LTI differential and algebraic equations. For example, consider the following linearized model of a continuous stirred-tank reactor (CSTR) involving an exothermic (heat-generating) reaction [9]

$$\frac{dC'_A}{dt} = a_{11}C'_A + a_{12}T' + b_{11}T'_c + b_{12}C'_{Ai}$$

$$\frac{dT'}{dt} = a_{21}C'_A + a_{22}T' + b_{21}T'_c + b_{22}C'_{Ai}$$

where C_A is the concentration of a key reactant, T is the temperature in the reactor, T_c is the coolant temperature, C_{Ai} is the reactant concentration in the reactor feed, and a_{ij} and b_{ij} are constants. See the process schematic in CSTR Schematic on page 2-6. The primes (e.g., C'_A) denote a deviation from the nominal steady-state condition at which the model has been linearized.



CSTR Schematic

Measurement of reactant concentrations is often difficult, if not impossible. Let us assume that T is a measured output, C_A is an unmeasured output, T_c is a manipulated variable, and C_{Ai} is an unmeasured disturbance.

The model fits the general state-space format

$$\frac{dx}{dt} = Ax + Bu$$

$$y = Cx + Du$$

where

$$x = \begin{bmatrix} C'_A \\ T' \end{bmatrix} \quad u = \begin{bmatrix} T'_c \\ C'_{Ai} \end{bmatrix} \quad y = \begin{bmatrix} T' \\ C'_A \end{bmatrix},$$

$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \quad B = \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} \quad C = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad D = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$$

The following code shows how to define such a model for some specific values of the a_{ij} and b_{ij} constants:

```

A = [-0.0285 -0.0014
      -0.0371 -0.1476];
B = [-0.0850  0.0238
      0.0802  0.4462];
C = [0 1
      1 0];
D = zeros(2,2);
CSTR = ss(A,B,C,D);

```

This defines a *continuous-time* state-space model. If you do not specify a sampling period, a default sampling value of zero applies. You can also specify discrete-time state-space models. You can specify delays in both continuous-time and discrete-time models. For more information, see the Control System Toolbox documentation.

Note In the CSTR example, the D matrix is zero and the output does not instantly respond to change in the input. The Model Predictive Control Toolbox software prohibits direct (instantaneous) feedthrough from a manipulated variable to an output. For example, the CSTR model could include direct feedthrough from the unmeasured disturbance, C_{A_i} , to either C_A or T but direct feedthrough from T_c to either output would violate this restriction. If the model had direct feedthrough from T_c , you can add a small delay at this input to circumvent the problem.

LTI Object Properties

The `ss` function in the last line of the above code creates a state space model, CSTR, which is an *LTI object*. The `tf` and `zpk` commands described in “Transfer Function Models” on page 2-4 and “Zero/Pole/Gain Models” on page 2-5 also create LTI objects. Such objects contain the model parameters as well as optional properties.

LTI Properties for the CSTR Example

The following code sets some of the CSTR model’s optional properties:

```

CSTR.InputName = {'T_c', 'C_A_i'};
CSTR.OutputName = {'T', 'C_A'};
CSTR.StateName = {'C_A', 'T'};

```

```
CSTR.InputGroup.MV = 1;
CSTR.InputGroup.UD = 2;
CSTR.OutputGroup.MO = 1;
CSTR.OutputGroup.UO = 2;
CSTR
```

The first three lines specify labels for the input, output and state variables. The next four specify the signal type for each input and output. The designations MV, UD, MO, and UO mean *manipulated variable*, *unmeasured disturbance*, *measured output*, and *unmeasured output*. (See “Plant Inputs and Outputs” on page 2-3 for definitions.) For example, the code specifies that input 2 of model CSTR is an unmeasured disturbance. The last line causes the LTI object to be displayed, generating the following lines in the MATLAB Command Window:

```
a =
      C_A      T
C_A -0.0285 -0.0014
T    -0.0371 -0.1476
```

```
b =
      T_c      C_Ai
C_A -0.085  0.0238
T    0.0802 0.4462
```

```
c =
      C_A      T
T      0      1
C_A    1      0
```

```
d =
      T_c      C_Ai
T      0      0
C_A    0      0
```

```
Input groups:
      Name      Channels
      MV        1
```

```

UD          2

```

Output groups:

```

Name      Channels
MO        1
UO        2

```

Continuous-time model

Input and Output Names

The optional `InputName` and `OutputName` properties affect the model displays, as in the above example. The toolbox also uses the `InputName` and `OutputName` properties to label plots and tables. In that context, the underscore character causes the next character to be displayed as a subscript.

Input and Output Types

General Case. As mentioned in “Plant Inputs and Outputs” on page 2-3, Model Predictive Control Toolbox software supports three input types and two output types. In a Model Predictive Control Toolbox design, designation of the input and output types determines the controller dimensions and has other important consequences.

For example, suppose your plant structure were as follows:

Plant Inputs	Plant Outputs
Two manipulated variables (MVs)	Three measured outputs (MOs)
One measured disturbance (MD)	Two unmeasured outputs (UOs)
Two unmeasured disturbances (UDs)	

The resulting controller has four inputs (the three MOs and the MD) and two outputs (the MVs). It includes feedforward compensation for the measured disturbance, and assumes that you wanted to include the unmeasured disturbances and outputs as part of the regulator design.

If you didn't want a particular signal to be treated as one of the above types, you could do one of the following:

- Eliminate the signal before using the model in controller design.
- For an output, designate it as unmeasured, then set its weight to zero (see "Output Weights" on page 3-26).
- For an input, designate it as an unmeasured disturbance, then define a custom state estimator that ignores the input (see "Disturbance Modeling and Estimation" on page 3-35).

Note By default, the toolbox assumes that unspecified plant inputs are manipulated variables, and unspecified outputs are measured. Thus, if you didn't specify signal types in the above example, the controller would have four inputs (assuming all plant outputs were measured) and five outputs (assuming all plant inputs were manipulated variables).

For model CSTR, default Model Predictive Control Toolbox assumptions are incorrect. You must set its `InputGroup` and `OutputGroup` properties, as illustrated in the above code, or modify the default settings when you load the model into the design tool.

Use `setmpcsignals` to make type definition. For example

```
CSTR = setmpcsignals(CSTR, 'UD', 2, 'UO', 2);
```

sets `InputGroup` and `OutputGroup` to the same values as in the previous example. The CSTR display would then include the following lines:

Input groups:

Name	Channels
Unmeasured	2
Manipulated	1

Output groups:

Name	Channels
------	----------


```

Unmeasured      2
Measured        1

```

Notice that `setmpcsignals` sets unspecified inputs to Manipulated and unspecified outputs to Measured.

See the Control System Toolbox documentation for additional information on LTI object properties.

LTI Model Characteristics

Control System Toolbox software provides functions for analyzing LTI models. Some of the more commonly used are listed below. Type the example code at the MATLAB prompt to see how they work for the CSTR example.

Example	Intended Result
<code>dcgain(CSTR)</code>	Calculate gain matrix for the CSTR model's input/output pairs.
<code>impulse(CSTR)</code>	Graph CSTR model's unit-impulse response.
<code>ltiview(CSTR)</code>	Open the LTI Viewer with the CSTR model loaded. You can then display model characteristics by making menu selections.
<code>pole(CSTR)</code>	Calculate CSTR model's poles (to check stability, etc.).
<code>step(CSTR)</code>	Graph CSTR model's unit-step response.
<code>zero(CSTR)</code>	Compute CSTR model's transmission zeros.

Specify Multi-Input Multi-Output (MIMO) Plants

Most Model Predictive Control Toolbox applications involve plants having multiple inputs and outputs. Model CSTR described in “CSTR Model” on page 2-15 is a MIMO plant, and the state space format extends naturally from single-input single-output (SISO) to MIMO plants.

You can also use the `tf` and `zpk` functions to build a MIMO plant model. For example, consider the following model of a distillation column [11], which has been used in many advanced control studies:

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} \frac{12.8e^{-s}}{16.7s+1} & \frac{-18.9e^{-3s}}{21.0s+1} & \frac{3.8e^{-8.1s}}{14.9s+1} \\ \frac{6.6e^{-7s}}{10.9s+1} & \frac{-19.4e^{-3s}}{14.4s+1} & \frac{4.9e^{-3.4s}}{13.2s+1} \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix}$$

Outputs y_1 and y_2 represent measured product purities. The control objective is to hold each at specified setpoints. To do so, the controller manipulates inputs u_1 and u_2 , the flow rates of reflux and reboiler steam, respectively. Input u_3 is a measured feed flow rate disturbance.

The model consists of six transfer functions, one for each input/output pair. Each transfer function is the first-order-plus-delay form often used by process control engineers.

The following code shows how to define the distillation column model for use in the toolbox:

```
g11 = tf( 12.8, [16.7 1], 'IODELAY', 1.0);
g12 = tf(-18.9, [21.0 1], 'IODELAY', 3.0);
g13 = tf( 3.8, [14.9 1], 'IODELAY', 8.1);
g21 = tf( 6.6, [10.9 1], 'IODELAY', 7.0);
g22 = tf(-19.4, [14.4 1], 'IODELAY', 3.0);
g23 = tf( 4.9, [13.2 1], 'IODELAY', 3.4);
DC = [g11 g12 g13
      g21 g22 g23];
DC.InputName = {'Reflux Rate', 'Steam Rate', 'Feed Rate'};
DC.OutputName = {'Distillate Purity', 'Bottoms Purity'};
DC = setmpcsignals(DC, 'MD', 3)
```

The code defines the individual transfer functions, and then forms a matrix in which each row contains the transfer functions for a particular output, and each column corresponds to a particular input. The code also sets the signal names and designates the third input as a measured disturbance. The resulting LTI object display is as follows:

Transfer function from input "Reflux Rate" to output...

$$\text{Distillate Purity: } \exp(-1*s) * \frac{12.8}{16.7 s + 1}$$

$$\text{Bottoms Purity: } \exp(-7*s) * \frac{6.6}{10.9 s + 1}$$

Transfer function from input "Steam Rate" to output...

$$\text{Distillate Purity: } \exp(-3*s) * \frac{-18.9}{21 s + 1}$$

$$\text{Bottoms Purity: } \exp(-3*s) * \frac{-19.4}{14.4 s + 1}$$

Transfer function from input "Feed Rate" to output...

$$\text{Distillate Purity: } \exp(-8.1*s) * \frac{3.8}{14.9 s + 1}$$

$$\text{Bottoms Purity: } \exp(-3.4*s) * \frac{4.9}{13.2 s + 1}$$

Input groups:

Name	Channels
Measured	3
Manipulated	1,2

Output groups:

Name	Channels
------	----------

Measured 1,2


CSTR Model

The linearized model of a continuous stirred-tank reactor (CSTR) involving an exothermic (heat-generating) reaction is represented by the following differential equations:[9]

$$\frac{dC'_A}{dt} = a_{11}C'_A + a_{12}T' + b_{11}T'_c + b_{12}C'_{Ai}$$

$$\frac{dT'}{dt} = a_{21}C'_A + a_{22}T' + b_{21}T'_c + b_{22}C'_{Ai}$$

where C_A is the concentration of a key reactant, T is the temperature in the reactor, T_c is the coolant temperature, C_{Ai} is the reactant concentration in the reactor feed, and a_{ij} and b_{ij} are constants. The primes (e.g., C'_A) denote a deviation from the nominal steady-state condition at which the model has been linearized.



Measurement of reactant concentrations is often difficult, if not impossible. Let us assume that T is a measured output, C_A is an unmeasured output, T_c is a manipulated variable, and C_{Ai} is an unmeasured disturbance.

The model fits the general state-space format

$$\frac{dx}{dt} = Ax + Bu$$

$$y = Cx + Du$$

where

$$x = \begin{bmatrix} C'_A \\ T' \end{bmatrix} \quad u = \begin{bmatrix} T'_c \\ C'_{Ai} \end{bmatrix} \quad y = \begin{bmatrix} T' \end{bmatrix},$$

$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \quad B = \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} \quad C = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad D = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$$

The following code shows how to define such a model for some specific values of the a_{ij} and b_{ij} constants:

```
A = [-0.0285  -0.0014  
     -0.0371  -0.1476];  
B = [-0.0850   0.0238  
     0.0802   0.4462];  
C = [0 1  
     1 0];  
D = zeros(2,2);  
CSTR = ss(A,B,C,D);
```

The following code sets some of the CSTR model's optional properties:

```
CSTR.InputName = {'T_c', 'C_A_i'};  
CSTR.OutputName = {'T', 'C_A'};  
CSTR.StateName = {'C_A', 'T'};  
CSTR.InputGroup.MV = 1;  
CSTR.InputGroup.UD = 2;  
CSTR.OutputGroup.MO = 1;  
CSTR.OutputGroup.UO = 2;
```

To view the properties of CSTR, enter:

```
CSTR
```

Linearize Simulink Models

Generally, real systems are nonlinear. To design an MPC controller of a nonlinear system, you must model the plant in Simulink.

Although an MPC controller can regulate a nonlinear plant, the model used within the controller must be linear. In other words, the controller employs a linear approximation of the nonlinear plant. The accuracy of this approximation significantly affects controller performance.

This raises the question: how to obtain the approximation? The usual approach is to *linearize* the nonlinear plant at a specified *operating point*. The Simulink environment provides two ways to accomplish this:

- “Linearization Using MATLAB Code” on page 2-17
- “Linearization Using Linear Analysis Tool in Simulink Control Design” on page 2-20

Note Simulink Control Design™ software must be installed to linearize nonlinear Simulink models.

Linearization Using MATLAB Code

This example shows how to programmatically obtain a linear model of a plant.

For this example, linearize the CSTR model, `CSTR_OpenLoop`. The model inputs are the coolant temperature (manipulated variable of the MPC controller), limiting reactant concentration in the feed stream, and feed temperature. The model states are the temperature and concentration of the limiting reactant in the product stream. You measure both states and use them for feedback control.

- 1 Obtain a steady-state operating point.

The operating point defines the nominal conditions at which you will linearize. It is usually a steady state condition.

Suppose you plan to operate the CSTR with the output concentration at 2 kmol/m³. The nominal feed concentration is 10 kmol/m³ and the nominal feed temperature is 300 K. Create an operating point specification object to define the steady state.

```
opspec = operspec('CSTR_OpenLoop');  
opspec = addoutputspec(opspec, 'CSTR_OpenLoop/CSTR', 2);  
opspec.Outputs(1).Known = true;  
opspec.Outputs(1).y = 2;
```

```
op1 = findop('CSTR_OpenLoop', opspec);
```

The calculated operating point is $C_A = 2$ kmol/m³ and $T_K = 373$ K. Notice that the coolant temperature at steady state is also given as 299 K, which is the nominal value of the manipulated variable of the model predictive controller.

To specify:

- Values of known inputs, use the `Input.Known` and `Input.u` fields of `opspec`
- Initial guesses for state values, use the `State.x` field of `opspec`

For example, the following code specifies the coolant temperature as 305K and initial guess values of the C_A and T_K states before calculating the steady-state operating point:

```
opspec = operspec('CSTR_OpenLoop');  
opspec.States(1).x = 1;  
opspec.States(2).x = 400;  
opspec.Inputs(1).Known = true;  
opspec.Inputs(1).u = 305;  
  
op2 = findop('CSTR_OpenLoop', opspec);
```

2 Specify linearization inputs/outputs.

If the linearization input /output signals are already defined in the model, as in `CSTR_OpenLoop`, then use the following to programmatically obtain the signal set:


```
io = getlinio('CSTR_OpenLoop');
```

Otherwise, specify the input/output signals as shown:

```
io(1) = linio('CSTR_OpenLoop/Feed Concentration', 1, 'input');
io(2) = linio('CSTR_OpenLoop/Feed Temperature', 1, 'input');
io(3) = linio('CSTR_OpenLoop/CoolantTemperature', 1, 'input');
io(4) = linio('CSTR_OpenLoop/CSTR', 1, 'output');
io(5) = linio('CSTR_OpenLoop/CSTR', 2, 'output');
```

3 Linearize the model at the specified operating point.

```
sys = linearize('CSTR_OpenLoop', op1, io)
```

```
sys
```

```
a =
```

	c_a	t_k
c_a	-5	-0.3427
t_k	47.68	2.785

```
b =
```

	coolant temp	feed concent	feed tempera
c_a	0	1	0
t_k	0.3	0	1

```
c =
```

	c_a	t_k
cstr/1	0	1
cstr/2	1	0

```
d =
```

	coolant temp	feed concent	feed tempera
cstr/1	0	0	0
cstr/2	0	0	0

sys is a continuous-time state-space model.

Linearization Using Linear Analysis Tool in Simulink Control Design

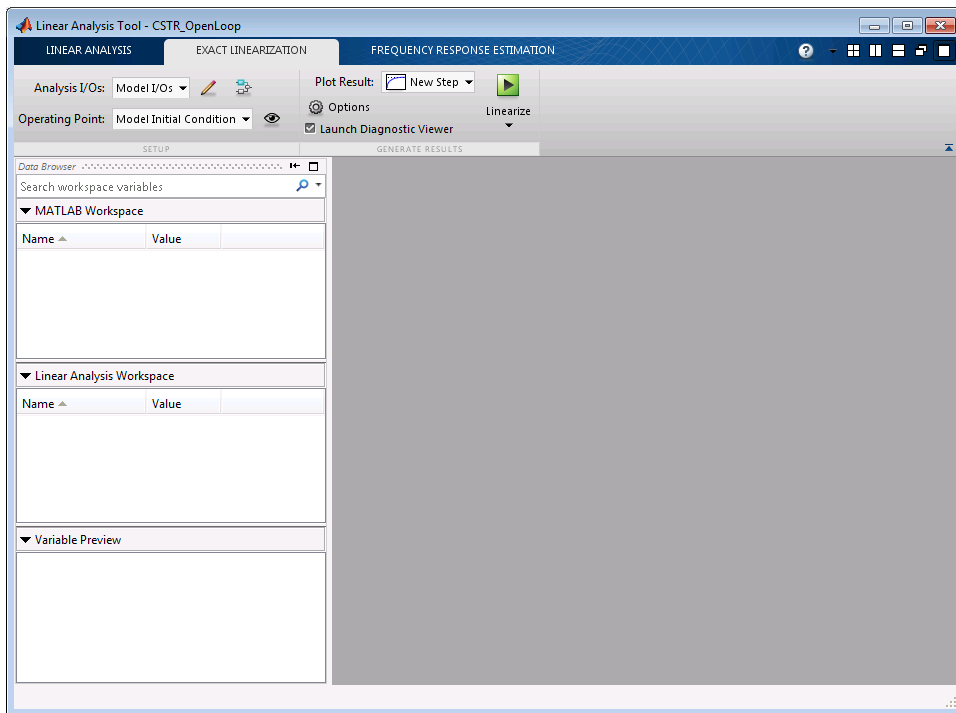
This example shows how to linearize a Simulink model using the Linear Analysis Tool. The Linear Analysis Tool, provided by Simulink Control Design, is a graphical interface for model linearization.

- 1 Open the Simulink model.

```
sys = 'CSTR_OpenLoop';  
open_system(sys)
```


- 2 Open the Linear Analysis Tool for the model.

In the Simulink model window, select **Analysis > Control Design > Linear Analysis**.



- 3 Specify linearization input/output signals.

- In the Simulink model window, right-click the output signal from the Feed Concentration block. Select **Linear Analysis Points > Input Perturbation**.
- In the Simulink model window, right-click the output signal from the Feed Temperature block. Select **Linear Analysis Points > Input Perturbation**.
- In the Simulink model window, right-click the output signal from the Coolant Temperature block. Select **Linear Analysis Points > Input Perturbation**.
- In the Simulink model window, right-click the input signal to the CSTR Temperature block. Select **Linear Analysis Points > Output Measurement**.
- In the Simulink model window, right-click the input signal to the Residual Concentration block. Select **Linear Analysis Points > Output Measurement**.



4 Specify the residual concentration as a trim constraint.

In the Simulink model window, right-click the CA output signal from the CSTR block. Select **Linear Analysis Points > Trim Output Constraint**.


5 Opens the **Specifications for trim** dialog box.

In the Linear Analysis Tool, click the **Exact Linearization** tab. In the **Operating Point** list, select Trim model. The **Trim Model** tab opens. Click **Specifications**.

6 Specify the residual concentration as a known output.

Select the **Outputs** tab and set the value of the **CSTR_OpenLoop/CSTR** (CA) to 2 kmol/m³.

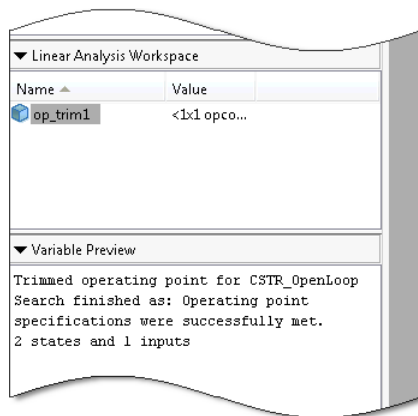
Select the **Known** check box.




7 Create and verify the operating point.

In the **Trim Model** tab, click **Trim**.

The operating point **op_trim1** appears in the **Linear Analysis Workspace**.



Double click **op_trim1** to view the resulting operating point. This action opens the Edit trim point dialog box. Select the **Input** tab.

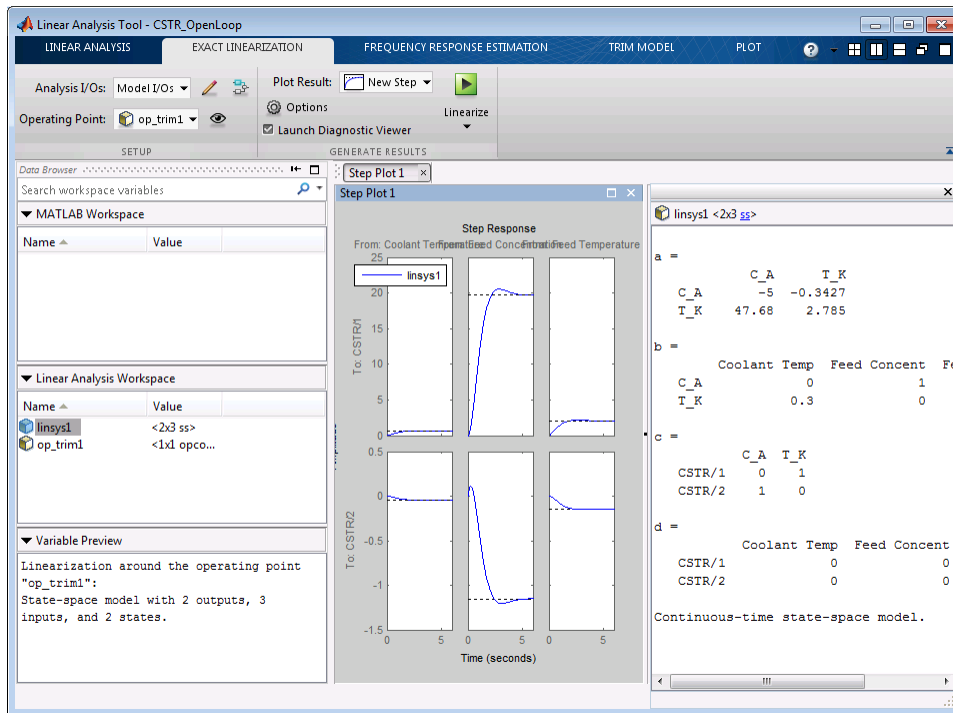


Input	Desired Value	Actual Value
CSTR_OpenLoop/Coolant Temperature		
Input - 1	[-Inf, Inf]	299.0349

The coolant temperature at steady state is 299 K, as desired.

- 8** In the **Exact Linearization** tab, select `op_trim1` from the **Operating Point** list.
- 9** Click **Linearize**.

This action creates the linear model `linsys1` in the **Linear Anaysis Workspace**. `linsys1` uses `op_trim1` as its operating point.



The step response from feed concentration to output CSTR/2 displays an interesting inverse response. An examination of the linear model shows that CSTR/2 is the residual CSTR concentration, C_A . When the feed concentration increases, C_A increases initially (more reactant is entering), which increases the reaction rate. This increases the reactor temperature (output CSTR/1), which further increases the reaction rate and C_A decreases dramatically.

10 Export the linearization result to the MATLAB workspace.

If necessary, you can repeat any of the above steps. Once you are satisfied with your linear model, drag and drop it from the **Linear Analysis Workspace** to the **MATLAB Workspace** in the Linear Analysis Tool. You may now use the linearization result with other tools, such as the controller design.

Identify Plant from Data

This example shows how to identify a linear plant model using measured data.

When you have measured plant input/output data, you can use System Identification Toolbox software to estimate a linear plant model. Then, you can use the estimated plant model to create a model predictive controller.

You can estimate the plant model either programmatically or by using the System Identification Tool. For information regarding using the System Identification Tool to identify a model, see “Identify Linear Models Using the GUI”.

Note This example requires a System Identification Toolbox license.

Load the measured data.

```
load dryer2
```

The variables `u2` and `y2`, which contain the data measured for a temperature-control application, are loaded into the MATLAB workspace. `u2` is the plant input, and `y2` is the plant output. The sampling period for `u2` and `y2` is 0.08 seconds.

Create an `iddata` object for the measured data.

Use an `iddata` object to store the measured values of the plant inputs and outputs, sampling time, channel names, etc.

```
Ts = 0.08;  
dry_data = iddata(y2,u2,Ts);
```

`dry_data` is an `iddata` object that contains the measured input/output data.

You can optionally assign channel names and units for the input/output signals. To do so, use the `InputName`, `OutputName`, `InputUnit` and `OutputUnit` properties of an `iddata` object. For example:

```
dry_data.InputName = 'Power';
```

```
dry_data.OutputName = 'Temperature';
```

Detrend the measured data.

Before estimating the plant model, preprocess the measured data to increase the accuracy of the estimated model. For this example, `u2` and `y2` contain constant offsets that you eliminate.

```
dry_data_detrended = detrend(dry_data);
```

Estimate a linear plant model.

You can use System Identification Toolbox software to estimate a linear plant model in one of the following forms:

- State-space model
- Transfer function model
- Polynomial model
- Process model
- Grey-box model

For this example, estimate a third-order, linear state-space plant model using the detrended data.

```
plant = ssest(dry_data_detrended,3);
```

To view the estimated parameters and estimation details, at the MATLAB command prompt, type `plant`.

See Also

`iddata` | `detrend` | `ssest`

Related Examples

- “Identify Linear Models Using the GUI”
- “Design Controller for Identified Plant” on page 2-29
- “Design Controller Using Noise Model of Identified Model” on page 2-31

Concepts

- “Handling Offsets and Trends in Data”
- “Ways to Design MPC Controllers for Identified Plant Models” on page 2-27
- “Working with Impulse-Response Models” on page 2-33

Ways to Design MPC Controllers for Identified Plant Models

When you have measured plant data, you can use the System Identification Toolbox software to identify a linear plant model. You use the `mpc` function to create a model predictive controller for the identified plant model.

An identified model has two types of inputs—measured inputs and unmeasured noise inputs. For more information regarding the structure of an identified model, see “Input Output Structure of Identified Model” on page 2-28.

There are two ways you create an MPC controller using the identified model:

- Use only the measured components of the identified model for the controller plant model. This approach is the preferred method for creating an MPC controller for an identified plant model.
 - Measured inputs of the identified model are treated as the manipulated variables of the model predictive controller.
 - Measured outputs of the identified plant model are treated as the measured outputs of the plant for the controller.

For an example, see “Design Controller for Identified Plant” on page 2-29.

- Use the noise model of the identified model for the controller plant model. When you have a reliable noise model of the identified plant, you can use this method to create an MPC controller.

The measured components of the identified model are treated just as in the first approach. However, this approach additionally treats the unmeasured noise inputs of the identified model as the unmeasured disturbances to the plant for the controller.

Unmeasured input disturbances influence the default controller design. A plant may yield different noise models for different experiments. If the controller works well during setpoint changes, but eliminates disturbances slowly or exhibits steady-state error, try modifying the unmeasured disturbance model.

For an example, see “Design Controller Using Noise Model of Identified Model” on page 2-31.

Input Output Structure of Identified Model

The input/output relationship for an identified model is:

$$y = Gu + He$$

- y is the measured output.
- u is the measured input.
- e is the *unmeasured noise* input and is Gaussian white noise of a specified variance.
- G is the *measured* transfer function that relates the measured input, u , to the output, y .
- H is the *noise* transfer function that describes the effect of output disturbances.

The noise channel of an identified model, e , is not used in calculating the input size of an identified model. Also, by default, e is not accounted for in traditional model analysis such as simulation, frequency response computation, and pole-zero maps.

See Also mpc

Related Examples

- “Design Controller for Identified Plant” on page 2-29
- “Design Controller Using Noise Model of Identified Model” on page 2-31
- “Identify Plant from Data” on page 2-25

Concepts

- “About Identified Linear Models”

Design Controller for Identified Plant

This example shows how to design a model predictive controller for a linear identified plant model.

This example uses only the measured input and output of the identified model for the model predictive controller plant model.

Note This example requires a System Identification Toolbox license.

Obtain an identified linear plant model.

```
load dryer2;
Ts = 0.08;
dry_data = iddata(y2,u2,Ts);
dry_data_detrended = detrend(dry_data);
plant_idss = ssest(dry_data_detrended,3);
```

`plant_idss` is a third-order, identified state-space model that contains one measured input and one unmeasured (noise) input.

Convert the identified plant model to a numeric LTI model.

You can convert the identified model to an `ss`, `tf`, or `zpk` model. For this example, convert the identified state-space model to a numeric state-space model.

```
plant_ss = ss(plant_idss);
```

`plant_ss` contains the measured input and output of `plant_idss`. The software discards the noise input of `plant_idss` when it creates `plant_ss`.

Design a model predictive controller for the numeric plant model.

```
controller = mpc(plant_ss,Ts);
```

`controller` is an `mpc` object. The software treats:

- The measured input of `plant_ss` as the manipulated variable of controller
- The output of `plant_ss` as the measured output of the plant for controller

To view the structure of the model predictive controller, type `controller` at the MATLAB command prompt.

See Also `mpc`

Related Examples

- “Identify Plant from Data” on page 2-25
- “Design Controller Using Noise Model of Identified Model” on page 2-31

Concepts

- “Ways to Design MPC Controllers for Identified Plant Models” on page 2-27

Design Controller Using Noise Model of Identified Model

This example shows how to design a model predictive controller that includes the noise model of an identified plant model.

Note This example requires a System Identification Toolbox license.

Obtain an identified linear plant model.

```
load dryer2;
Ts = 0.08;
dry_data = iddata(y2,u2,Ts);
dry_data_detrended = detrend(dry_data);
plant_idss = ssest(dry_data_detrended,3);
```

`plant_idss` is a third-order, identified state-space model that contains one measured input and one unmeasured (noise) input.

Design a model predictive controller for the identified plant model.

```
controller = mpc(plant_idss,Ts);
```

`controller` is an `mpc` object. The software treats:

- The measured input of `plant_ss` as the manipulated variable of `controller`
- The unmeasured noise input of `plant_ss` as the unmeasured disturbance of the plant for `controller`
- The output of `plant_ss` as the measured output of the plant for `controller`

To view the structure of the model predictive controller, at the MATLAB command prompt, type `controller`.

Tip You can change the treatment of a plant input in one of two ways:

- Programmatic — Use the `setmpcsignals` command to modify the signal .
 - Model Predictive Control Toolbox design tool — Use the **Input signal properties** table to modify the plant model signal types.
-

You can also design a model predictive controller using:

```
plant_ss = ss(plant_idss, 'augmented');  
controller = mpc(plant_ss, Ts);
```

When you use the 'augmented' input argument, `ss` creates two input groups, `Measured` and `Noise`, for the measured and noise inputs of `plant_idss`. `mpc` handles the measured and noise components of `plant_ss` and `plant_idss` identically.

See Also `mpc`

Related Examples

- “Identify Plant from Data” on page 2-25
- “Design Controller for Identified Plant” on page 2-29

Concepts

- “Ways to Design MPC Controllers for Identified Plant Models” on page 2-27

Working with Impulse-Response Models

You can use System Identification Toolbox software to estimate finite step-response or finite impulse-response (FIR) plant models using measured data. Such models, also known as *nonparametric models* (see [6] for example), are easy to determine from plant data ([3] and [7]) and have intuitive appeal.

You use the `impulseest` function to estimate an FIR model from measured data. The function returns an identified transfer function model, `idtf`. To design a model predictive controller for the plant, you can convert the identified FIR plant model to a numeric LTI model. However, this conversion usually yields a high-order plant, which can degrade the controller design. This result is particularly an issue for MIMO systems. For example, the estimator design can be affected by numerical precision issues with high-order plants.

Model predictive controllers work best with low-order parametric models (see [5] for example). Therefore, to design a model predictive controller using measured plant data, you can use a parametric estimator, such as `ssest`. Then estimate a low-order parametric plant model. Alternatively, you can initially identify a nonparametric model using the data and then estimate a low-order parametric model for the nonparametric model's response. (See [10] for an example.)

See Also

`impulseest` | `ssest`

Related Examples

- “Identify Plant from Data” on page 2-25
- “Design Controller for Identified Plant” on page 2-29
- “Design Controller Using Noise Model of Identified Model” on page 2-31

Concepts

- “Ways to Design MPC Controllers for Identified Plant Models” on page 2-27

Bibliography

- [1] Allgower, F., and A. Zheng, *Nonlinear Model Predictive Control*, Springer-Verlag, 2000.
- [2] Camacho, E. F., and C. Bordons, *Model Predictive Control*, Springer-Verlag, 1999.
- [3] Cutler, C., and F. Yocum, “Experience with the DMC inverse for identification,” *Chemical Process Control — CPC IV* (Y. Arkun and W. H. Ray, eds.), CACHE, 1991.
- [4] Kouvaritakis, B., and M. Cannon, *Non-Linear Predictive Control: Theory & Practice*, IEE Publishing, 2001.
- [5] Maciejowski, J. M., *Predictive Control with Constraints*, Pearson Education POD, 2002.
- [6] Prett, D., and C. Garcia, *Fundamental Process Control*, Butterworths, 1988.
- [7] Ricker, N. L., “The use of bias least-squares estimators for parameters in discrete-time pulse response models,” *Ind. Eng. Chem. Res.*, Vol. 27, pp. 343, 1988.
- [8] Rossiter, J. A., *Model-Based Predictive Control: A Practical Approach*, CRC Press, 2003.
- [9] Seborg, D. E., T. F. Edgar, and D. A. Mellichamp, *Process Dynamics and Control*, 2nd Edition, Wiley, 2004, pp. 34–36 and 94–95.
- [10] Wang, L., P. Gawthrop, C. Chessari, T. Podsiadly, and A. Giles, “Indirect approach to continuous time system identification of food extruder,” *J. Process Control*, Vol. 14, Number 6, pp. 603–615, 2004.
- [11] Wood, R. K., and M. W. Berry, *Chem. Eng. Sci.*, Vol. 28, pp. 1707, 1973.

Designing Controllers Using the Design Tool GUI

- “Design Controller Using the Design Tool” on page 3-2
- “Test Controller Robustness” on page 3-50
- “Design Controller for Plant with Delays” on page 3-53
- “Design Controller for Nonsquare Plant” on page 3-58
- “Nonlinear Plants” on page 3-60

Design Controller Using the Design Tool

In this section...

“Start the Design Tool” on page 3-2

“Load a Plant Model” on page 3-3

“Navigate Using the Tree View” on page 3-6

“Perform Linear Simulations” on page 3-10

“Customize Response Plots” on page 3-18

“Change Controller Settings” on page 3-23

“Define Soft Output Constraints” on page 3-42

“Save Your Work” on page 3-46

“Load Your Saved Work” on page 3-48

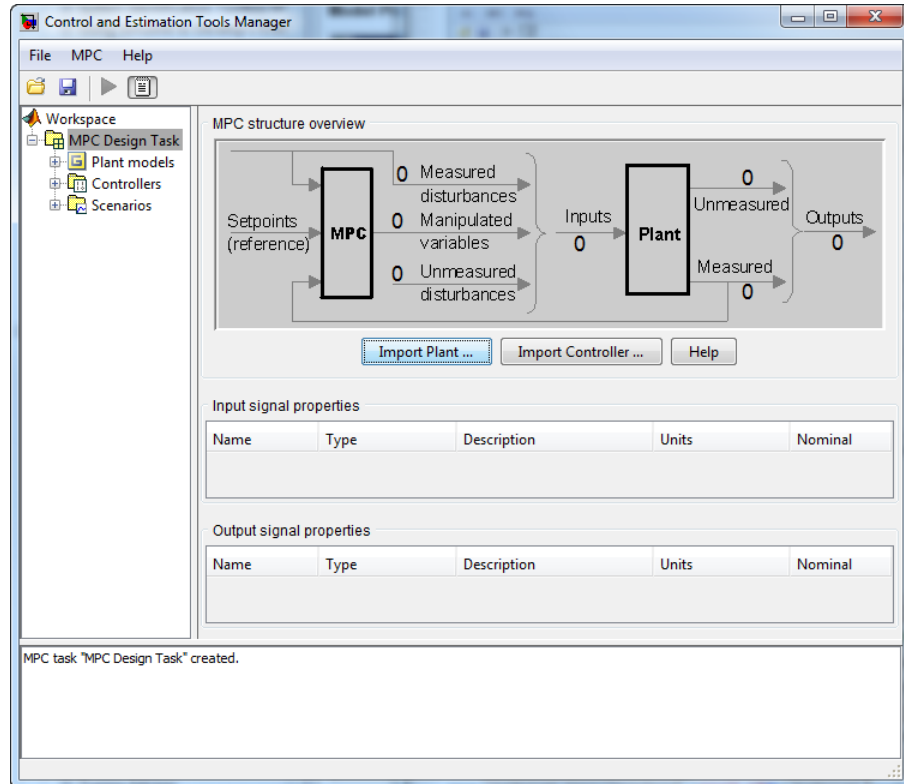
The Model Predictive Control Toolbox design tool is a graphical user interface for controller design. This GUI is part of the Control and Estimation Tools Manager GUI.

Start the Design Tool

Start the design tool by typing the MATLAB command

```
mpctool
```

The Control and Estimation Tools Manager window appears, as shown below. By default, it contains a Model Predictive Control Toolbox task called **MPC Design Task** (listed in the tree view on the left side of the window), which is selected, causing the view shown on the right to appear.



Model Predictive Control Toolbox Design Tool Initial View


Load a Plant Model

The first step in the design is to load a plant model. Its dimensions and signal characteristics set the context for the remaining steps. You can either load the model directly, as described in this section, or indirectly by importing a controller or a saved design (see “Load Your Saved Work” on page 3-48).

The following example uses the CSTR model described in “CSTR Model” on page 2-15. Verify that the LTI object CSTR is in your MATLAB workspace.

Plant Model Importer Dialog Box

Click the **Import Plant** button in the design tool's initial view (see Model Predictive Control Toolbox Design Tool Initial View on page 3-3). The Plant Model Importer dialog box appears.

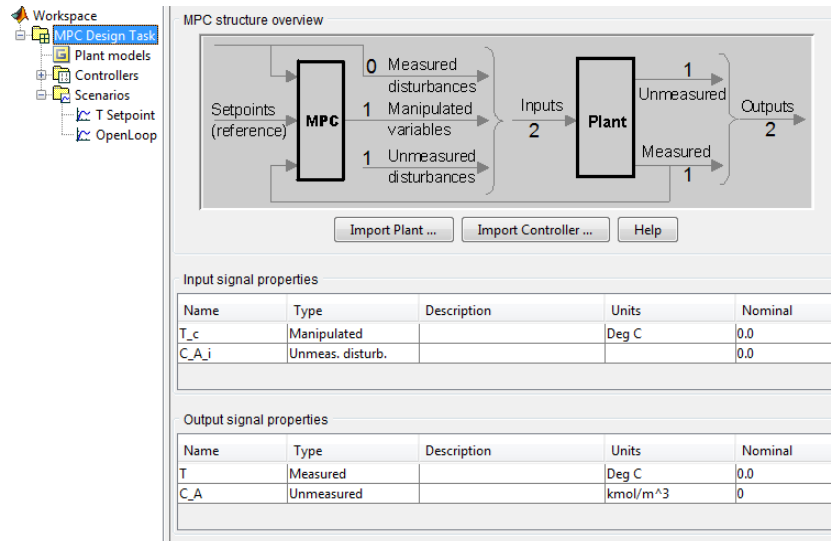


Plant Model Importer Dialog Box

The **Import from MATLAB workspace** option button should be selected by default, as shown. The **Items in your workspace** table lists your LTI models. If CSTR doesn't appear, define it as discussed in "State-Space Models" on page 2-5, then reopen this dialog box.

After CSTR appears, select it. The **Properties** list displays the number of inputs and outputs, their names and signal types, etc.

Click the **Import** button. This loads CSTR into the design tool. Then click the **Close** button (otherwise the dialog box remains visible in case you want to import another model). The design tool should appear as in Model Predictive Control Toolbox Design Tool's Signal Definition View on page 3-5.



Model Predictive Control Toolbox Design Tool's Signal Definition View

Signal Property Specifications

The figure's graphical display indicates that you've imported a plant model by showing the number of inputs and outputs, and the number in each subclass: measured disturbance, manipulated variables, etc.). Also, the tables labeled **Input signal properties** and **Output signal properties** fill with data:

- The **Name** entries are from the CSTR model's InputName and OutputName properties (the design tool assigns defaults if necessary). You can edit these at any time.
- The **Type** entries are from the CSTR model's InputGroup and OutputGroup properties. (The design tool defaults all unspecified inputs to manipulated variables and all unspecified outputs to measured.)

Note Once you leave this view, if you subsequently change a signal type, you will have to restart the design. Be sure the signal types are correct at the beginning.

- The **Description** and **Unit** entries are optional. You can enter the values shown in Model Predictive Control Toolbox Design Tool's Signal Definition View on page 3-5 manually. As you will see, the design tool uses them to label plots and other tables.
- The **Nominal** entries are initial conditions for simulations. The design tool default is 0.0.

Navigate Using the Tree View

The *tree* in the left-hand frame of Model Predictive Control Toolbox Design Tool's Signal Definition View on page 3-5 shows that the default **MPC Design Task** node (see Model Predictive Control Toolbox Design Tool Initial View on page 3-3) has been renamed **CSTRcontrol** by clicking the name, waiting for the usual edit box to appear, typing the new name, and pressing **Enter** to finalize the choice.

Model Predictive Control Toolbox Design Tool's Signal Definition View on page 3-5 also shows three new nodes below **CSTRcontrol**. These activate once you've imported a plant model (or controller).

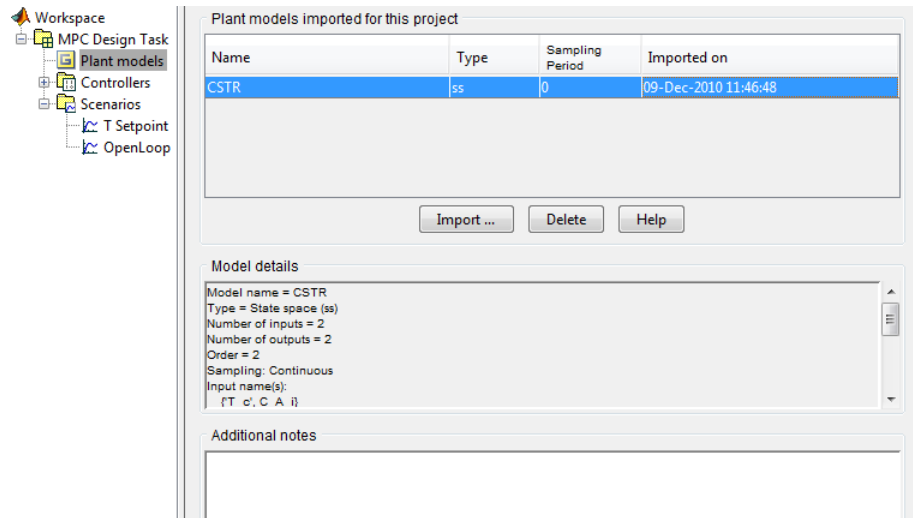
In general, clicking a node displays a *view* supporting a particular design activity.

Project View (Signal Properties Tables)

For example, clicking the *project* node (**CSTRcontrol** in Model Predictive Control Toolbox Design Tool's Signal Definition View on page 3-5) allows you to review and edit the signal properties tables.

Listing Your Plant Models


Select the **Plant models** node to list the plant models you've imported, as shown in Plant Models View with CSTR Model Selected on page 3-7. (Each model name is editable.) The **Model details** section displays properties of the selected model. There is also a space to enter notes describing the model's special features. Buttons allow you to import a new model or delete one you no longer need.



Plant Models View with CSTR Model Selected

Viewing Your Controllers

Next, select **Controllers**. The view shown in Controllers View on page 3-8 appears. A + sign to the left of **Controllers** indicates that it contains subnodes. You can click a + sign to expand the tree (as shown in Controllers View on page 3-8, where the + sign has changed to a – sign).



Controllers View

The table at the top of Controllers View on page 3-8 lists all the controllers you've defined. The design tool automatically creates a controller containing Model Predictive Control Toolbox defaults, naming it **MPC1**. It is a subnode of **Controllers**.

Note If you define additional controllers, they will appear here. For example, you might want to test several options, saving each as a separate controller, making it easy to switch from one to another during testing.

The table **Controllers defined in this project** allows you to edit the controller name and gives you quick access to three important design parameters: **Plant Model**, **Control Interval**, and **Prediction Horizon**. All are editable, but leave them at their default values for this example.

The buttons shown in Controllers View on page 3-8 let you do the following:

- Import a controller designed previously and stored either in your workspace or in a MAT-file.
- Export the selected controller to your workspace.

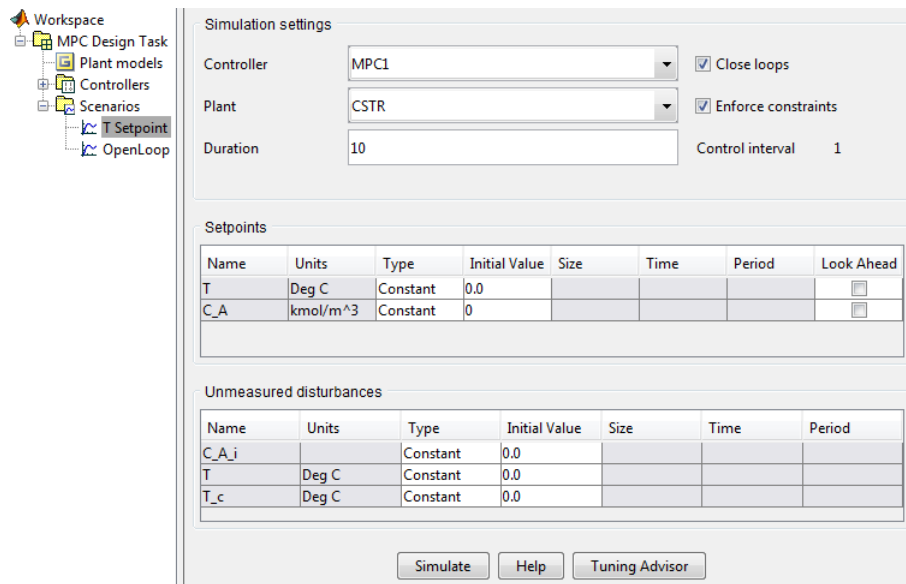
- Create a new controller initialized to Model Predictive Control Toolbox defaults.
- Copy the selected controller, creating a duplicate you can modify.
- Delete the selected controller.

You can also right-click the **Controllers** node to access menu options **New**, **Import**, and **Export**, or one of its subnodes to access menu options **Copy**, **Rename**, **Export**, and **Delete**.

Select the **MPC1** node to display Model Predictive Control Toolbox default controller settings. (“Change Controller Settings” on page 3-23 covers this view in detail).

Viewing Simulation Scenarios

A *scenario* is a set of conditions defining a simulation. The design tool creates a default scenario and names it **Scenario1**. To view it, click the **+** symbol next to the **Scenarios** node, and select the **Scenario1** subnode. You should see a view like that shown in CSTR Temperature Setpoint Change Scenario on page 3-12.



Scenarios View

Whenever you select the **Scenarios** node, you see a table summarizing your current scenarios (not shown). Its function is similar to the **Controllers** view described previously.

Perform Linear Simulations

You usually want to test your controller in simulations. The Model Predictive Control Toolbox design tool makes it easy to run closed-loop simulations involving a Model Predictive Control Toolbox controller and an LTI plant model. This plant can differ from that used in the controller design, allowing you to test your controller's sensitivity to prediction errors (see "Test Controller Robustness" on page 3-50).

This section covers the following topics:

- "Defining Simulation Conditions" on page 3-11
- "Running a Simulation" on page 3-12
- "Open-Loop Simulations" on page 3-15


Defining Simulation Conditions

To define simulation conditions, select an existing scenario node or create a new one, and then edit its tabular fields to define your conditions.

CSTR Temperature Setpoint Change Scenario on page 3-12 shows the result of renaming the default **Scenario1** node to **T Setpoint** and editing its default conditions. The required editing steps are as follows:

- Increase **Duration** from 10 to 30.
- Locate the tabular data defining the reactor temperature setpoint (first row of the upper table in CSTR Temperature Setpoint Change Scenario on page 3-12).
 - Click the **Type** table cell and select **Step** from the list of choices.
 - Change **Size** from 1.0 to 2.
 - Change **Time** from 1.0 to 5.

Note The **Control Interval** is a property of the controller being used (MPC1 in this case). To change it, select the controller node in the tree, and then edit the value on the **Model and Horizons** tab (see “Model and Horizons” on page 3-24). Such a change would apply to all simulations involving that controller.



CSTR Temperature Setpoint Change Scenario


Running a Simulation

To run a simulation, do *one* of the following:


- Select the scenario you want to run, and click its **Simulate** button (see the bottom of CSTR Temperature Setpoint Change Scenario on page 3-12).
- Click the toolbar's **Simulation** button, which is the triangular icon shown on the top left of CSTR Temperature Setpoint Change Scenario on page 3-12.

The toolbar method runs the *current scenario*, i.e., the one most recently selected or modified.

Try running the **T Setpoint** scenario. This should generate the two response plot windows shown in Plant Outputs for T Setpoint Scenario with Added Data Markers on page 3-13 and Plant Inputs for the T Setpoint Scenario on page 3-14.



Plant Outputs for T Setpoint Scenario with Added Data Markers



Plant Inputs for the T Setpoint Scenario

Plant Outputs for T Setpoint Scenario with Added Data Markers on page 3-13 shows that the reactor temperature setpoint increases suddenly by 2 degrees at $t = 5$, as you specified when defining the scenario in “Defining Simulation Conditions” on page 3-11. Unfortunately, the temperature does not track the setpoint very well, and there is a persistent error of about 1.6 degrees at the end of the simulation.

Also, the controller requests a sudden jump in the coolant temperature (see the upper graph in Plant Inputs for the T Setpoint Scenario on page 3-14), which might be difficult to deliver in practice. (The lower graph in Plant Inputs for the T Setpoint Scenario on page 3-14 shows that the feed concentration, C_{A_i} , remains constant, as specified in the scenario.)

See “Change Controller Settings” on page 3-23 for ways to overcome these deficiencies.

Note Plant Outputs for T Setpoint Scenario with Added Data Markers on page 3-13 has *data markers*. To add these, left-click the curve to create the data marker. Drag a marker to relocate it. Left-click in a graph's white space to erase its markers. For more information on data markers, see the Control System Toolbox documentation.

Open-Loop Simulations

By default, scenarios are *closed loop*, i.e., an active controller adjusts the manipulated variables entering your plant based on feedback from the plant outputs. You can also run *open-loop* simulations that test the plant model without feedback control.

For example, you might want to check your plant model's response to a particular input without opening another tool. You might also want to display unmeasured disturbance signals before using them in a closed-loop simulation.

To see how this works, create a new scenario by right-clicking the **T Setpoint** node in the tree, and selecting **Copy Scenario** in the resulting menu. Rename the copy **OpenLoop**.

Select **OpenLoop** in the tree. On its scenario view, change Duration to 100, and turn off (clear) **Close loops**.

Open-loop simulations ignore the **Setpoints** table settings, so there's no need to modify them.

If CSTR had a measured disturbance input, the pane would contain another table allowing you to specify it.


For this example, focus on the **Unmeasured disturbances** table. Configure it as shown below.

Unmeasured disturbances						
Name	Units	Type	Initial Value	Size	Time	Period
C_A_i	kmol/m ³	Constant	1			
T	Deg C	Constant	0.0			
T_c	Deg C	Constant	0.0			

The C_{A_i} input's nominal value is 0.0 (see Model Predictive Control Toolbox Design Tool's Signal Definition View on page 3-5), so the above models a sudden increase to 1 at the beginning of the simulation. The following is an equivalent setup using the Step type.

Unmeasured disturbances						
Name	Units	Type	Initial Value	Size	Time	Period
C_A_i	kmol/m ³	Step	0	1	0	
T	Deg C	Constant	0.0			
T_c	Deg C	Constant	0.0			

Using one of these, simulate the scenario (click its **Simulate** button). The output response plot should be as shown below.



This is the CSTR model's open-loop response to a unit step in the C_{A_i} disturbance input.

You could also set up the table as shown below.

Unmeasured disturbances						
Name	Units	Type	Initial Value	Size	Time	Period
C_A_i	kmol/m ³	Constant	0.0			
T	Deg C	Constant	0.0			
T_c	Deg C	Constant	1			


This simulation would display the open-loop response to a unit step in the T_c manipulated variable input (try it).

Finally, set it up as follows.

Unmeasured disturbances						
Name	Units	Type	Initial Value	Size	Time	Period
C_A_i	kmol/m ³	Constant	0.0			
T	Deg C	Pulse	0.0	0.95	10	20
T_c	Deg C	Constant	0.0			

This adds a pulse to the T output. The pulse begins at time $t = 10$, and lasts 20 time units. Its height is 0.95 degrees.

Run the simulation. The output response plot displays the pulse (see Response Plot Showing Open-Loop Pulse Disturbance on page 3-17). In this case, the T output's nominal value is zero, so you only see the pulse. (If the T output had a nonzero nominal value, the pulse would add to that.)




Response Plot Showing Open-Loop Pulse Disturbance

If you were to run a *closed-loop* simulation with this same T disturbance, the controller would attempt to hold T at its setpoint, and the result would differ

from that shown in Response Plot Showing Open-Loop Pulse Disturbance on page 3-17.

Customize Response Plots

Each time you simulate a scenario, the design tool plots the corresponding plant input and output responses. The graphic below shows such a *response plot* for a plant having two outputs (the corresponding input response plot is not shown).



By default, each plant signal plots in its own graph area (as shown above). If the simulation is closed loop, each output signal plot include the corresponding setpoint.

The following sections describe response plot customization options:


- “Data Markers” on page 3-19
- “Displaying Multiple Scenarios” on page 3-20
- “Viewing Selected Variables” on page 3-21
- “Grouping Variables in a Single Plot” on page 3-22

- “Normalizing Response Amplitudes” on page 3-22

Data Markers

You can use data markers to label a curve or to display numerical details.

Adding a Data Marker. To add a data marker, click the desired curve at the location you want to mark. The following graph shows a marker added to each output response and its corresponding setpoint.



Data Marker Contents. Each data marker provides information about the selected point, as follows:

- **Response** – The *scenario* that generated the curve.
- **Time** – The time value at the data marker location.
- **Amplitude** – The signal value at the data marker location.
- **Output** – The plant variable name (plant outputs only).
- **Input** – Variable name for plant inputs and setpoints.

Changing a Data Marker's Alignment. To relocate the data marker's label (without moving the marker), right-click the marker, and select one of the four **Alignment** menu options. The above example shows three of the possible four alignment options.

Relocating a Data Marker. To move a marker, left-click it (holding down the mouse key) and drag it along its curve to the desired location.

Deleting Data Markers. To delete all data markers in a plot, click in the plot's white space.

To delete a single data marker, right-click it and select the **Delete** option.

Right-Click Options. Right-click a data marker to use one of the following options:

- **Alignment** – Relocate the marker's label.
- **Font Size** – Change the label's font size.
- **Movable** – On/off option that makes the marker movable or fixed.
- **Delete** – Deletes the selected marker.
- **Interpolation** – Interpolate linearly between the curve's data points, or locate at the nearest data point.
- **Track Mode** – Changes the way the marker responds when you drag it.

Displaying Multiple Scenarios

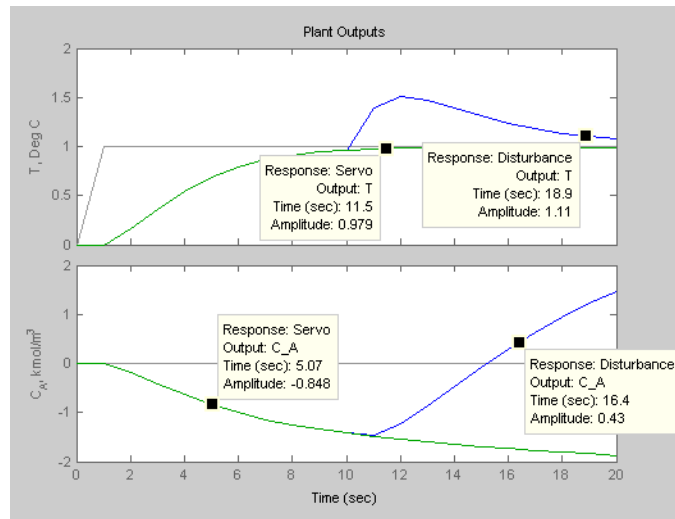
By default the response plots include all the scenarios you've simulated. The example below shows a response plot for a plant with two outputs. The data markers indicate the two scenarios being plotted: "Accurate Model" and "Perturbed Model". Both scenarios use the same setpoints (not marked—the lighter solid lines).

Viewing Selected Scenarios. If your plots are too cluttered, you can hide selected scenarios. To do so:

- Right-click in the plot's white space.
- Select **Responses** from the resulting context menu.

- Toggle a response on or off using the submenu.

Note This selection affects all variables being plotted.



Revising a Scenario. If you modify and recalculate a scenario, its data are replotted, replacing the original curves.

Viewing Selected Variables

By default, the design tool plots all plant inputs in a single window, and plots all plant outputs in another. If your application involves many signals, the plots of each may be too small to view comfortably.

Therefore, you can control the variables being plotted. To do so, right-click in a plot's white space and select **Channel Selector** from the resulting menu. A dialog box appears, on which you can opt to show or hide each variable.

Grouping Variables in a Single Plot


By default, each variable appears in its own plot area. You can instead choose to display variables together in a single plot. To do so, right-click in a plot's white space, and select **Channel Grouping**, and then select **All**.

To return to the default mode, use the **Channel Grouping: None** option.

Normalizing Response Amplitudes

When you're using the **Channel Grouping: All** option, you might find that the variables have very different scales, making it difficult to view them together. You can choose to *normalize* the curves, so that each expands or contracts to fill the available plot area.

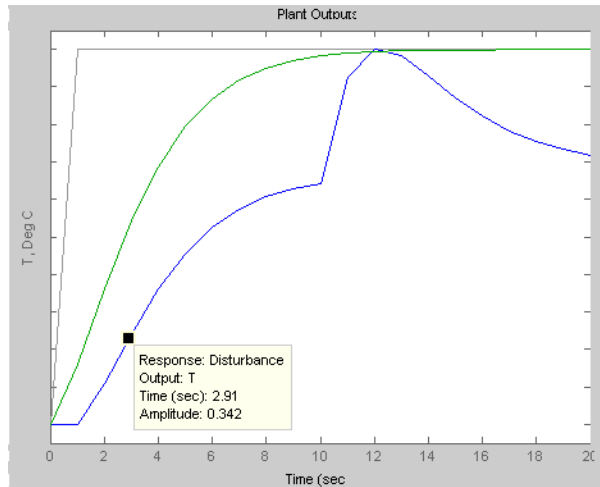
For example, the plot below shows two plant outputs together (**Channel Grouping: All** option). The outputs have very different magnitudes. When plotted together, it's hard to see much detail in the smaller response.



The plot below shows the normalized version, which displays each curve's variations clearly.

The y-axis scale is no longer meaningful, however. If you want to know a normalized signal's amplitude, use a data marker (see "Adding a Data

Marker” on page 3-19). Note that the two data markers on the plot below are at the same normalized y-axis location, but correspond to very different amplitudes in the original (unnormalized) coordinates.



Change Controller Settings

The simulations shown in Plant Outputs for T Setpoint Scenario with Added Data Markers on page 3-13 and Plant Inputs for the T Setpoint Scenario on page 3-14 uses the default controller settings. These often work well, but the CSTR application is an exception. The following discussion covers the main controller options in detail, and shows how to tune a controller for better performance.

This section covers the following topics:


- “Model and Horizons” on page 3-24
- “Weight Tuning” on page 3-25
- “Blocking” on page 3-30
- “Defining Manipulated Variable Constraints” on page 3-33
- “Disturbance Modeling and Estimation” on page 3-35

Model and Horizons

Select your **MPC1** controller in the tree. The view shown in Controller Options — Model and Horizons Tab on page 3-24 should appear. If necessary, click the **Model and Horizons** tab to bring it to the front. This tab contains the following controller options:

- **Plant model** specifies the LTI model to be used for controller predictions.
- **Control interval** sets the elapsed time between successive adjustments of the controller’s manipulated variables.
- **Prediction horizon** is the number of control intervals over which the outputs are to be optimized.
- **Control horizon** sets the number of control intervals over which the manipulated variables are to be optimized.
- Selecting the **Blocking** option gives you more control over the way in which the controller’s *moves* are allocated.

Leave all of the **Model and Horizons** tab settings at their default values for now.



Controller Options — Model and Horizons Tab

Weight Tuning

Click the **Weight Tuning** tab. The view shown in Controller Options — Weight Tuning Tab on page 3-25 appears.

The screenshot shows the 'Weight Tuning' tab in the Design Tool. The workspace tree on the left includes 'MPC Design Task', 'Plant models', 'Controllers' (with 'MPC1' selected), 'Scenarios' (with 'T Setpoint' selected), and 'OpenLoop'. The main panel has tabs for 'Model and Horizons', 'Constraints', 'Weight Tuning', and 'Estimation (Advanced)'. The 'Weight Tuning' tab is active, showing an 'Overall' section with a slider for 'More robust' to 'Faster response' set to 0.8. Below are 'Input weights' and 'Output weights' tables.

Name	Description	Units	Weight	Rate Weight
T_c		Deg C	0	0.3

Name	Description	Units	Weight
T		Deg C	1.0
C_A		kmol/m ³	0


Controller Options – Weight Tuning Tab

Make the following changes (already done in the above view):

- In the **Input weights** section, change the coolant temperature's **Rate Weight** from the default 0.1 to 0.3.
- In the **Output weights** section, change the reactant concentration's **Weight** (last entry in the second row) from the default 1.0 to 0.

Test these changes using the **T Setpoint** scenario (click the toolbar's **Simulation** button).

Improved Setpoint Tracking for CSTR Temperature on page 3-26 shows that the CSTR temperature now tracks the setpoint change smoothly, reaching the new value in about 10 time units with no overshoot.



Improved Setpoint Tracking for CSTR Temperature

On the other hand, the reactant concentration, C_A , exhibits a larger deviation from its setpoint, which is being held constant at zero (compare to Plant Outputs for T Setpoint Scenario with Added Data Markers on page 3-13, where the final deviation is about a factor of 4 smaller).

This behavior reflects an unavoidable trade-off. The controller has only one adjustment at its disposal: the coolant temperature. Therefore, it can't satisfy setpoints on both outputs.

Output Weights. The output weights let you dictate the accuracy with which each output must track its setpoint. Specifically, the controller predicts deviations for each output over the prediction horizon. It multiplies each deviation by the output's weight value, and then computes the weighted sum of squared deviations, $S_y(k)$, as follows

$$S_y(k) = \sum_{i=1}^P \sum_{j=1}^{n_y} \{w^y_j [r_j(k+i) - y_j(k+i)]\}^2$$

where k is the current sampling interval, $k+i$ is a future sampling interval (within the prediction horizon), P is the number of control intervals in the

prediction horizon, n_y is the number of plant outputs, w_j^y is the weight for output j , and the term $[r_j(k+i) - y_j(k+i)]$ is a predicted deviation for output j at interval $k+1$.

The weights must be zero or positive. If a particular weight is large, deviations for that output dominate $S_y(k)$. One of the controller's objectives is to *minimize* $S_y(k)$. Thus, a large weight on a particular output causes the controller to minimize deviations in that output (relative to outputs having smaller weights).

For example, the default values used to produce Plant Outputs for T Setpoint Scenario with Added Data Markers on page 3-13 specify equal weights on each output, so the controller is trying to eliminate deviations in both, which is impossible. On the other hand, the design of Improved Setpoint Tracking for CSTR Temperature on page 3-26 uses a weight of zero on the second output, so it is able to eliminate deviations in the first output.


Note The second output is unmeasured. Its predictions rely on the plant model and the temperature measurements. If the model were reliable, we could hold the predicted concentration at a setpoint, allowing deviations in the reactor temperature instead. In practice, it would be more common to control the temperature as done here, using the predicted reactant concentration as an auxiliary indicator.

You might expect equal output weights to result in equal output deviations at steady state. Plant Outputs for T Setpoint Scenario with Added Data Markers on page 3-13 shows that this is not the case. The reason is that the controller is trying to satisfy several additional objectives simultaneously.

Rate Weights. One is to minimize the weighted sum of controller adjustments, calculated according to

$$S_{\Delta u}(k) = \sum_{i=1}^M \sum_{j=1}^{n_{my}} \{w_j^{\Delta u} \Delta u_j(k+i-1)\}^2$$

where M is the number of intervals in the control horizon, n_{mv} is the number of manipulated variables, $\Delta u_j(k+i-1)$ is the predicted adjustment in manipulated variable j at future (or current) sampling interval $k+i-1$, and $w_j^{\Delta u}$ is the weight on this adjustment, called the *rate weight* because it penalizes the incremental change rather than the cumulative value. Increasing this weight forces the controller to make smaller, more cautious adjustments.



Plant Inputs for Modified Rate Weight

Plant Inputs for Modified Rate Weight on page 3-28 shows that increasing the rate weight from 0.1 to 0.3 decreases the move sizes significantly, especially the initial move (compare to Plant Inputs for the T Setpoint Scenario on page 3-14).

Setting the rate weight to 0.1 yields a faster approach to the T setpoint with a small overshoot, but the initial T_c move is about six times larger than needed to achieve the new steady state, which would be unacceptable in most applications (not shown; try it).

Note The controller minimizes the sum $S_y(k) + S_{\Delta u}(k)$. Changes in the coolant temperature have unequal effects on the two outputs, so the steady-state output deviations won't be equal, even when both output weights are unity as in Plant Outputs for T Setpoint Scenario with Added Data Markers on page 3-13.

Input Weights. The controller also minimizes the weighted sum of manipulated variable deviations from their nominal values, computed according to

$$S_u(k) = \sum_{i=1}^M \sum_{j=1}^{n_{mv}} \{w_j^u [u_j(k+i-1) - \bar{u}_j]\}^2$$

where w_j^u is the *input weight* and \bar{u}_j is the nominal value for input j . In the above simulations, you used the default, $w_j^u = 0$. This is the usual choice.

When a sustained disturbance or setpoint change occurs, the manipulated variable must deviate permanently from its nominal value (as shown in Plant Inputs for the T Setpoint Scenario on page 3-14 and Plant Inputs for Modified Rate Weight on page 3-28). Using a nonzero input weight forces the corresponding input back toward its nominal value. Test this by running a simulation in which you set the input weight to 1. The final T_c value is closer to its nominal value, but this causes T to deviate from the new setpoint (not shown).

Note Some applications involve more manipulated variables than plant outputs. In such cases, it is common to define nonzero input weights on certain manipulated variables in order to hold them near their most economical values. The remaining manipulated variables eliminate steady-state error in the plant outputs.

Blocking

The section “Weight Tuning” on page 3-25 used *penalty weights* to shape the controller’s response. This section covers the following topics:

- An alternative to penalty weighting, called *blocking*
- Side-by-side controller comparisons

To begin, select **Controllers** in the tree, and click the **New** button, creating a controller initialized to the Model Predictive Control Toolbox default settings. Rename this controller **Blocking 1** by editing the appropriate table cell.

Select **Blocking 1** in the tree, select its **Weight Tuning** tab, and set the **Weight** for output **C_A** to 0 (see “Weight Tuning” on page 3-25 to review the reason for this). Leave other weights at their defaults.

Now select the **Model and Horizons** tab, and select its **Blocking** check box. This activates the blocking options. It also deactivates the **Control Horizon** option (the blocking options override it).

Set **Number of moves computed per step** to 2. Verify that **Blocking allocation within prediction horizon** is set to Beginning, the default.

Select **Controllers** in the tree, and use its **Copy** button to create two controllers based on **Blocking 1**. Rename these **Blocking 2** and **Blocking 3**. Edit their blocking options, setting **Blocking allocation within prediction horizon** to Uniform for **Blocking 2**, and to End for **Blocking 3**.

Select **Scenarios** in the tree. Rename the **T Setpoint** scenario to **T Setpoint 1**, and set its **Controller** option to **Blocking 1**.

Create two copies of **T Setpoint 1**, naming them **T Setpoint 2** and **T Setpoint 3**. Set their **Controller** options to **Blocking 2** and **Blocking 3**, respectively. Now you should have three scenarios, identical except for the controller being used.

Delete the **MPC1** controller and select **T Setpoint 1** in the tree. Your view should resemble T Setpoint 1 Scenario on page 3-31, with three controllers and three scenarios in the tree.

The screenshot displays the MPC Design Tool interface. On the left is a workspace tree with the following structure:

- Workspace
 - MPC Design Task
 - Plant models
 - Controllers
 - Blocking 1
 - Blocking 2
 - Blocking 3
 - Scenarios
 - T_Setpoint 1
 - T_Setpoint 2
 - T_Setpoint 3

The main panel shows the following settings:

Simulation settings

- Controller: Blocking 1
- Plant: CSTR
- Duration: 30
- Control interval: 1
- Close loops
- Enforce constraints

Setpoints

Name	Units	Type	Initial Val...	Size	Time	Period	Look Ahe...
T	Deg C	Step	0.0	2	5		<input type="checkbox"/>
C_A	kmol/m ³	Constant	0				<input type="checkbox"/>


Unmeasured disturbances

Name	Units	Type	Initial Value	Size	Time	Period
C_A_i		Constant	0.0			
T	Deg C	Constant	0.0			
T_c	Deg C	Constant	0.0			


T Setpoint 1 Scenario

If any simulation plot windows are open, close them. This forces subsequent simulations to generate new plots.

Simulate each of the three scenarios. When you run the first, new plot windows open. Leave them open when you run the other two scenarios so all three results appear together, as shown in Blocking Comparison, Outputs on page 3-32 and Blocking Comparison, Manipulated Variable on page 3-32.



Blocking Comparison, Outputs



Blocking Comparison, Manipulated Variable

The numeric annotations on these figures refer to the three scenarios. Recall that **T Setpoint 1** uses the default blocking, which usually results in faster setpoint tracking but larger manipulated variable moves. The blocking options in **T Setpoint 2** and **T Setpoint 3** reduce the move size but make setpoint tracking more sluggish.

Results for **T Setpoint 3** are very similar to those shown in Improved Setpoint Tracking for CSTR Temperature on page 3-26 and Plant Inputs for Modified Rate Weight on page 3-28, where a penalty *rate weight* reduced the move sizes. If rate weights and blocking achieve the same ends, why does the toolbox provide both features? One difference not evident in this simple

problem is that blocking applies to all the manipulated variables in your application, but each rate weight affects one only.

Note To obtain the dashed lines shown for **T Setpoint 2**, activate the plot window and select **Property Editor** from the **View** menu. The Property Editor appears at the bottom of the window. Then select the curve you want to edit. The Property Editor lets you change the line type, thickness, color, and symbol type. Select the axis labels to see additional options.

By default, the toolbox plots each scenario on the same plot. If you recalculate a revised scenario, it replots that result but doesn't change any others.

If you don't want to see a particular scenario, right-click the plot and use the **Responses** menu option to hide it. (You can also close the plot window and recalculate selected responses in a fresh window.)

Defining Manipulated Variable Constraints

Physical devices have limited ranges and rates of change. For example, the CSTR model's coolant might be restricted to a 20 degree range (from -10 to 10) and its maximum rate of change might be ± 4 degrees per control interval. If these are true physical restrictions, it's good practice to include them in the controller design. Otherwise the controller might attempt an unrealistic adjustment.

To compare constrained and unconstrained performance for the CSTR example, select the **Blocking 1** controller in the tree, and rename it **Unconstrained**.

Select its **Model and Horizons** tab and turn off (clear) its **Blocking** option. Increase **Control horizon** to 3, and reduce **Control Interval** to 0.25.

Delete the **Blocking 2** and **Blocking 3** controllers. (Click **Yes** or **OK** to dismiss the resulting warning messages.)

Copy the **Unconstrained** controller. Name this copy **MVconstraints**, and select its **Constraints** tab. Then enter the manipulated variable constraints shown in Entering CSTR Manipulated Variable Constraints on page 3-34.

Model and Horizons					
Constraints					
Weight Tuning					
Estimation (Advanced)					
Constraints on manipulated variables					
Name	Units	Minimum	Maximum	Max Down Rate	Max Up Rate
T_c	Deg C	-10	10	-4	4

Entering CSTR Manipulated Variable Constraints

If any simulation plot windows are open, close them (to force fresh plots).

Select the **T Setpoint 1** scenario. If necessary, set its **Controller** option to **Unconstrained**. Change **Duration** to 15, and simulate the scenario.


Select the **T Setpoint 2** scenario, set its **Controller** option to **MVconstraints**, change its **Duration** to 15, and simulate it. The results appear in CSTR Outputs, Unconstrained (1) and MVconstraints (2) on page 3-35 and CSTR Manipulated Variable, Unconstrained (1) and MVconstraints (2) on page 3-35.

The larger control horizon and smaller control interval cause the unconstrained controller to make larger moves (see CSTR Manipulated Variable, Unconstrained (1) and MVconstraints (2) on page 3-35, curve 1). The output settles at the new setpoint in about 5 time units rather than the 10 required previously (compare curve 1 in CSTR Outputs, Unconstrained (1) and MVconstraints (2) on page 3-35 to curve 1 in Blocking Comparison, Outputs on page 3-32).


CSTR Manipulated Variable, Unconstrained (1) and MVconstraints (2) on page 3-35 (curve 2) shows that the **Max Up Rate** constraint limits the size of the first two moves to 4 degrees. The third move hits the **Maximum** constraint at 10 degrees. The coolant temperature remains saturated at its upper limit for the next 7 control intervals, then slowly moves back down to its final value.

CSTR Outputs, Unconstrained (1) and MVconstraints (2) on page 3-35 (curve 2) shows that the output response is slower, but still settles at the

new setpoint smoothly within about 5 time units. This demonstrates the *anti-windup* protection provided automatically by the Model Predictive Control Toolbox controller.



CSSTR Outputs, Unconstrained (1) and MVconstraints (2)



CSSTR Manipulated Variable, Unconstrained (1) and MVconstraints (2)

Disturbance Modeling and Estimation

The previous sections tested the controller's response to setpoint changes. In process control, disturbance rejection is often more important. The Model Predictive Control Toolbox product allows you to tailor the controller's disturbance response. The following example assumes that you have just completed the previous section, and the design tool is still open.

Select the first controller in the tree and rename it **InputSteps**. (Its settings should be identical to the **Unconstrained** controller of the previous section.)

Copy this controller. Rename the copy **OutputSteps**. Click its **Estimation** tab. The initial view should be as in Default Output Disturbance Settings for CSTR on page 3-37. Note the following:

- Model Predictive Control Toolbox default settings are being used. (The **Use Model Predictive Control Defaults** button restores these settings if you modify them.)
- The **Output Disturbances** tab is selected, and the **Signal-by-signal** option is selected. The graphic shows that the output disturbances add to each output.
- The tabular entries show, however, that these disturbance magnitudes are currently zero.

Model and Horizons | Constraints | Weight Tuning | Estimation (Advanced)

Overall estimator gain

Low gain High gain

Value: 0.5

Output Disturbances | Input Disturbances | Measurement Noise

Name	Units	Type	Magnitude
T	Deg C	White	0.0
C_A	kmol/m ³	White	0.0

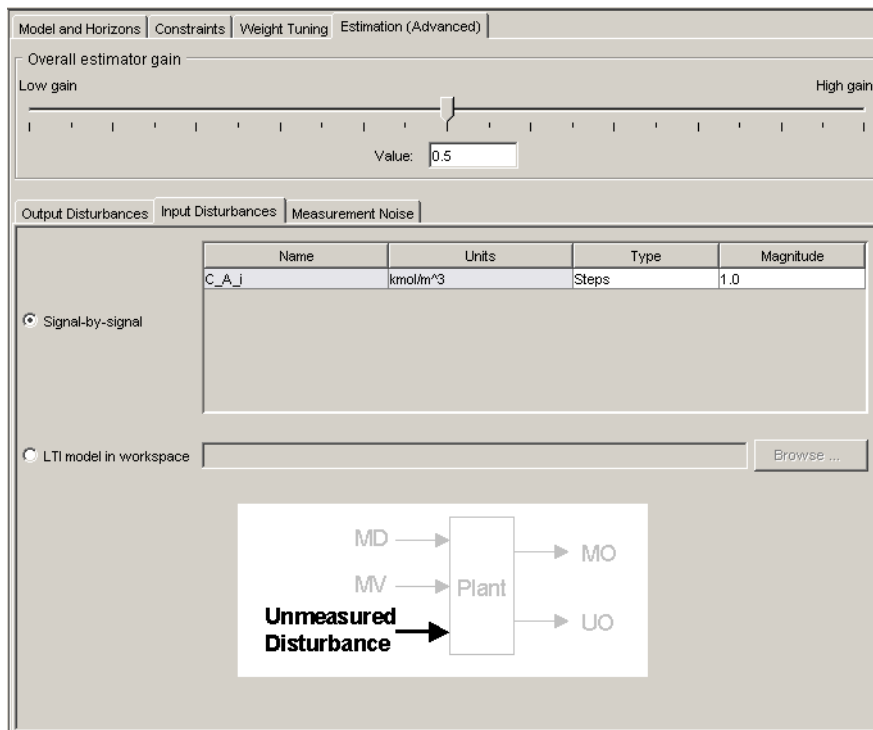
Signal-by-signal

LTI model in workspace Browse ...

The diagram shows a block labeled 'Plant' with three input arrows from the left labeled MD, MV, and UD. The output of the plant is a summing junction with a plus sign. An arrow labeled 'Unmeasured Disturbance' points down to the summing junction with a plus sign. The output of the summing junction is labeled 'Outputs'.

Default Output Disturbance Settings for CSTR

Click the **Input Disturbances** tab. (This would be inactive if the plant model had no unmeasured disturbances.) The view should change to that shown in Default Input Disturbance Settings for CSTR on page 3-38.



Default Input Disturbance Settings for CSTR

In this case the disturbance magnitude is nonzero, and the disturbance type is **Steps**. Thus, the controller assumes that disturbances enter as random steps (integrated white noise) at the plant's unmeasured disturbance input.

Click the **Measurement Noise** tab, verifying that the controller is assuming white noise, magnitude 1 (not shown).

The following summarizes Model Predictive Control Toolbox default disturbance modeling assumptions for this case:

- Additive output disturbances: none
- Unmeasured input disturbances: random steps (integrated white noise)
- Measurement noise: white

In general, if your plant model includes unmeasured disturbance inputs, the toolbox default strategy will assume that they are dominant and sustained, as in the above example. This forces the controller to include an integrating mode, intended to eliminate steady-state error.

If the plant model contains no unmeasured input disturbances, the toolbox assumes sustained (integrated white noise) disturbances at the measured plant outputs.

If there are more measured outputs than unmeasured input disturbances, it assumes sustained disturbances in both locations according to an algorithm described in the product's online documentation.

In any case, the design tool displays the assumptions being used.

To modify the estimation strategy in the `OutputSteps` controller, do the following:

- Click the **Input Disturbances** tab. Set the disturbance **Type** to `White`, and its **Magnitude** to 0.
- Click the **Output Disturbances** tab. For the T output, set the disturbance **Type** to `Steps`, and its magnitude to 1.

This *reverses* the default assumptions, i.e., the **OutputSteps** controller assumes that sustained disturbances enter at the plant output rather than at the unmeasured disturbance input. The **InputSteps** controller is still using the original (default) assumptions.

Next, select the first scenario in the tree. Rename it **Disturbance 1**, set its **Duration** to 30, define all setpoints as constant zero values, and define a unit-step disturbance in the unmeasured input, `C_A_i`. If necessary, set its **Controller** option to `InputSteps`. CSTR Disturbance 1 Scenario on page 3-40 shows the final **Disturbance 1** scenario.

Simulation settings

Controller: Close loops:

Plant: Enforce constraints:

Duration: Control interval: 0.25

Setpoints

Name	Units	Type	Initial Value	Size	Time	Period	Look Ahead
T	Deg C	Constant	0.0				<input type="checkbox"/>
C_A	kmol/m ³	Constant	0.0				<input type="checkbox"/>

Unmeasured disturbances

Name	Units	Type	Initial Value	Size	Time	Period
C_A_i	kmol/m ³	Step	0.0	1.0	5	
T	Deg C	Constant	0.0			
T_c	Deg C	Constant	0.0			


CSTR Disturbance 1 Scenario

Copy **Disturbance 1**. Rename the copy **Disturbance 2**, and set its **Controller** option to OutputSteps.


If necessary, close any open simulation plot windows. Simulate both scenarios. CSTR Outputs for Disturbance Scenarios 1 and 2 on page 3-41 and CSTR Inputs for Disturbance Scenarios 1 and 2 on page 3-41 show the results.

CSTR Outputs for Disturbance Scenarios 1 and 2 on page 3-41 shows that default controller (case 1) returns to the setpoint in less than one third the time required by the modified controller (case 2). Its maximum deviation from the setpoint is also 10% smaller. CSTR Inputs for Disturbance Scenarios 1 and 2 on page 3-41 shows that in both cases the input moves are smooth and of reasonable magnitude. (It also shows the input disturbance.)

The default controller expects unmeasured disturbances to enter as defined in the scenarios, so it's not surprising that the default controller performs better. The point is that the difference can be large, so it merits design consideration.




CSTR Outputs for Disturbance Scenarios 1 and 2



CSTR Inputs for Disturbance Scenarios 1 and 2

For comparison, reset the two scenarios so that the only disturbance is a one-degree step increase added to the measured reactor temperature. The modified controller (case 2) is designed for such disturbances, and CSTR Outputs, Output Disturbance Scenarios 1 and 2 on page 3-42 shows that it performs better, but the difference is less dramatic than in the previous scenario. The default controller is likely to be best if the real process has multiple dominant disturbance sources.



CSTR Outputs, Output Disturbance Scenarios 1 and 2

Define Soft Output Constraints

The discussion in “Weight Tuning” on page 3-25, defined temperature control as the primary goal for the CSTR application. The predicted (but unmeasured) reactant concentration, C_A , could vary freely.


Suppose this were acceptable provided that C_A stayed below a specified maximum (above which unwanted reactions would occur). You can use an *output constraint* to enforce this specification.

Start with a single controller identical to the **InputSteps** controller described in “Disturbance Modeling and Estimation” on page 3-35. Rename it **Unconstrained**.

Right-click **Unconstrained** in the tree and select **Copy**. Rename the copy **Yhard**. Make another copy, naming it **Ysoft**.

Similarly, start with a single scenario identical to CSTR Disturbance 1 Scenario on page 3-40, except that its **Controller** setting should be Unconstrained. Name this scenario **None**.

Right-click **None** in the tree and select **Copy**. Rename the copy **Hard**, and change its **Controller** setting to Yhard. Make another copy, naming it **Soft** and changing its **Controller** setting to Ysoft. Your tree should be as shown below.



Select your **Yhard** controller. On the **Constraints** tab, set the maximum for C_A to 3 as shown below.

Constraints on output variables			
Name	Units	Minimum	Maximum
T	Deg C		
C_A	kmol/m ³		3

Click the **Constraint Softening** button to open the dialog box in Constraint Softening Dialog Box on page 3-44. The **Input constraints** section shows the bounds on the inputs and their rates, and *relaxation bands*, which let you *soften* these constraints. By default, input constraints are *hard*, meaning that the controller tries to prevent any violation.

Specify relaxation bands

Input constraints

Name	Units	Minimum	Min Band	Maximum	Max Band	Max Down...	Max Down...	Max Up Rate	Max Up Ba...
T_c	Deg C	-10		10		-4		4	

Output constraints

Name	Units	Minimum	Min Band	Maximum	Max Band
T	Deg C				
C_A	kmol/m ³			3	0

Overall constraint softness

Soft constraints Hard constraints


Value:

Constraint Softening Dialog Box

The **Output constraints** section lists the output limits and their relaxation bands. By default, the output constraints are *soft*. Make the C_A upper limit hard by entering a zero as its relaxation band (as in Constraint Softening Dialog Box on page 3-44).

Select the **Ysoft** controller. Define a soft upper bound on C_A by using the same settings shown in Constraint Softening Dialog Box on page 3-44, but with a relaxation band of 100 instead of 0.

Simulate the three scenarios in the order they appear in the tree, i.e., **None**, **Hard**, **Soft**. The resulting output responses appear in Constraint Softening Scenarios: 1 = None, 2 = Hard, 3 = Soft on page 3-45.



Constraint Softening Scenarios: 1 = None, 2 = Hard, 3 = Soft

Curve 1 is without output constraints, which is identical to curve 1 in CSTR Outputs for Disturbance Scenarios 1 and 2 on page 3-41. This controller allows the C_A output to vary freely, but the controlled temperature returns to its setpoint within 10 time units after the disturbance happens.

Curve 2 shows the behavior with a hard upper limit at $C_A = 3$. Once C_A reaches this bound, the controller must use its one manipulated variable (T_c) to satisfy the constraint, so it's no longer able to control the temperature.

Curve 3 shows the result for a soft constraint. The controller reaches a compromise between the competing objectives: temperature control and constraint satisfaction. As you'd expect, performance lies between the curve 1 and curve 2 extremes.

The numerical value of the relaxation band represents a *relative tolerance* for constraint violations, not a strict limit (if it were the latter, it would be a hard constraint). If you were to increase its relaxation band (currently set at 100), performance would move toward Case 1, and vice versa.

If you have multiple constraints, you can harden or soften them simultaneously using the slider at the bottom of the controller's constraint softening dialog box (see Constraint Softening Dialog Box on page 3-44).

In general, you'll have to experiment to determine the settings that provide appropriate trade-offs for your application. In particular, the relaxation band settings interact with those on the controller's **Weight Tuning** tab (see "Weight Tuning" on page 3-25).

Another important factor is the expected numerical range for each variable. For example, if a particular variable stays within ± 0.1 of its nominal value, it should have a small relaxation band relative to another variable having a range of ± 100 .

For details on the Model Predictive Control Toolbox constraint softening formulation, see "Optimization Problem".

Save Your Work

You'll usually want to save your design so you can reuse or revise it. You can save individual controllers or an entire project.

When you close the design tool, you'll be prompted to save new or modified designs. You can also save manually to preserve an intermediate state or guard against an unexpected shutdown.

This section covers the following topics:

- "Exporting a Controller" on page 3-46
- "Saving a Project" on page 3-47

Exporting a Controller

To save a controller, *export* it to your MATLAB workspace or to a MAT-file. The former allows you to use the exported controller in command-line functions or a Simulink block.

Note Your workspace disappears when you exit MATLAB. A MAT-file is permanent, and you can reload it in a subsequent MATLAB session.

The following example assumes that the design tool is open and in the state described in the previous section. Suppose you want to export the controller

to your workspace. Expand the tree if necessary, right-click **MPC2**, and select **Export Controller** from the resulting menu. The following dialog box appears.



The default behavior is to export the selected controller to the workspace. Click **Export** to confirm. You can verify the export by typing

`whos`


at the MATLAB prompt. The resulting list should include an `mpc` object named `MPC2`. Type

`MPC2`


to display the object's properties.

Saving a Project

To save your entire project, click the toolbar's **Save** button.



The following dialog box appears.



Dialog Box for Saving a Controller Design Project

The default behavior saves the current project (named `Project - CSTR_MPC` in this case) in a MAT-file (called `CSTR_MPC` here). If the design tool had contained other projects, they would appear in the list, and you could select the ones you wanted to save.

The MAT-file will be saved in the default folder. To verify the location, click the **Browse (...)** button and change the folder if necessary. When ready to save, click **OK**.

Load Your Saved Work

The following assumes that you've saved a project as described in the previous section. To reload this project, close the design tool if it's open. Also clear any `mpc` objects from your workspace. (Type `whos` at the MATLAB prompt for a list of objects.) For example, if `MPC1` and `MPC2` are in your workspace, type

```
clear MPC1 MPC2
```

to clear (remove) them.

If you've closed the `CSTR_MPC` model, open it. (Simulink Model for CSTR Control on page 3-61 shows the model diagram). Double-click the MPC Controller block to open its mask, and verify that the **MPC Controller** parameter is set to `MPC2`.

Note If you were to attempt to run the CSTR_MPC model at this stage, an error dialog box would indicate that the MPC Controller block was unable to initialize. The MPC2 object specified in the block mask must be loaded into your workspace or be part of an active design tool task.


You could define the required MPC2 object in one of the following ways:

- Import MPC2 from a MAT-file (assuming you had saved it as explained in “Exporting a Controller” on page 3-46).
- Load the model’s project file, which contains a copy of MPC2.

To use the second approach, open the design tool by typing

```
mpctool
```

in the MATLAB Command Window. This creates a blank Model Predictive Control Toolbox project called **MPC Design Task**. Click **Load** on the toolbar.



This opens a dialog box similar to that shown in Dialog Box for Saving a Controller Design Project on page 3-48. Use it to select the project file you’ve saved, and then click **OK** to load the project. It should appear in the tree. Verify that it contains a controller named MPC2.

Run the CSTR_MPC model in Simulink. The block mask automatically retrieves **MPC2** from the design tool, and the simulation runs. In other words, loading the project automatically restores the link between the design tool and the MPC Controller block.

Test Controller Robustness

It's good practice to test your controller's sensitivity to prediction errors. Classical phase and gain margins are one way to quantify robustness for a SISO application. Robust Control Toolbox™ software provides sophisticated approaches for MIMO systems. It can also be helpful to run simulations. The following example illustrates the simulation approach.

Plant Model Perturbation

Create a perturbed version of a model.

For this example, use the nonlinear CSTR model discussed in “CSTR Model” on page 2-15.

```
CSTRp = CSTR;  
CSTRp.a=[ -0.0303   -0.0113  
          -0.0569   -0.1836];  
CSTRp.b=[ -0.0857   0.0191  
          0.1393    0.4241];
```

This creates a copy of CSTR called CSTRp, then replaces the state space A and B matrices with perturbed versions (compare to the originals defined in “State-Space Models” on page 2-5). Use the following command to compare the two step responses:

```
step(CSTR, CSTRp)
```

Observe the difference in the responses (not shown).

Select **Plant models** in the tree. Click the **Import** button and import the CSTRp model.

Simulation Tests

Delete all controllers except the first one in the tree. If necessary, make its settings identical to **Unconstrained** (see “Defining Manipulated Variable Constraints” on page 3-33).

Delete all scenarios except the first, naming that **Accurate Model**. Define its properties as shown in Robustness Test, Accurate Plant Model Scenario on page 3-51. The scenario begins with a step change in the temperature setpoint, followed 25 time units later by a step disturbance in the reactant entering the CSTR.

Copy **Accurate Model**. Rename the copy **Perturbed Model**, and set its **Plant** option to CSTRp. Thus, both scenarios use the same controller, which is based on the CSTR model, but the **Perturbed Model** scenario uses a different model to represent the “real” plant. This tests the controller’s robustness to a change in plant parameters.

Simulation settings

Controller: Unconstrained Close loops:

Plant: CSTR Enforce constraints:

Duration: 50 Control interval: 0.25

Setpoints

Name	Units	Type	Initial Value	Size	Time	Period	Look Ahead
T	Deg C	Step	0.0	1	1		<input type="checkbox"/>
C_A	kmol/m ³	Constant	0.0				<input type="checkbox"/>


Unmeasured disturbances

Name	Units	Type	Initial Value	Size	Time	Period
C_A_i	kmol/m ³	Step	0.0	1.0	25	
T	Deg C	Constant	0.0			
T_c	Deg C	Constant	0.0			

Robustness Test, Accurate Plant Model Scenario

Simulate the two scenarios. Robustness Test, Accurate Model (1) and Perturbed Model (2) on page 3-52 shows the output responses. As expected, setpoint tracking degrades when the model is inaccurate, but performance is still acceptable.

The disturbance rejection appears to *improve* with the perturbed model. This is a consequence of the perturbations used. The gain for the T/C_{Ai} output/input pair is about 15% smaller in the CSTRp model, which has two beneficial effects: the actual impact of the disturbance is reduced, and the controller is aggressive because it expects a larger impact.



Robustness Test, Accurate Model (1) and Perturbed Model (2)

Note MIMO applications are usually more sensitive to model error than SISO applications, so robustness testing is especially recommended for MIMO cases.

Design Controller for Plant with Delays

Unlike many controller design approaches, the Model Predictive Control Toolbox can handle models that include delays. A typical example is the distillation column model, DC, introduced in “Specify Multi-Input Multi-Output (MIMO) Plants” on page 2-12. The model includes a delay in each input/output channel.

The presence of delays will influence controller performance, and your controller specifications should account for them. The following example provides some guidelines and covers the following topics:

- “Importing the Plant Model” on page 3-53
- “Specifying Controller Horizons” on page 3-54

Importing the Plant Model

To learn how to create the DC model, see “Specify Multi-Input Multi-Output (MIMO) Plants” on page 2-12. You must import the DC model into the MATLAB workspace.

Start the Model Predictive Control Toolbox design tool (type `mpctool` at the MATLAB command line). Import the DC model (see “Load a Plant Model” on page 3-3).

Select **Plant models** in the tree. The DC model should be the only one listed. Scroll the **Model details** view to show the last few lines, which should appear as in DC Model Details, Maximum Delay Values on page 3-53.



DC Model Details, Maximum Delay Values

Specifying Controller Horizons

The model's maximum I/O delay is 6.7 minutes. It is good practice to specify the prediction and control horizons such that


$$P - M \geq t_{d,max} / \Delta t$$

where P is the prediction horizon, M is the control horizon, $t_{d,max}$ is the maximum delay, and Δt is the control interval.


Select **MPC1** (the default controller name) in the tree. Click the **Model and Horizons** tab, and set **Control interval** to 1, a reasonable choice if the closed-loop response time is to be of order 5-10 minutes.

Given the amount of plant delay and the specified control interval, the default horizons, $P = 10$, $M = 2$, would be a poor choice. Instead, set **Prediction horizon** to 30, and **Control horizon** to 5.

Select **Scenario1** in the tree. Set **Duration** to 50. Define a constant setpoint of 1 for the first output (the distillate purity). Define a step increase of 1 in the second output's setpoint, occurring at $t = 25$. All other signals should be zero. Simulate the scenario. DC Setpoint Response Scenario, Outputs on page 3-55 and DC Setpoint Response Scenario, Inputs on page 3-55 show the results.



DC Setpoint Response Scenario, Outputs



DC Setpoint Response Scenario, Inputs

As seen in DC Setpoint Response Scenario, Outputs on page 3-55, the first output cannot respond for a minimum of one minute, the delay in the y_1/u_1 transfer function. After that, it reaches the setpoint in two minutes and

settles quickly. Similarly, y_2 cannot respond for a minimum of three minutes, the delay in the y_2/u_2 transfer function, but settles rapidly thereafter. Changing one setpoint disturbs the other output, but the magnitude of this *interaction* is less than 10%.


DC Setpoint Response Scenario, Inputs on page 3-55 shows that the initial input moves are more than five times the final change. Also, there are periodic pulses in the control action as the controller attempts to counteract the delayed effects of each input on the two outputs.

You can moderate these effects using the weights (see “Weight Tuning” on page 3-25). Instead, define a custom blocking strategy as illustrated in DC Model, Custom Blocking Strategy on page 3-56. This uses five moves as before, but allocates them more uniformly over the prediction horizon.


Output Responses for Setpoint Scenario with Blocking on page 3-57 and Input Moves for Setpoint Scenario with Blocking on page 3-57 show the corresponding simulation results. The initial input moves are much smaller, and the moves are less oscillatory overall. The trade-off is a slower output response with about 20% interaction.

The screenshot shows the 'Weight Tuning' tab of the Design Tool GUI. The 'Plant model' is set to 'DC'. Under the 'Horizons' section, the 'Control interval (time units)' is 1.0, the 'Prediction horizon (intervals)' is 30, and the 'Control horizon (intervals)' is 5. The 'Blocking' checkbox is checked. In the 'Blocking' section, the 'Blocking allocation within prediction horizon' is set to 'Custom', the 'Number of moves computed per step' is 5, and the 'Custom move allocation vector' is [5 5 5 5 10].

DC Model, Custom Blocking Strategy



Output Responses for Setpoint Scenario with Blocking



Input Moves for Setpoint Scenario with Blocking

Design Controller for Nonsquare Plant

A *nonsquare* plant has an unequal number of manipulated variables and output variables. This is common in practice, and the Model Predictive Control Toolbox software supports an excess of manipulated variables or outputs. In such cases you will usually need to modify default toolbox settings.

This section covers the following topics:

- “More Outputs Than Manipulated Variables” on page 3-58
- “More Manipulated Variables Than Outputs” on page 3-59

More Outputs Than Manipulated Variables

When there are excess outputs, you can’t hold each at a setpoint. You have the following options:

- Enforce setpoints on all outputs, in which case all will deviate from their setpoints to some extent.
- Specify that certain outputs need not be held at setpoints by setting their weights to zero on the controller’s **Weight Tuning** tab.

The initial test of the CSTR controller used option 1 (the default), which caused both outputs to deviate from their setpoints (see Plant Outputs for T Setpoint Scenario with Added Data Markers on page 3-13). You can adjust the offset in each output by changing the output weights. Increasing an output weight decreases the offset in that output (at the expense of increased offset in other outputs).

The modified CSTR controller used option 2 (see the discussion in “Weight Tuning” on page 3-25). In general, if the application has N_e more outputs than manipulated variables, setting N_e output weights to zero should allow the remaining outputs to be held at setpoints (unless the manipulated variables are constrained). This was the case for the modified CSTR controller (see Improved Setpoint Tracking for CSTR Temperature on page 3-26).

Outputs that have been “sacrificed” by setting their weights to zero can still be useful. If measured, they can help the controller to estimate the plant’s state,

thereby improving its predictions. They can also be used as indicators, or as variables to be held within an operating region defined by output constraints.

More Manipulated Variables Than Outputs

In this situation, default Model Predictive Control Toolbox settings should provide offset-free output-setpoint tracking, but the manipulated variables are likely to drift.

One way to avoid this is to use *manipulated variable setpoints*. If there are N_e excess manipulated variables and you hold N_e of them at target values, the rest should not drift. Rather, they will attain the values needed to eliminate output offset.

To define a manipulated variable setpoint:

- 1** Enter the setpoint value in the **Nominal** field in the signal properties view – see Model Predictive Control Toolbox Design Tool’s Signal Definition View on page 3-5.
- 2** Assign a nonzero *input weight* using the **Weight** entry on the controller’s **Weight Tuning** tab – see Controller Options — Weight Tuning Tab on page 3-25.

In step 2, the magnitude of the input weight determines the extent to which the manipulated variable can deviate from its target during a transient. See “Input Weights” on page 3-29 for more discussion and mathematical details.

You might want to allow such deviations temporarily in order to provide better output setpoint tracking. In that case, use a relatively small input weight. If you want the manipulated variable to stay near its target value at all times, increase its input weight.

Another way to avoid drift is to constrain one or more manipulated variables to a narrow operating region. You can even hold an MV constant by setting its lower and upper bounds to the same value (in which case its nominal value should also be set to this value), or by setting both of its rate constraints to zero. To define constraints, use the controller’s **Constraints** tab (see “Defining Manipulated Variable Constraints” on page 3-33).

Nonlinear Plants

You can control a Simulink plant modeled using a Model Predictive Control Toolbox Simulink block. This section illustrates the procedure using the nonlinear CSTR model introduced earlier (see “Linearize Simulink Models” on page 2-17).

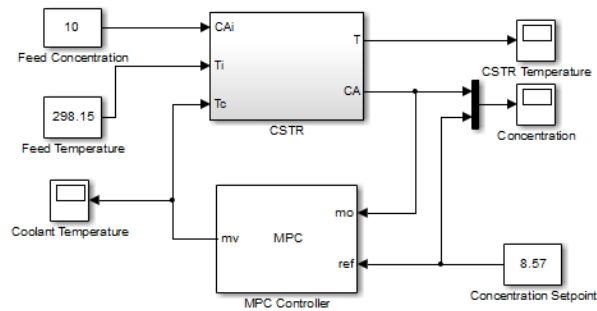
In this section...
“MPC Controller Block” on page 3-60
“Initiating the Controller Design” on page 3-61
“Validating the Linearized Model” on page 3-63
“Modifying the Linearized Model” on page 3-65
“Linear Simulation Tests” on page 3-67
“Nonlinear Simulation Tests” on page 3-68
“Modifying the Controller Using the Design Tool” on page 3-70
“Exiting the Design Tool” on page 3-70

MPC Controller Block

Simulink Model for CSTR Control on page 3-61 shows the CSTR block with an added MPC Controller block. The intent is to hold the residual concentration at a setpoint by adjusting the coolant temperature. The feed concentration and temperature inputs are disturbances, and the reactor temperature is an auxiliary, uncontrolled output. The concentration has a setpoint (reference), currently fixed at 8.57.

To open the model, type the following:

```
open_system('CSTR_MPC');
```




Simulink Model for CSTR Control

Note If you have Model Predictive Control Toolbox software and Simulink installed, the MPC Controller block appears in your Simulink block library. If necessary, see the Simulink documentation for details on how to use library blocks to construct a model.

Initiating the Controller Design

Once you've installed an MPC Controller block in your Simulink model, as in Simulink Model for CSTR Control on page 3-61, you need to configure it. Click it to open its *block mask*, and verify that it is as shown below.



MPC Controller Block Mask

If you had used the design tool to create a controller and had saved it in your workspace, you could type its name in the **MPC controller** box and you'd be ready to test it.

Instead, create the controller using the Model Predictive Control Toolbox design tool. To do so, click the **Design** button. A dialog box (not shown) asks you to enter the number of *manipulated variables*, which are the plant inputs the controller can adjust to achieve its objectives. The coolant temperature is the only manipulated variable. Verify that the dialog box entry is 1 and click **OK**. A progress window marks the completion of each of the following controller initialization steps:


- 1 Open the Control and Estimation Tools Manager, and install a *project* having the same name as the Simulink model (CSTR_MPC in this case).
- 2 Install a *linearization task* within the project and determine the plant's I/O points, i.e., the signals connected to the MPC Controller block's mv (manipulated variable) and mo (measured output) ports.
- 3 Determine a default *operating point*. This requires assumptions that might be incorrect, as demonstrated in the next section. In particular, the default outputs are zero.
- 4 Calculate a *linearized plant model* at this operating point using the linearization tool in Simulink® Control Design (see “Linearization Using Linear Analysis Tool in Simulink Control Design” on page 2-20). The controller is open loop during this step and involves only the blocks between the plant inputs and outputs (as determined in step 2).
- 5 Install a Model Predictive Control Toolbox design task in the project. Use the linearized model from step 4 to define a default controller (named MPC1), and enter its name as the block mask's **MPC controller** parameter.

When the last step has been completed, click **OK** to close the progress window.

Validating the Linearized Model

Automatic linearization requires assumptions that might be incorrect. You should always validate a linearized model before using it in a controller.

As shown below, the controller task is the **MPC Controller** node, the name of the corresponding Simulink block.



The screenshot shows the Simulink workspace on the left and a table of states on the right. The workspace tree includes:

- Workspace
 - Project - CSTR_MPC
 - Operating Points
 - Default Operating Point
 - MPC open loop plant 1
 - MPC Design Task - MPC Controller
 - Plant models
 - MPC open loop plant 1
 - Operating Point
 - Controllers
 - Scenarios

The table on the right shows the states of the MPC Controller block:

State	Desired Value	Actual Value	Desired dx	Actual dx
CSTR_MPC/CSTR/Integrator				
State - 1	[0 , Inf]	229.3546	0	-0.00038873
CSTR_MPC/CSTR/Integrator1				
State - 1	[0 , Inf]	5	0	4.9991
CSTR_MPC/MPC Controller/sfunction				
State - 1	[-Inf , Inf]	0.00011662	0	0

Note In general, a Simulink model can contain multiple controllers, in which case the project would include multiple controller design tasks, each with a unique name.


Start by examining the default operating point. Expand the tree until it appears as above, and then select **Operating Point**. The **States** tab should be active as shown (if not, click it).

By default, the Simulink Control Design linearization tool tries to find a steady-state point, i.e., it sets the **Desired dx** column to zero. As shown above, the **Actual dx** column contains one value that is far from zero (4.9991), i.e., it has failed to achieve this goal.

Note Don't be concerned if you see numerical values that differ slightly from those shown.

Next, click the **Outputs** tab. By default, the linearization tool sets the desired plant output to zero. In the CSTR, this would require 100% conversion of the reactant, which is impossible, and the **Outputs** tab confirms that the desired values were not achieved.

Finally, export the linearized model to your workspace. Right-click the **MPC open loop plant 1** node and choose **Export** to open the following dialog box.




Rename **Linearized Model** to **Plant1** as shown, clear the **Operating Point** check box (because there's no reason to export it), and click **OK**. This exports the linearized model as an LTI object named **Plant1**.

At the MATLAB command line, type:

step(Plant1)

You should obtain the following plot.



Thus, the linearized model predicts that a unit step *increase* in the coolant temperature will *decrease* the residual concentration, which is qualitatively correct. (Increasing the coolant temperature increases the reactor temperature. This increases the reaction rate, and the residual concentration decreases.)

The predicted magnitude is very small, however: of order 10^{-6} . It should be of order 10^{-2} (you can verify this by removing the controller block from the diagram and running an open-loop step-response simulation). The incorrect assumptions used to generate the default operating point are the cause. If this incorrect model were used in the controller, the coolant-temperature adjustments would be far too large, probably resulting in unstable behavior.

Modifying the Linearized Model

Therefore, the next step is to define a reasonable operating point and calculate a corresponding linearized model.

In the navigation tree, click the **MPC Task - MPC Controller** node. Notice that the nominal input and output signal values are zero by default. As discussed in “Linearize Simulink Models” on page 2-17, the plant is at steady state at a coolant temperature of 298.15 K and a residual concentration of 8.57 kmol/m³. Thus, in the **Input signal properties** table, set the coolant temperature’s nominal value to **298.15**, as shown below, and in the **Output signal properties** table, set the concentration’s nominal value to **8.57**. Also assign more descriptive signal names, as shown.

Input signal properties				
Name	Type	Description	Units	Nominal
Coolant Temperature	Manipulated			298.15
Output signal properties				
Name	Type	Description	Units	Nominal
Concentration	Measured			8.57

Click the **Import Plant** button to open the Plant Model Importer dialog box. Then do the following:

- 1** Click the **Linearized Plant from Simulink** tab.
- 2** Change the **Linearization model name** to Plant2.
- 3** Select **Create a new operating condition from MPC I/O values**, and then click **OK**. A new model node named **Plant2** and its operating point will appear within your design task. Verify that its actual dx values are now all close to zero and the outputs are essentially equal to their desired values.
- 4** Expand the **Controllers** node to expose the **MPC1** node and select it. On its **Model and Horizons** pane, set **Plant model** to Plant2. This replaces the default (invalid) model with the modified one.


Leave the other controller settings at their default values.

Linear Simulation Tests

It's good practice to test a controller in linear simulations before trying it on the nonlinear plant. Expand the **Scenarios** node and select **Scenario1**. Set **Plant** to **Plant2**, and set **Duration** to 40. Finally, define a step increase of 0.5 units in the concentration setpoint starting at time $t = 1$, as shown below.

Setpoints							
Name	Units	Type	Initial Value	Size	Time	Period	Look Ahead
Concentration		Step	8.57	0.5	1.0		<input type="checkbox"/>

Click the scenario's **Simulate** button. This linear test predicts a smooth, rapid approach to the new setpoint with minimal overshoot, as shown below.



The corresponding coolant temperature adjustments are reasonable (not shown). If anything, you might want to make the controller more aggressive by adjusting its tuning weights (see “Weight Tuning” on page 3-25).


Note This plant's steady-state gain is of order 0.01. Therefore, the default tuning weights lead to a relatively sluggish response. In general, you must adjust the tuning weights to compensate for the plant's natural input/output response magnitudes.

You can also run tests to verify that the controller responds rapidly to either of the two unmeasured disturbances (not shown).

Nonlinear Simulation Tests

The controller seems to be performing well in linear tests, so try it with the nonlinear plant. As will be demonstrated below, you can run simulations that automatically use your latest design from the design tool. If the performance is unsatisfactory, you can easily change your design and retest.


In the CSTR_MPC model window, change the concentration setpoint to 9.07. This simulates a 0.5-unit step increase in the setpoint at time $t = 0$. Run the Simulink simulation. The MPC Controller block automatically obtains the most recent **MPC1** definition from the design tool. If necessary, open the model's concentration scope block. The result should be as shown below.



The nonlinear response is more sluggish than the linear prediction. The controller reduces the coolant temperature to about 284.6 K (verify this by opening the model's coolant temperature scope), whereas the linear simulation predicts a reduction to 289.4 K. In other words, when plant

moves in this direction, the linearized model's gain is too large. Still, the concentration goes to the setpoint rapidly.

Next, set the concentration setpoint to 8.07, i.e., a step-change of equal magnitude in the opposite direction. Run the simulation to obtain the concentration scope response shown below.



In this direction, the response is underdamped. You can verify that the controller changes the coolant temperature to 304.1 K, an increase of 6.0 K (recall the decrease of 13.5 K when the change was in the opposite direction). In other words, the controller's linear model underpredicts the effect of a coolant temperature increase. If the setpoint were reduced significantly, the closed-loop system would become unstable (try 7.0, for example).

Thus, the controller's effective operating range is limited. If you wanted to operate at a low concentration, you'd need to determine a linearized model at that condition and use it to design another controller.

Note This is typical of strongly nonlinear plants. If you needed to operate a conventional controller over such a wide range, you might consider gain scheduling. Equivalently, you could define predictive controllers for several operating points and switch from one to another depending on the measured concentration (see “Simulate Controller with Nonlinear Plant” on page 4-9 for an example of this).

For the tested range, however, the oscillations die out quickly. You can verify that the controller responds equally well to small, sustained disturbances (i.e., ± 0.5 in feed concentration or ± 3 in feed temperature).

Modifying the Controller Using the Design Tool

Keeping the design tool open makes it easy to test controller modifications. Each time you run a Simulink simulation, the MPC Controller block automatically obtains the latest settings for its **MPC controller** parameter (**MPC1** in the above tests).

It's also easy to compare several controllers. Suppose you wanted to try a different control interval (the controller designed above uses a 1-second interval). Create a new controller by right-clicking the design tool's **MPC1** controller node, and selecting **Copy Controller** in the menu. This creates a duplicate controller called **MPC1_copy**. Rename it **MPC2**. Select **MPC2** in the tree, activate its **Model and Horizons** tab and set **Control interval** to 2.

Now return to the MPC Controller block mask (if necessary, double-click the block to open the mask) and change **MPC controller** from **MPC1** to **MPC2**. Close the mask or click the **Apply** button to notify the controller of this change. Run a simulation with a disturbance or setpoint change, and then open the **Coolant Temperature** scope. You should see that the controller is now making step-wise adjustments every 2 minutes. As would be expected, controller performance degrades with less frequent measurement feedback, but not much in this case.

To restore the original controller, just change **MPC2** to **MPC1** in the block mask and click **Apply**.

Exiting the Design Tool

When you close the design tool, you'll be prompted to save any projects you've created or edited. For details, see "Saving a Project" on page 3-47.

Once the design tool has closed, the MPC Controller block mask expects its controller object to be in your workspace. Thus, you'll also be prompted to export such controller objects. You need to do this if you plan to run your Simulink model after you've closed the design tool.

Designing Controllers Using the Command Line

- “Design Controller at the Command Line” on page 4-2
- “Simulate Controller with Nonlinear Plant” on page 4-9
- “Control Based On Multiple Plant Models” on page 4-15
- “Compute Steady-State Gain” on page 4-24
- “Extract Controller” on page 4-25
- “Bibliography” on page 4-27

Design Controller at the Command Line

In this section...
“Create a Controller Object” on page 4-2
“View and Alter Controller Properties” on page 4-3
“Perform Linear Simulations” on page 4-6
“Save Calculated Results” on page 4-7

“Design Controller Using the Design Tool” on page 3-2 shows how to use the Model Predictive Control Toolbox design tool to create a controller and test it. You might prefer to use functions instead. They allow access to options not available in the design tool, as well as automation of repetitive tasks using scripts and customized plotting.

Create a Controller Object

This topic uses the CSTR model described in “CSTR Model” on page 2-15 as an example.

Use the `mpc` function to create a controller. For example, type

```
Ts = 1;  
MPCobj = mpc(CSTR, Ts);
```

to create one based on the CSTR model with a control interval of 1 time unit and all other parameters at their default values.

Note MPCobj is an MPC object. It contains a complete controller definition for use with Model Predictive Control Toolbox software.

To display the controller’s properties in the Command Window, type

```
display(MPCobj)
```

or type the object’s name without a trailing semicolon.

View and Alter Controller Properties

Once you've defined an MPC object, it's easy to alter its properties. For a description of the editable properties, type:

```
mpcprops
```

To display the list of properties and their current values, type

```
get(MPCobj)
```

For the CSTR example this displays:

```
ManipulatedVariables (MV): [1x1 struct]
    OutputVariables (OV): [1x2 struct]
    DisturbanceVariables (DV): [1x1 struct]
        Weights (W): [1x1 struct]
            Model: [1x1 struct]
                Ts: 1
            Optimizer: [1x1 struct]
    PredictionHorizon (P): 10
    ControlHorizon: 2
        History: [2e+003 7 21 20 18 20.1]
        Notes: {}
    UserData: []
```

(Your History entry will differ.) To alter one of these properties, you can use the syntax

```
ObjName.PropName = value;
```

where ObjName is the object name, and PropName is the property you want to set. For example, to change the prediction horizon from 10 (the default) to 15, type:

```
MPCobj.P = 15;
```

Note You can abbreviate property names provided that the abbreviation is unambiguous.

As shown above, many of the properties are MATLAB *structures* containing additional properties. For example, type

```
MPCobj.MV
```

which displays:

```
        Min: -Inf
        Max: Inf
    MinECR: 0
    MaxECR: 0
    RateMin: -Inf
    RateMax: Inf
RateMinECR: 0
RateMaxECR: 0
    Target: 'nominal'
      Name: 'T_c'
     Units: ''
```

This shows that the default controller has no constraints on the manipulated variable. To include constraints as shown in [Entering CSTR Manipulated Variable Constraints](#) on page 3-34, you could type

```
MPCobj.MV.Min = -10;
MPCobj.MV.Max = 10;
MPCobj.MV.RateMin = -4;
MPCobj.MV.RateMax = 4;
MPCobj.MV.Units = 'Deg C';
```

or use the set command:

```
set(MPCobj, 'MV', struct('Min', -10, 'Max', 10, ...
    'RateMin', -4, 'RateMax', 4, 'Units', 'Deg C'));
```

Note The `Units` property is for display purposes only and is optional.

There are two outputs in this case, so `MPCobj.OV` is a 1-by-2 structure. To set measurement units to the values shown in Controller Options — Weight Tuning Tab on page 3-25, you could type

```
MPCobj.OV(1).Units = 'Deg C';  
MPCobj.OV(2).Units = 'kmol/m^3';
```

Finally, check the default weights by typing

```
MPCobj.W
```

which displays:

```
ManipulatedVariables: 0  
ManipulatedVariablesRate: 0.1000  
OutputVariables: [1 1]  
ECR: 100000
```

Change to the values shown in Controller Options — Weight Tuning Tab on page 3-25 by typing:

```
MPCobj.W.ManipulatedVariablesRate = 0.3;  
MPCobj.W.OutputVariables = [1 0];
```

You can also specify time-varying weights and constraints. The time-varying weights and constraints are defined for the prediction horizon, which shifts at each time step. This implies that as long as the property is not changed, the set of time-varying parameters is the same at each time step. To learn how to specify time-varying constraints and weights in the GUI, see “Constraints Tab” and “Weight Tuning Tab”.

The time-varying weights modify the tuning of the unconstrained controller response. To specify a different weight for each step in the prediction horizon, modify the `Weight` property. For example,

```
MPCobj.W.OutputVariables = [0.1 0; 0.2 0; 0.5 0; 1 0];
```

deemphasizes setpoint tracking errors early in the prediction horizon. The default weight of 1 is used for the fourth step and beyond.

Constraints can also be time varying. The time-varying constraints have a nonlinear effect when they are active. For example,

```
MPCobj.MV.RateMin=[ -4; -3.5; -3; -2.5]  
MPCobj.MV.RateMax=[ 4; 3.5; 3; 2.5]
```

forces MV to change more and more slowly along the prediction horizon. The constraint of -2.5 and 2.5 is used for the fourth step and beyond.

You could also alter the controller's disturbance rejection characteristics using functions that parallel the design tool's disturbance modeling options (described in "Disturbance Modeling and Estimation" on page 3-35). See the reference pages for the `setestim`, `setindist`, and `setoutdist` functions.

Perform Linear Simulations

To run a linear simulation, use the `sim` function. For example, given the MPCobj controller defined in the previous section, type:

```
T = 26;  
r = [2 0];  
sim(MPCobj, T, r);
```

This simulates the closed-loop response for a duration of 26 control intervals with a setpoint of 2 for the first output (the reactor temperature) and 0 for the second output (the residual concentration). Recall that the second output's tuning weight is zero (see the discussion in "Output Weights" on page 3-26), so its setpoint is ignored.

By default, the same linear model is used for controller predictions and the plant, i.e., there is no plant/model mismatch. You can alter this as shown in "Simulation Options" on page 4-7.

When you use the above syntax (no output variables), `sim` automatically plots the plant inputs and outputs (not shown, but see Improved Setpoint Tracking for CSTR Temperature on page 3-26 and Plant Inputs for Modified Rate Weight on page 3-28 for results of a similar scenario).

Simulation Options

You can modify simulation options using the `mpcsimopt` function. For example, the code

```
MPCOpts = mpcsimopt;  
MPCOpts.Constraints = 'off';  
sim(MPCobj, T, r, MPCOpts)
```

runs an unconstrained simulation. Comparing to the case described earlier, the controller's first move now exceeds 4 units (the specified rate constraint).

Other options include the addition of a specified noise sequence to the manipulated variables or measured outputs, open-loop simulations, a look-ahead option for better setpoint tracking or measured disturbance rejection, and plant/model mismatch.

For example, the following code defines a new plant model having gains 50% larger than those in the CSTR model used in the controller, then repeats the above simulation:

```
Plant = 1.5*CSTR;  
MPCOpts.Model = Plant;  
sim(MPCobj, T, r, MPCOpts)
```

In this case, the plant/model mismatch degrades controller performance, but only slightly. Degradation can be severe and must be tested on a case-by-case basis.

Save Calculated Results

If you'd like to save simulation results in your workspace, use the following `sim` function format:

```
[y, t, u] = sim(MPCobj, T, r);
```

This suppresses automatic plotting, instead creating variables `y`, `t`, and `u`, which hold the computed outputs, time, and inputs, respectively. A typical use is to create customized plots. For example, to plot both outputs on the same axis versus time, you could type:

```
plot(t, y)
```


Simulate Controller with Nonlinear Plant

You can use `sim` to simulate a closed-loop system consisting of a linear plant model and an MPC controller.

If your plant is a nonlinear Simulink model, you must linearize the plant (see “Linearization Using Linear Analysis Tool in Simulink Control Design” on page 2-20) and design a controller for the linear model (see “Nonlinear Plants” on page 3-60). To simulate the system, specify the controller in the MPC block parameter **MPC Controller** field and run the closed-loop Simulink model.

Alternatively, your nonlinear model might be a MEX-file, or you might want to include features unavailable in the MPC block, such as a custom state estimator. The `mpcmove` function is the Model Predictive Control Toolbox computational engine, and you can use it in such cases. The disadvantage is that you must duplicate the infrastructure that the `sim` function and the MPC block provide automatically.

The rest of this section covers the following topics:

- “Nonlinear CSTR Application” on page 4-9
- “Example Code for Successive Linearization” on page 4-10
- “CSTR Results and Discussion” on page 4-12

Nonlinear CSTR Application

The CSTR model described in “Linearize Simulink Models” on page 2-17 is a strongly nonlinear system. As shown in “Nonlinear Plants” on page 3-60, a controller can regulate this plant, but degrades (and might even become unstable) if the operating point changes significantly.

The objective of this example is to redefine the predictive controller at the beginning of each control interval so that its predictive model, though linear, represents the latest plant conditions as accurately as possible. This will be done by linearizing the nonlinear model repeatedly, allowing the controller to adapt as plant conditions change. See references [1] and [2] for more details on this approach.

Example Code for Successive Linearization

In the following code, the simulation begins at the CSTR model's nominal operating point (concentration = 8.57) and moves to a low concentration (= 2) where the reaction rate is much higher. The required code is as follows:

```
[sys, xp] = CSTR_INOUT([],[],[], 'sizes');
up = [10 298.15 298.15];
u = up(3);
tsave = []; usave = []; ysave = []; rsave = [];
Ts = 1;
t = 0;
while t < 40
    yp = xp;
    % Linearize the plant model at the current conditions
    [a,b,c,d]=linmod('CSTR_INOUT', xp, up );
    Plant = ss(a,b,c,d);
    Plant.InputGroup.ManipulatedVariables = 3;
    Plant.InputGroup.UnmeasuredDisturbances = [1 2];
    Model.Plant = Plant;

    % Set nominal conditions to the latest values
    Model.Nominal.U = [0 0 u];
    Model.Nominal.X = xp;
    Model.Nominal.Y = yp;

    dt = 0.001;

    simOptions.StartTime = num2str(t);
    simOptions.StopTime = num2str(t+dt);
    simOptions.LoadInitialState = 'on';
    simOptions.InitialState = 'xp';
    simOptions.SaveTime = 'on';
    simOptions.SaveState = 'on';
    simOptions.LoadExternalInput = 'on';
    simOptions.ExternalInput = '[t up; t+dt up]';

    simOut = sim('CSTR_INOUT',simOptions);

    T = simOut.get('tout');
    XP = simOut.get('xout');
```



```
YP = simOut.get('yout');

Model.Nominal.DX = (1/dt)*(XP(end,:) - xp(:));

% Define MPC Toolbox controller for the latest model
MPCobj = mpc(Model, Ts);
MPCobj.W.Output = [0 1];

% Ramp the setpoint
r = max([8.57 - 0.25*t, 2]);

% Compute the control action
if t <= 0
    xd = [0; 0];
    x = mpcstate(MPCobj, xp, xd, [], u);
end

u = mpcmove(MPCobj, x, yp, [0 r], []);

% Simulate the plant for one control interval
up(3) = u;

simOptions.StartTime = num2str(t);
simOptions.StopTime = num2str(t+Ts);
simOptions.InitialState = 'xp';
simOptions.ExternalInput = '[t up; t+Ts up]';

simOut = sim('CSTR_INOUT',simOptions);

T = simOut.get('tout');
XP = simOut.get('xout');
YP = simOut.get('yout');

% Save results for plotting
tsave = [tsave; T];
ysave = [ysave; YP];
usave = [usave; up(ones(length(T),1),:)]';
rsave = [rsave; r(ones(length(T),1),:)]';

xp = XP(end,:)';
```



```
        t = t + Ts;
    end

    figure(1)
    plot(tsave,[ysave(:,2) rsave])
    title('Residual Concentration')
    figure(2)
    plot(tsave,usave(:,3));
    title('Coolant Temperature')
```

CSTR Results and Discussion

The plotted results appear below. Note the following points:

- The setpoint is being ramped from the initial concentration to the desired final value (see the step-wise changes in the reactor concentration plot below). The reactor concentration tracks this ramp smoothly with some delay (see the smooth curve), and settles at the final state with negligible overshoot. The controller works equally well (and achieves the final concentration more rapidly) for a step-wise setpoint change, but it makes unrealistically rapid changes in coolant temperature (not shown).
- The final steady state requires a coolant temperature of 305.20 K (see the coolant temperature plot below). An interesting feature of this nonlinear plant is that if one starts at the initial steady state (coolant temperature = 298.15 K), stepping the coolant temperature to 305.20 and holding will not achieve the desired final concentration of 2. In fact, under this simple strategy the reactor concentration stabilizes at a final value of 7.88, far from the desired value. A successful controller must increase the reactor temperature until the reaction “takes off,” after which it must reduce the coolant temperature to handle the increased heat load. The relinearization approach provides such a controller (see following plots).



- Function `linearize` relinearizes the plant as its state evolves. This function was discussed previously in “Linearization Using MATLAB Code” on page 2-17.
- The code also resets the linear model’s nominal conditions to the latest values. Note, however, that the first two input signals, which are

unmeasured disturbances in the controller design, always have nominal zero values. As they are unmeasured, the controller cannot be informed of the true values. A non-zero values would cause an error.

- Function `mpc` defines a new controller based on the relinearized plant model. The output weight tuning ignores the temperature measurement, focusing only on the concentration.
- At $t = 0$, the `mpcstate` function initializes the controller's extended state vector, x , which is an *mpcstate object*. Thereafter, the `mpcmove` function updates it automatically using the controller's default state estimator. It would also be possible to use an Extended Kalman Filter (EKF) as described in [1] and [2], in which case the EKF would reset the `mpcstate` input variables at each step.
- The `mpcmove` function uses the latest controller definition and state, the measured plant outputs, and the setpoints to calculate the new coolant temperature at each step.
- The Simulink `sim` function simulates the nonlinear plant from the beginning to the end of the control interval. Note that the final condition from the previous step is being used as the initial plant state, and that the plant inputs are being held constant during each interval.

Remember that a conventional feedback controller or a fixed Model Predictive Control Toolbox controller tuned to operate at the initial condition would become unstable as the plant moves to the final condition. Periodic model updating overcomes this problem automatically and provides excellent control under all conditions.

Control Based On Multiple Plant Models

The “Nonlinear CSTR Application” on page 4-9 shows how updates to the prediction model can improve MPC performance. In that case the model is nonlinear and you can obtain frequent updates by linearization.

A more common situation is that you have several linear plant models, each of which applies at a particular operating condition and you can design a controller based on each linear model. If the models cover the entire operating region and you can define a criterion by which you switch from one to another as operating conditions change, the controller set should be able to provide better performance than any individual controller.

The Model Predictive Control Toolbox includes a Simulink block that performs this function. It is the *Multiple MPC Controllers* block. The rest of this section is an illustrative example organized as follows:

- “A Two-Model Plant” on page 4-15
- “Designing the Two Controllers” on page 4-17
- “Simulating Controller Performance” on page 4-18

A Two-Model Plant

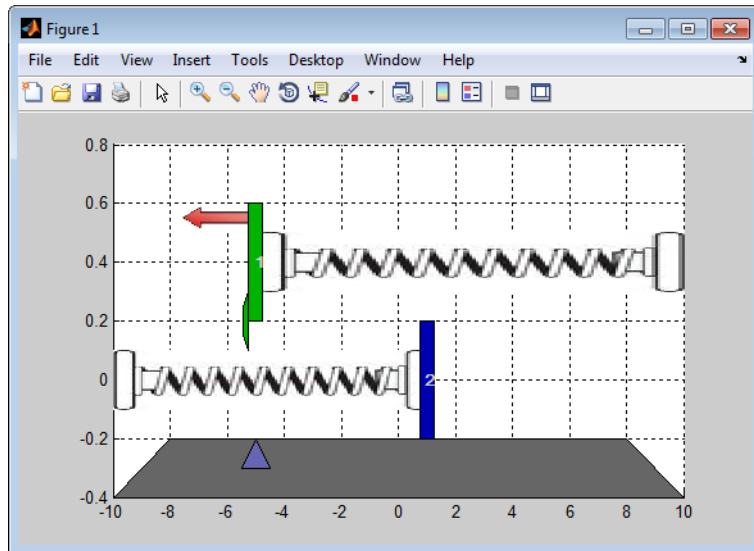
Note Switching MPC Controllers with Multiple MPC Controllers Block provides an animated version of the plant described below.

Animation of the Multi-Model Example on page 4-16 is a stop-action snapshot of the subject plant. It consists of two masses, M_1 and M_2 . A spring connects M_1 to a rigid wall and pulls it to the right. An applied force, shown as a red arrow in Animation of the Multi-Model Example on page 4-16, opposes this spring, pulling M_1 to the left.

When the two masses are detached, as in Animation of the Multi-Model Example on page 4-16, mass M_2 is uncontrollable and responds only to the spring pulling it to the left.

If the two masses collide, however, they stick together (the collision is completely inelastic) until a change in the applied force separates them.

The control objective is to move M_1 in response to a command signal. The blue triangle in Animation of the Multi-Model Example on page 4-16 represents the desired location. At the instant shown, the desired location is -5 .



Animation of the Multi-Model Example

In order to achieve its objective, the controller can adjust the applied force magnitude (the length of the red arrow). It receives continuous feedback on the M_1 location. There is also a contact sensor to signal collisions. The M_2 location is unmeasured.

If M_1 were isolated, this would be a routine control problem. The challenge is that the relationship between the applied force and the M_1 movement changes dramatically when M_2 attaches to M_1 .

The following code defines the model. First define the system parameters as follows:

```
%% Model Parameters
M1=1;      % mass
```

```

M2=5;      % mass
k1=1;      % spring constant
k2=0.1;    % spring constant
b1=0.3;    % friction coefficient
b2=0.8;    % friction coefficient
yeq1=10;   % wall mount position
yeq2=-10;  % wall mount position

```

Next define a model of M_1 when the masses are separated. Its states are the M_1 position and velocity. Its inputs are the applied force, which will be the controller's manipulated variable, and a spring constant calibration signal, which is a measured disturbance input.

```

A1=[0 1; -k1/M1 -b1/M1];
B1=[0 0; -1/M1 k1*yeq1/M1];
C1=[1 0];
D1=[0 0];
sys1=ss(A1,B1,C1,D1);
sys1=setmpcsignals(sys1, 'MV', 1, 'MD', 2);

```

The `setmpcsignals` command specifies the input type for the two inputs.

We need another model (with the same input/output structure) to predict movement when the two masses are joined, as follows:

```

A2=[0 1; -(k1+k2)/(M1+M2) -(b1+b2)/(M1+M2)];
B2=[0 0; -1/(M1+M2) (k1*yeq1+k2*yeq2)/(M1+M2)];
C2=[1 0];
D2=[0 0];
sys2=ss(A2,B2,C2,D2);
sys2=setmpcsignals(sys2, 'MV', 1, 'MD', 2);

```

Designing the Two Controllers

Next we define controllers for each case. Both use a 0.2 second sampling period, a prediction horizon of $P = 20$, a control horizon of $M = 1$, and the default values for all other controller design parameters. The only difference in the controllers is the prediction model.

```

Ts=0.2;    % sampling time
p=20;      % prediction horizon

```

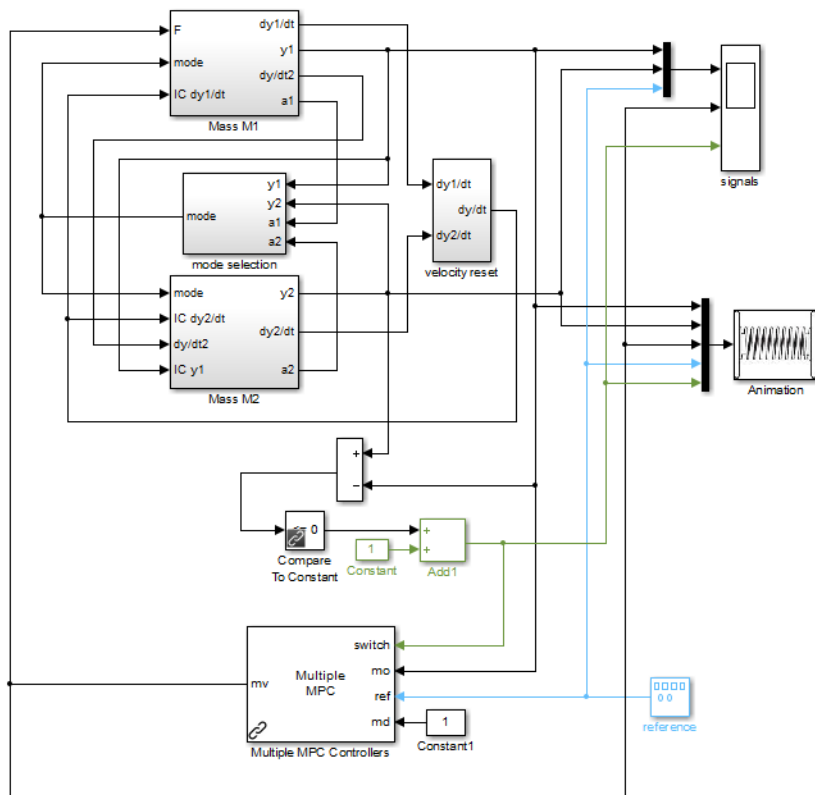
```
m=1;          % control horizon
MPC1=mpc(sys1,Ts,p,m); % Controller for M1 detached from M2
MPC2=mpc(sys2,Ts,p,m); % Controller for M1 connected to M2
```

The applied force also has the same constraints in each case. Its lower bound is zero (it can't reverse direction), and its maximum rate of change is 1000 per second (increasing or decreasing).

```
MPC1.MV=struct('Min',0,'RateMin',-1e3,'RateMax',1e3);
MPC2.MV=struct('Min',0,'RateMin',-1e3,'RateMax',1e3);
```

Simulating Controller Performance

Block Diagram of the Two-Model Example on page 4-19 shows the Simulink block diagram for this example. The upper portion simulates the movement of the two masses, plots the signals as a function of time, and animates the example.




Block Diagram of the Two-Model Example

The lower part contains three key elements:

- A pulse generator that supplies the desired M_1 position (the controller reference signal). Its output is a square wave varying between -5 and 5 with a frequency of 0.015 per second.
- A simulation of a contact sensor. When the two masses have the same position, the Compare to Constant block evaluates to true, and the Add1 block converts this to a 2. Otherwise, the Add1 output is 1.
- The Multiple MPC Controller block. It has four inputs. The measured output (mo), reference (ref), and measured disturbance (md) inputs are as

for a standard MPC Controller block. The distinctive feature is the switch input.

The figure below shows the Multiple MPC Controller block mask for this example (obtained by double clicking on the controller block).




When the switch input is 1 the block automatically activates the first controller listed (MPC1), which is appropriate when the masses are separated. When the switch input is 2 the block automatically enables the second controller (MPC2).

The following code simulates the controller performance

```
Tstop=100; % Simulation time
y1initial=0; % Initial M1 and M2 Positions
y2initial=10;
open('mpc_switching');
sim('mpc_switching',Tstop);
```

The figure below shows the signals scope output.




In the upper plot, the cyan curve is the desired position. It starts at -5 . The M_1 position (yellow) starts at 0 and under the control of MPC1, M_1 moves rapidly toward the desired position. M_2 (magenta) starts at 10 and begins moving in the same direction. At about $t = 13$ seconds, M_2 collides with M_1 . The switching signal (lower plot) changes at this instant from 1 to 2, so controller MPC2 has taken over.

The collision moves M_1 away from its desired position and M_2 remains joined to M_1 . Controller MPC2 adjusts the applied force (middle plot) so M_1 quickly returns to the desired position.

When the desired position changes step-wise to 5, the two masses separate briefly (with appropriate switching to MPC1) but for the most part move together and settle rapidly at the desired position. The transition back to -5 is equally well behaved.

Now suppose we force MPC2 to operate under all conditions. The figure below shows the result. When the masses are separated, as at the start, MPC2 applies excessive force and then over-compensates, resulting in oscillatory behavior. Once the masses join, the movement smooths out, as would be expected.

The oscillations are especially severe in the last transition. The masses collide frequently and M_1 never reaches the desired position.



If we put MPC1 in charge exclusively, we instead see sluggish movements that fail to settle at the desired position before the next transition occurs (not shown but you can run `mpcswitching`).

In this case, at least, two controllers are better than one.

Compute Steady-State Gain

This example shows how to analyze a Model Predictive Controller using `cloffset`. The function computes the closed-loop, steady-state gain for each output when subjected to a sustained, 1-unit disturbance added to each output. It assumes that no constraints will be encountered.

For example, consider the controller operating at the final steady-state of the nonlinear CSTR of the previous section. To compute its gain, type

```
cloffset(MPCobj)
```

which gives the result:

```
ans =
```

```
    1.0000    14.5910  
   -0.0000    -0.0000
```

The interpretation is that the controller doesn't react to a sustained disturbance of 1 unit in the first output (the reactor temperature). Recall that we assigned zero weight to this output in the controller design, so the controller ignores deviations from its setpoint. The same disturbance has no effect on the second output (the 2,1 element is zero).

If there is a 1-unit disturbance in the second output, the controller reacts, and the first output increases 14.59 units. This is again due to the zero weight on this output. The second output stays at its setpoint (the 2,2 element is zero).

Extract Controller

This example shows how to obtain an LTI representation of an unconstrained Model Predictive Control Toolbox controller using `ss`. You can use this to analyze the controller's closed-loop frequency response, etc.

For example, consider the controller designed in “Create a Controller Object” on page 4-2. To extract the controller, you could type:

```
MPCss = ss(MPCobj);
```

You could then construct an LTI model of the closed-loop system using the feedback function (see Control System Toolbox for details) by typing:

```
CSTRd = c2d(CSTR, MPCss.Ts);  
Feedin = 1;  
Feedout = 1;  
Sign = 1;  
CLsys = feedback(CSTRd, MPCss, Feedin, Feedout, Sign);
```

Note The CSTR model must be converted to discrete form with the same control interval as the controller.

Recall that the CSTR plant has two inputs and two outputs. The first input is the manipulated variable and the other is an unmeasured disturbance. The first output is measured for feedback and the other is not. The `Feedin` and `Feedout` parameters specify the input and output to be used for control. The `Sign` parameter signifies that the MPC object uses positive feedback, i.e., the measured outputs enter the controller with no sign change. Omission of this would cause the `feedback` command to use negative feedback by default and would almost certainly lead to an unstable closed-loop system.

You could then type

```
eig(CLsys)
```

to verify that all closed-loop poles are within the unit circle, or

```
bode(CLsys)
```

to compute a closed-loop Bode plot.

Bibliography

[1] Lee, J. H. and N. L. Ricker, "Extended Kalman Filter Based Nonlinear Model Predictive Control," *Ind. Eng. Chem. Res.*, Vol. 33, No. 6, pp. 1530–1541 (1994).

[2] Ricker, N. L., and J. H. Lee "Nonlinear Model Predictive Control of the Tennessee Eastman Challenge Process," *Computers & Chemical Engineering*, Vol. 19, No. 9, pp. 961–981 (1995).

Designing and Testing Controllers in Simulink

- “Design Controller in Simulink” on page 5-2
- “Test an Existing Controller” on page 5-17
- “Schedule Controllers at Multiple Operating Points” on page 5-18

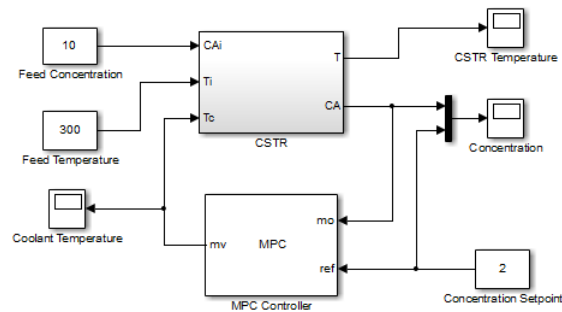
The Model Predictive Control Toolbox provides a controller block library for use in Simulink. This allows you to test controller performance in nonlinear simulations. The procedure depends on whether you have already designed a controller. You also have the option to coordinate multiple Model Predictive Controllers.

Design Controller in Simulink

This example shows how to design a Model Predictive Controller in Simulink.

1 Open the Simulink model.

```
sys = 'CSTR_ClosedLoop';
open_system(sys);
```



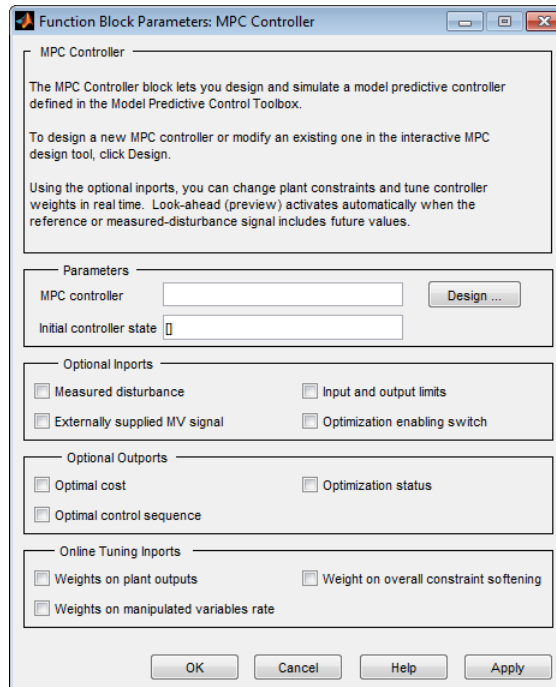
The controller adjusts a coolant temperature, which is one of three plant inputs. The other inputs are the concentration of the limiting reactant in the CSTR feed stream and the temperature of this stream. The CSTR states are the temperature and concentration of the limiting reactant in the product stream.

The controller is connected such that the residual concentration is the measured and controlled output variable (MO) and the coolant temperature is the manipulated variable (MV). The feed concentration and temperature are unmeasured disturbances and the reactor temperature is ignored. In practice, the reactor temperature is usually measured and controlled, but the configuration here presents a more challenging nonlinear control problem.

Tip To add a MPC Controller block to your model, open the Simulink Library Browser and drag and drop the block from the **Model Predictive Control Toolbox** library.

The desired CSTR concentration is 2 kmol/m^3 and the above model includes this constant reference signal. It also includes scope plots of the MV, OV (with reference), and CSTR temperature.

- 2 Double-click the MPC Controller block to view its options.

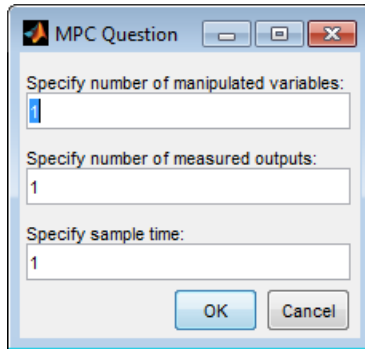


Measured disturbance is unchecked and the **MPC Controller** box is blank, which indicates that no controller has been designed yet.

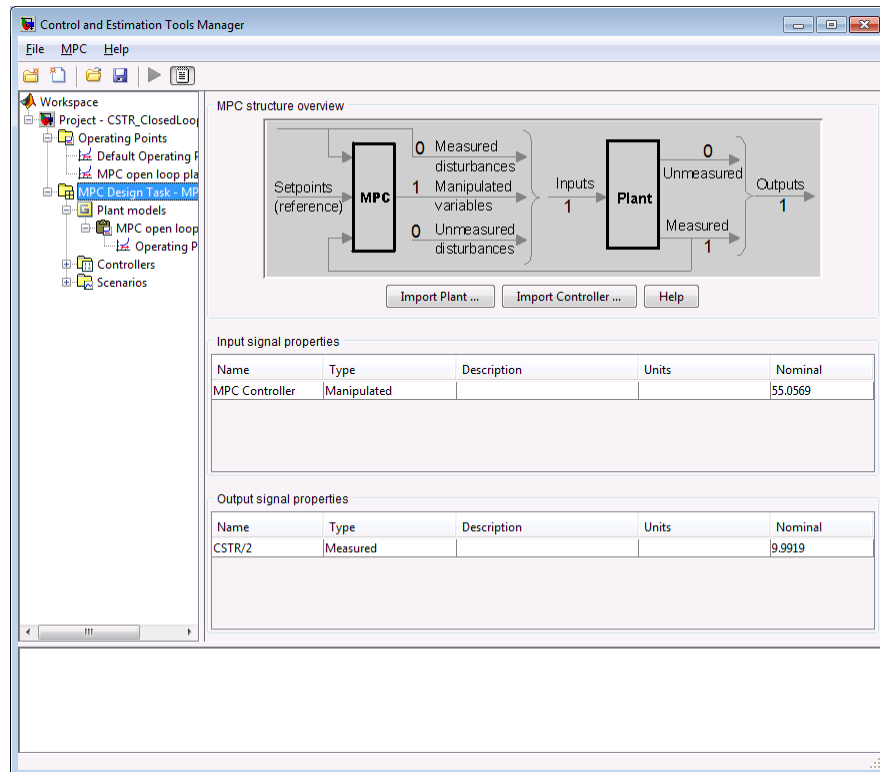
- 3 Click **Design** in the block parameters dialog box.

This action opens the MPC Question dialog box that asks you for information about the plant.

In this case, there is one MV and one OV. The controller has a sample time of 1 minute.



Click **OK**. This action linearizes the Simulink plant model using the signals connected to the controller. The Control and Estimation Tools Manager window also opens, which contains an **MPC Design Task** node with the linearized plant model loaded in it.



The nominal values for the **Input signal properties** and **Output signal properties** are shown. In the **MPC Design Task** node, select **MPC open loop plant 1** to see the linearized prediction model.

4 Evaluate the linearized plant model.

The automated linearization procedure uses the default equilibrium operating point that is likely different from the desired operating point. It is important to verify whether the operating point is appropriate. Else, the linear model may be unsuitable for model predictive control at the target conditions.


In the **MPC Design Task**, click **Operating Points** to see the default operating point. The linearization procedure is able to find a steady-state operating point (derivatives are close to zero), but it is far from the target

condition. The reactor temperature is very low (244 K) and consequently little reaction occurs — the concentration with the reactor is essentially equal to the feed concentration. Since the model has strong nonlinearities, the resulting linear model is probably a poor design basis.

5 Re-linearize the plant.

a Add output constraints.

In the Simulink model, right-click the MO signal entering the controller. Select **Linear Analysis Points**> **Trim Output Constraint**. Repeat for the MV signal of the MPC Controller block. All constrained signals are marked with **T**.




You must specify output constraints on all the OV/MV/MD signals connected to the MPC Controller because their steady-state values are used to compute the nominal values used by the model predictive controller.

b Select the **Compute Operating Points** tab in the **Operating Points** node of the Control and Estimation Tools Manager.

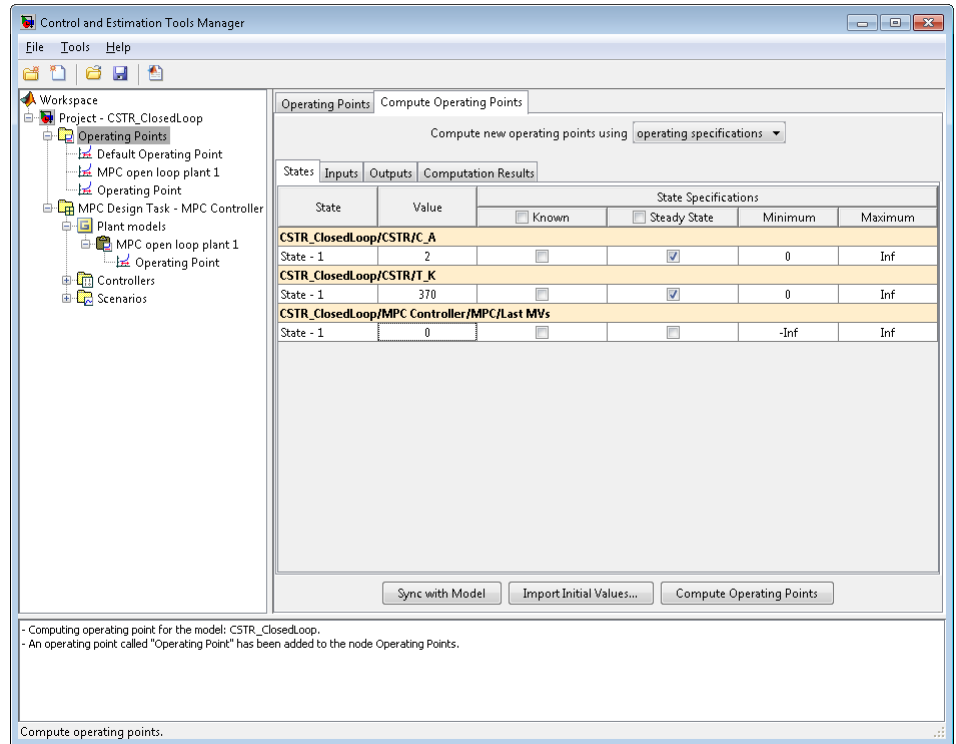
By default, the software computes a steady-state operating point.

c Specify known outputs.

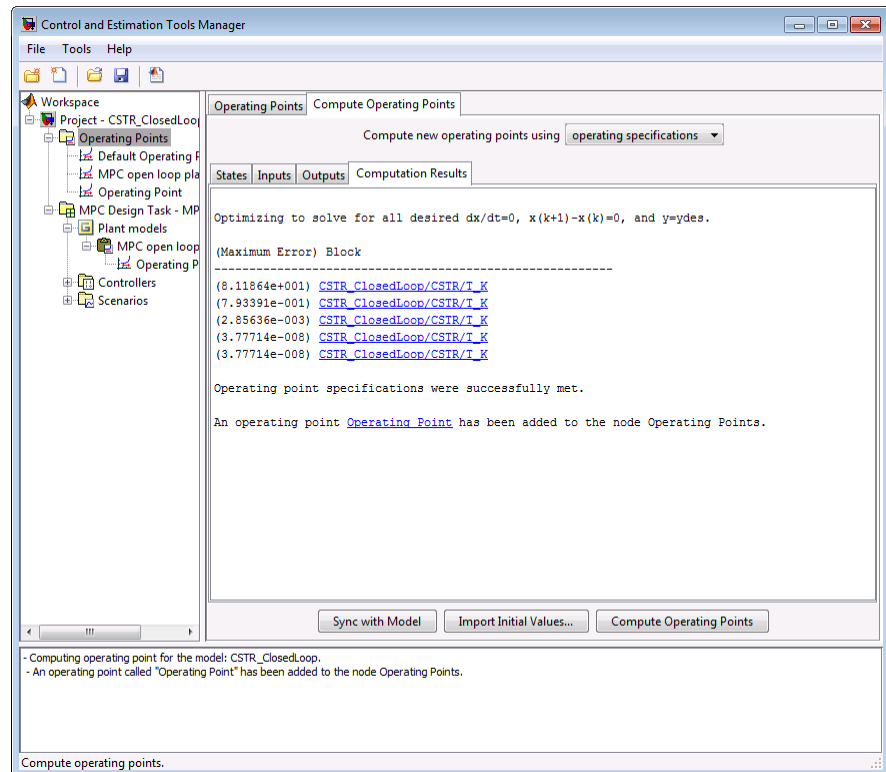
The CSTR concentration must be 2 kmol/m³. In the **Outputs** tab, select the **Known** check box for it and specify 2 in the **Value** box.



- d Configure initial estimates of the states. Otherwise, the numerical search for the new operating point may fail, depending on the plant's nonlinearity.
- The CSTR temperature estimate is 370 K. Specify 370 in the **Value** box for the CSTR_ClosedLoop/CSTR/T_K block.
 - The CSTR concentration is 2 kmol/m³. Specify 2 in the **Value** box for the CSTR_ClosedLoop/CSTR/C_A block.
 - Clear the **Steady State** check box for the CSTR_ClosedLoop/MPC Controller/MPC/Last MVs block.



e Click Compute Operating Points.




The new operating point, **Operating Point**, appears in the navigation pane.

Tip If an operating point computation fails, you can:


- Specify better initial estimates of the unknown quantities.
- Select a different optimization algorithm for the operating point search. To specify a different optimization algorithm, in the Controls and Estimation Tools Manager, select **Options** in the **Tools** menu and select **Operating Point Search**.

- a Select the new operating point in the navigation pane to evaluate it.




The new operating point is at steady state (all the derivatives are near zero). Also, the actual concentration value is 2 kmol/m^3 , as expected.

- b** Linearize at the new operating point.
 - Select **MPC Design Task** in the navigation pane.
 - Open the Plant Model Importer dialog box.
Click **Import Plant**.
 - Select the **Linearized Plant from Simulink** tab.
 - Select the operating point for the linearization.
In the **Use existing operating point** list, select **Operating Point**



- Click **OK**.
- Click **Close**.
- Confirm the new plant model.

Select **MPC Design Task** in the navigation pane. The linearization procedure has entered the new nominal conditions.



- d** (Optional) Delete the model computed at default operating point.

In the navigation pane, select **Plant models**. In the **Plant models imported for this project** list box, select MPC open loop plant 1. Click **Delete**.

Deleting MPC open loop plant 1 forces the controller to use the new model and eliminates the potential for confusion.

- 6** Export the controller to the Workspace.

By default, the software creates a controller based on the model obtained by automatic linearization. The default controller is named **MPC1**. The Control and Estimation Tools Manager automatically exports it to the Workspace and enters the name MPC1 in the **MPC Controller** field of the MPC Controller block parameters dialog box.

In general, you should use the design tool or commands to improve the design, testing initially with linear simulations. When finished with modifications, export the controller object to the Workspace.

Right-click the controller **MPC1** in the **Controllers** node and select **Export**. You can optionally rename the controller when you export it.

Tip Before you export the controller, you can examine it for design errors and stability problems using **review**.

7 Verify steady state for the plant/controller combination.

The MPC Controller must contain the name of the controller that you want to use from the base workspace. In this example, you also need to export the new operating point and initialize the Simulink model with it.

Right-click **Operating Point** in **mpc open loop plant 2** and select **Export to workspace**.

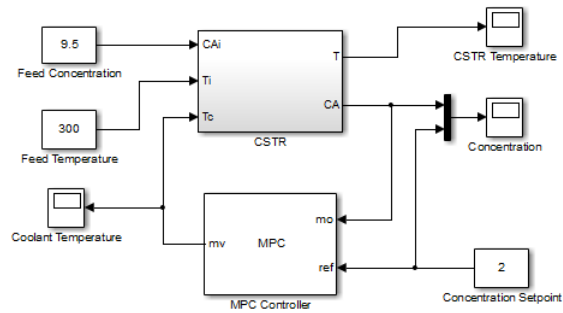
This action opens the Export to workspace dialog box.

Select the **Use the operating point to initialize model** check box.

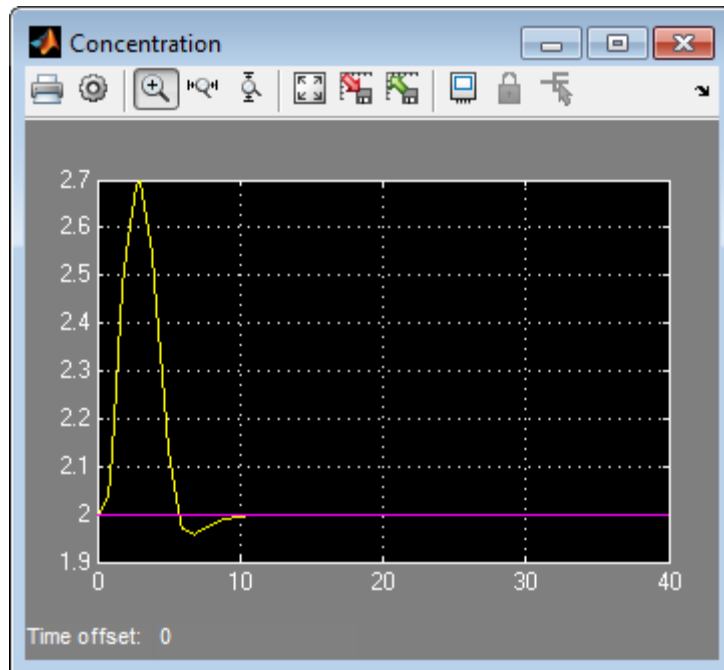
Run the Simulink model, including the installed controller to verify that both have been initialized to the correct nominal condition. If there are no disturbances, all signals would remain essentially constant. Else, make sure that the states within the plant model and the nominal controller output are set to the values calculated in the previous steps.


8 Test controller performance and revise the design if necessary.


For example, the diagram below shows the model configured to test the response to a feed concentration of 9.5 kmol/m^3 . The nominal value is 10 kmol/m^3



Run the Simulink model to obtain the trend plots shown below.







The decrease in feed concentration reduces heat generation. If the controller were absent, the reactor temperature would drop significantly, reducing the reduction rate, and the CSTR concentration would increase to roughly 3 kmol/m^3 . To counteract this, the controller raises the coolant temperature, returning the CSTR temperature to a value slightly below the nominal condition.

Test an Existing Controller

If you have already designed a model predictive controller to use with a Simulink plant, to test it in Simulink:

- 1 Open Simulink model.
- 2 Add an MPC Controller block to the model and connect it.
- 3 Specify the controller.

Open the MPC Controller and enter the name of your controller in the **MPC Controller** box. The string you enter here must be the name of an `mpc` object in the workspace.

- 4 (Optional) Modify the controller.

To use `mpctool`, click **Design** with the MPC Controller open. Change the controller parameters and then export the modified controller to the workspace.

Alternatively, use commands to modify the controller object in the workspace.

- 5 Run the Simulink model.

Schedule Controllers at Multiple Operating Points

In this section...
“A Two-Model Plant” on page 5-18
“Animation of the Multi-Model Example” on page 5-19
“Designing the Two Controllers” on page 5-21
“Simulating Controller Performance” on page 5-21


If your plant is nonlinear, a controller designed to operate in a particular target region may perform poorly in other regions. The common way to compensate is to design multiple controllers, each intended for a particular combination of operating conditions. Switch between them in real-time as conditions change. Gain scheduling is a conventional example of this technique. The following example shows how to coordinate multiple model predictive controllers for this purpose.

The rest of this section is an illustrative example organized as follows:

- “A Two-Model Plant” on page 5-18 describes the example application involving two distinct operating conditions.
- “Animation of the Multi-Model Example” on page 5-19 explains how to simulate the plant.
- “Designing the Two Controllers” on page 5-21 shows how to obtain the controllers for each operating condition.
- “Simulating Controller Performance” on page 5-21 shows how to use the Multiple MPC Controllers block in Simulink to coordinate the two controllers.

A Two-Model Plant

The figure below is a stop-action snapshot of the plant. It consists of two masses, M1 (upper) and M2 (lower). A spring connects M1 to a rigid wall and pulls it to the right. An applied force, shown as a red arrow, opposes the spring pulling M1 to the left.



In the above, mass M2 is uncontrollable. It responds solely to the spring pulling it to the left.

If the two masses collide, however, they stick together until a change in the applied force separates them.

The control objective is to make M1 track a reference signal. The blue triangle in the above indicates the desired location, which varies with time. At the instant shown, the desired numeric location is -5.

For an animated version of the plant, see *Scheduling Controllers for a Plant with Multiple Operating Points*.

Animation of the Multi-Model Example

In order to achieve its objective, the controller can adjust the applied force magnitude (the length of the red arrow). It receives continuous measurements of the M1 location. There is also a contact sensor to signal collisions. The M2 location is unmeasured.

If M1 were isolated, this would be a routine control problem. The challenge is that the relationship between the applied force and M1's movement changes dramatically when M2 attaches to M1.

Define the model.

Define the system parameters.

```
M1 = 1; % mass
M2 = 5; % mass
k1 = 1; % spring constant
k2 = 0.1; % spring constant
b1 = 0.3; % friction coefficient
b2 = 0.8; % friction coefficient
yeq1 = 10; % wall mount position
yeq2 = -10; % wall mount position
```

Define a model of M1 when the masses are separated. Its states are the position and velocity of M1. Its inputs are the applied force, which will be the controller's manipulated variable (MV), and a spring constant calibration signal, which is a measured disturbance (MD) input.

```
A1 = [0 1; -k1/M1 -b1/M1];
B1 = [0 0; -1/M1 k1*yeq1/M1];
C1 = [1 0];
D1 = [0 0];
sys1 = ss(A1,B1,C1,D1);
sys1 = setmpcsignals(sys1, 'MV', 1, 'MD', 2);
```

The call to `setmpcsignals` specifies the input type for the two inputs.

Define another model (with the same input/output structure) to predict movement when the two masses are joined.

```
A2 = [0 1; -(k1+k2)/(M1+M2) -(b1+b2)/(M1+M2)];
B2 = [0 0; -1/(M1+M2) (k1*yeq1+k2*yeq2)/(M1+M2)];
C2 = [1 0];
D2 = [0 0];
sys2 = ss(A2,B2,C2,D2);
sys2 = setmpcsignals(sys2, 'MV', 1, 'MD', 2);
```

Designing the Two Controllers

Next, define controllers for each case. Both controllers employ a 0.2 second sampling period (T_s), a prediction horizon of $p = 20$, a control horizon of $m = 1$ and default values for all other controller parameters.

The only property that distinguishes the two controllers is the prediction model.


```
Ts = 0.2;  
p = 20;  
m = 1;  
MPC1 = mpc(sys1,Ts,p,m); % Controller for M1 detached from M2  
MPC2 = mpc(sys2,Ts,p,m); % Controller for M1 connected to M2
```

The applied force also have the same constraints in each case. The lower bound for the applied force is zero because it cannot reverse direction. The maximum rate of change for the applied force is 1000 per second (increasing or decreasing).

```
MPC1.MV = struct('Min',0,'RateMin',-1e3,'RateMax',1e3);  
MPC2.MV = struct('Min',0,'RateMin',-1e3,'RateMax',1e3);
```

Simulating Controller Performance

The following is the Simulink block diagram for this example.

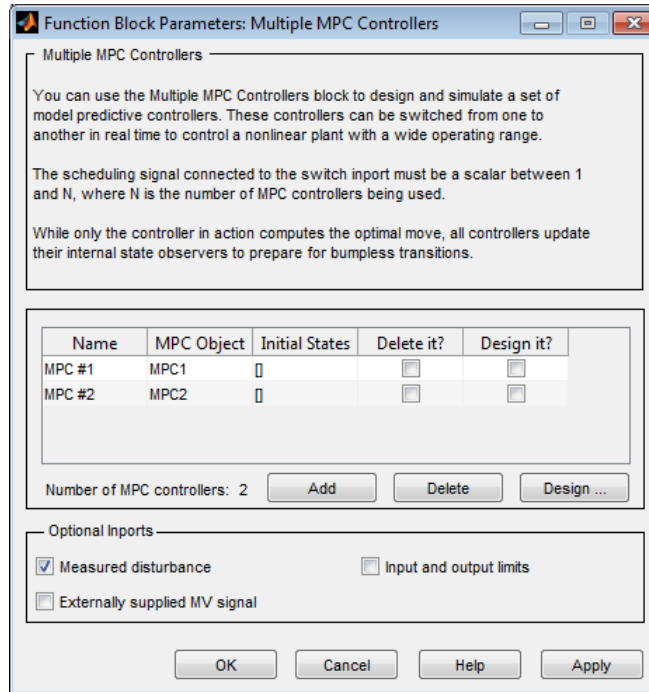


The upper portion simulates the movement of the two masses, plots the signals as a function of time and animates the plant.

The lower part contains the following key elements:

- A pulse generator that supplies the desired M1 position (the controller reference signal). The output of the pulse generator is a square wave that varies between -5 and 5 , with a frequency of 0.015 per second.
- Simulation of a contact sensor. When the two masses have the same position, the Compare to Constant block evaluates to true and the Add1 block converts this to a 2. Else, the output of Add1 is 1.
- The Multiple MPC Controllers block, which has four inputs. The measured output (mo), reference (ref) and measured disturbance (md) inputs are as for the MPC Controller block. The distinctive feature of the Multiple MPC Controllers block is the switch input.

The figure below shows the Multiple MPC Controllers block mask for this example.




When the switch input is 1, the block automatically activates the first controller that is listed (MPC1), which is appropriate when the masses are separated. When the switch input is 2, the block automatically enables the second controller (MPC2).

Simulate the controller performance using Simulink commands.

```
Tstop = 100; % Simulation time
y1initial = 0; % Initial M1 and M2 Positions
y2initial = 10;
open('mpc_switching');
sim('mpc_switching',Tstop);
```

The figure below shows the signals scope output.



In the upper plot, the cyan curve is the desired position. It starts at -5. The M1 position (yellow) starts at 0. Under the control of MPC1, M1 moves

rapidly towards the desired position. M2 (magenta) starts at 10 and begins moving in the same direction.


At about $t = 13$ seconds, M2 collides with M1. The switching signal (lower plot) changes at this instant from 1 to 2, so MPC2 takes over.

The completely inelastic collision moves M1 away from its desired position and M2 remains joined to M1. Controller MPC2 adjusts the applied force (middle plot) so M1 quickly returns to the desired position.

When the desired position changes step-wise to 5, the two masses separate briefly, with the controller switching to MPC1 appropriately. But, for the most part, the two masses move together and settle rapidly at the desired location. The transition back to -5 is equally well behaved.

Suppose we force MPC2 to operate under all conditions.

The figure below shows the result.



When the masses are disconnected, as at the start, MPC2 applies excessive force and then over-compensates, resulting in oscillatory behavior. Once the masses join, the move more smoothly, as would be expected.

The last transition causes especially severe oscillations. The masses collide frequently and M1 never reaches the desired position.

If MPC1 is in charge exclusively, the masses move sluggishly and fail to settle at the desired position before the next transition occurs. For more information, see the Scheduling Controllers for a Plant with Multiple Operating Points.

In this application, at least, two controllers are, indeed, better than one.

B

blocking 3-30

C

`cloffset` function 4-24

closed-loop

 design tool simulations 3-15

constraints

 input 3-33

 manipulated variable 3-33

 output 3-42

 softening 3-42

Control and Estimation Tools Manager 3-2

control horizon

 control specification 3-24

control interval

 control specification 3-24

controller settings

 delay compensation 3-54

 disturbance estimation 3-35

 get function 4-3

 input constraints 3-33

 input weights 3-29

 move suppression 3-27

 output weights 3-26

 prediction horizon 3-24

 rate weights 3-27

 set function 4-3

 weights 3-25

controllers

 comparisons 3-30

 creation 4-2

 design tool 3-23

 design tool list 3-7

 exporting 3-46

 LTI object extraction 4-25

 mpc function 4-2

 property settings 4-3

 robustness 3-50

CSTR (Continuous Stirred-Tank Reactor)

 linear model 2-5

D

data markers 3-19

delays

 controller settings for 3-54

description

 signal specification 3-5

design tool 3-6

 loading models 3-3

 navigation in tree view 3-6

 starting 3-2

See also views

direct feedthrough 2-7

disturbance estimation 3-35

disturbances

 specification in simulations 3-15

documentation 1-3

E

estimation

 disturbances 3-35

G

gain

 closed-loop 4-24

H

horizons

 specification 3-24

I

import

 plant model 3-4

importing

- controller 3-48
- InputGroup property 2-7
- InputName property 2-9
- inputs
 - manipulated variables 2-3
 - measured disturbances 2-3
 - names 2-9
 - nominal value 3-5
 - type specification 2-9
 - types 2-3
 - unmeasured disturbances 2-3

L

- linearization
 - repeated 4-10
- loading
 - controllers 3-48
 - project 3-48
- LTI objects
 - properties 2-7

M

- manipulated variables
 - definition 2-3
- measured disturbances
 - definition 2-3
- measured outputs
 - definition 2-3
- Model Predictive Control (MPC)
 - related products 1-4
- models
 - characteristics 2-11
 - delays 3-53
 - design tool list 3-6
 - disturbances 3-35
 - loading into design tool 3-3
 - LTI 2-4
 - LTI chemical reactor 2-5

- nonlinear 2-17
- nonlinear CSTR example 3-60
- nonlinear simulation 4-9
- nonsquare 3-58
- perturbation 3-50
- plant 2-2
- prediction error 3-50
- relinearizing 3-65
- state-space 2-5
- transfer function 2-4
- zero/pole/gain 2-5

mpc objects

- conversion to state-space 4-25
- creating 4-2

- mpcmove function 4-10

- mpcstate objects 4-12

- MV (manipulated variables). *See* manipulated variables

N

- names
 - specification using design tool 3-5
- nonlinear plants 3-60
- nonsquare plants 3-58

O

objects

- LTI 2-7

- mpc 4-2

- mpc conversion to ss 4-25

- mpcstate 4-12

open-loop

- design tool simulations 3-15

- OutputGroup property 2-7

- OutputName property 2-9

outputs

- measured 2-3

- names 2-9

- nominal value 3-5
- setpoint definition 3-11
- type specification 2-9
- types 2-3
- unmeasured 2-3

P

- plant/model mismatch 3-50
- plants
 - input/output signals 2-3
 - MIMO plant 2-12
 - nonlinear 3-60
 - nonlinear simulation 4-9
 - nonsquare plant 3-58
- plots
 - data markers 3-19
 - normalizing 3-22
 - simulation responses 3-18
 - variable selection 3-22
- properties
 - mpc object 4-3

R

- response plots 3-18 to 3-19
 - data markers 3-12
 - See also* plots
- robustness
 - controller prediction error 3-50
 - testing 3-50

S

- saving
 - controllers 3-46
 - projects 3-47
- scenarios
 - comparing 3-20
 - design tool list 3-9
 - specification 3-11

- setpoints
 - specification 3-11
- signal labels
 - description 3-5
 - name 3-5
 - units 3-5
- signals
 - inputs 2-3
 - outputs 2-3
- sim function
 - options 4-7
 - simulations 4-6
- simulations 3-9
 - closed-loop 3-15
 - conditions 3-11
 - linear 3-10
 - nonlinear 3-68
 - open loop 3-15
 - options 4-7
 - response plots 3-18
 - sim function 4-6
 - starting in design tool 3-12
 - toolbar icon 3-12
 - See also* scenarios
- soft constraints 3-42
- ss function 2-6 2-15

T

- tf function 2-4
- transfer functions
 - specification 2-4
- tree
 - navigation 3-6
- types
 - specification using design tool 3-5

U

- unmeasured outputs

definition 2-3

V

views 3-5

constraint softening 3-44

controller settings 3-23

controllers 3-7

plant models 3-6

scenarios 3-9

signal properties 3-6

See also design tool

W

weights

move suppression 3-27

output 3-26

rate 3-27

setpoint tracking 3-26

Z

zpk function 2-5